

VMware vSphere 8.0

Table of Contents

Release Notes.....	74
VMware vSphere 8.0 Release Notes.....	74
VMware vCenter Server Photon OS Security Patches.....	89
VMware Host Client 2.5.0 Release Notes.....	102
ESXi Update and Patch Release Notes.....	104
VMware ESXi 8.0 Update 3g Release Notes.....	104
VMware ESXi 8.0 Update 3f Release Notes.....	164
VMware ESXi 8.0 Update 3e Release Notes.....	177
VMware ESXi 8.0 Update 3d Release Notes.....	257
VMware ESXi 8.0 Update 3c Release Notes.....	269
VMware ESXi 8.0 Update 3b Release Notes.....	284
VMware ESXi 8.0 Update 3 Release Notes.....	333
VMware ESXi 8.0 Update 2e Release Notes.....	351
VMware ESXi 8.0 Update 2d Release Notes.....	365
VMware ESXi 8.0 Update 2c Release Notes.....	378
VMware ESXi 8.0 Update 2b Release Notes.....	392
VMware ESXi 8.0 Update 2 Release Notes.....	440
VMware ESXi 8.0 Update 1d Release Notes.....	458
VMware ESXi 8.0 Update 1c Release Notes.....	471
VMware ESXi 8.0 Update 1a Release Notes.....	507
VMware ESXi 8.0 Update 1 Release Notes.....	522
VMware ESXi 8.0f Release Notes.....	535
VMware ESXi 8.0e Release Notes.....	543
VMware ESXi 8.0d Release Notes.....	551
VMware ESXi 8.0c Release Notes.....	560
VMware ESXi 8.0b Release Notes.....	573
VMware ESXi 8.0a Release Notes.....	616
vCenter Server Update and Patch Release Notes.....	629
VMware vCenter 8.0 Update 3g Release Notes.....	629
VMware vCenter Server 8.0 Update 3e Release Notes.....	647
VMware vCenter Server 8.0 Update 3d Release Notes.....	668
VMware vCenter Server 8.0 Update 3c Release Notes.....	682
VMware vCenter Server 8.0 Update 3b Release Notes.....	696
VMware vCenter Server 8.0 Update 3a Release Notes.....	712
VMware vCenter Server 8.0 Update 3 Release Notes.....	728
VMware vCenter Server 8.0 Update 2e Release Notes.....	751

VMware vCenter Server 8.0 Update 2d Release Notes.....	753
VMware vCenter Server 8.0 Update 2c Release Notes.....	766
VMware vCenter Server 8.0 Update 2b Release Notes.....	778
VMware vCenter Server 8.0 Update 2a Release Notes.....	797
VMware vCenter Server 8.0 Update 2 Release Notes.....	809
VMware vCenter Server 8.0 Update 1e Release Notes.....	824
VMware vCenter Server 8.0 Update 1d Release Notes.....	834
VMware vCenter Server 8.0 Update 1c Release Notes.....	844
VMware vCenter Server 8.0 Update 1b Release Notes.....	857
VMware vCenter Server 8.0 Update 1a Release Notes.....	868
VMware vCenter Server 8.0 Update 1 Release Notes.....	877
VMware vCenter Server 8.0c Release Notes.....	892
VMware vCenter Server 8.0b Release Notes.....	900
VMware vCenter Server 8.0a Release Notes.....	912
VMware Host Client Release Notes.....	922
VMware Host Client 2.18.0 Release Notes.....	922
VMware Host Client 2.14.0 Release Notes.....	925
VMware Host Client 2.12.0 Release Notes.....	927
ESXi Installation and Setup.....	930
About VMware ESXi Installation and Setup.....	930
Introduction to vSphere Installation and Setup.....	930
How to Install and Set Up vSphere.....	930
About ESXi Evaluation and Licensed Modes.....	931
What is VMware vSphere Distributed Services Engine®	932
High Availability with VMware vSphere Distributed Services Engine.....	934
Error Handling, Failover, and Rollback for VMware vSphere Distributed Services Engine.....	934
Installing and Setting Up ESXi.....	935
ESXi Requirements.....	935
ESXi System Storage Overview.....	936
ESXi Hardware Requirements.....	940
Using Remote Management Applications.....	942
Recommendations for Enhanced ESXi Performance.....	943
Incoming and Outgoing Firewall Ports for ESXi Hosts.....	944
Required Free Space for System Logging.....	944
VMware Host Client System Requirements.....	945
ESXi Passwords and Account Lockout.....	945
Preparing for Installing ESXi.....	947
Download the ESXi Installer.....	947
Required Information for ESXi Installation.....	947
Media Options for Booting the ESXi Installer.....	948

Customizing Installations with vSphere ESXi Image Builder.....	951
How the vSphere ESXi Image Builder Works.....	951
Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects.....	956
Configure vSphere ESXi Image Builder.....	960
Using VMware.Image Builder Cmdlets.....	960
ESXi Image Profile Tasks.....	962
vSphere ESXi Image Builder Workflows with PowerCLI Cmdlets.....	976
Installing ESX.....	981
Installing ESXi Interactively.....	981
Installing ESXi by Using a Script.....	984
How to Boot an ESXi Host from a Network Device.....	1001
Installing ESXi Using vSphere Auto Deploy.....	1013
Troubleshooting vSphere Auto Deploy.....	1081
Setting Up ESX.....	1086
Initial ESXi Configuration.....	1086
Enable ESXi Shell and SSH Access with the Direct Console User Interface.....	1089
Set the Password for the Administrator Account.....	1090
Configuring the BIOS Boot Settings.....	1090
Configuring Network Settings.....	1091
Test the Management Network.....	1095
Restart the Management Agents.....	1096
Restart the Management Network.....	1096
Test Connectivity to Devices and Networks.....	1096
Restoring the Standard Switch.....	1096
Configuring System Logging.....	1097
Set the Host Image Profile Acceptance Level.....	1115
Remove All Custom Packages on ESXi.....	1115
Modify ESXi Configuration Files.....	1116
Deactivate Support for Non-ASCII Characters on ESXi.....	1116
Reset the System Configuration.....	1116
After You Install and Set Up ESXi.....	1117
Licensing ESXi Hosts.....	1117
Recording the License Key of an ESXi Host.....	1117
View the License Keys of ESXi Hosts from the vSphere Client.....	1118
Access the ESXi License Key from the Direct Console.....	1118
View System Logs.....	1118
Troubleshooting ESXi Booting.....	1118
Host Stops Unexpectedly at Bootup When Sharing a Boot Disk with Another Host.....	1118
Host Fails to Boot After You Install ESXi in UEFI Mode.....	1119
Decommission an ESXi Host.....	1120

VMware ESXi Upgrade.....	1121
vSphere Upgrade Options.....	1121
Overview of the vSphere Upgrade Process.....	1121
Upgrading ESXi Hosts.....	1126
ESXi Requirements.....	1126
ESXi System Storage Overview.....	1126
ESXi Hardware Requirements.....	1131
Using Remote Management Applications.....	1133
Recommendations for Enhanced ESXi Performance.....	1134
Incoming and Outgoing Firewall Ports for ESXi Hosts.....	1135
Required Free Space for System Logging.....	1135
VMware Host Client System Requirements.....	1136
ESXi Passwords and Account Lockout.....	1136
Before Upgrading ESXi Hosts.....	1138
Upgrading Hosts That Have Third-Party Custom VIBs.....	1139
Upgrading ESXi Hosts in an Environment With VMware NSX.....	1139
Media Options for Booting the ESXi Installer.....	1139
Download the ESXi Installer.....	1142
ESXi Storage Device Names and Identifiers.....	1142
Upgrade Hosts Interactively.....	1144
Installing or Upgrading Hosts by Using a Script.....	1145
Enter Boot Options to Run an Installation or Upgrade Script.....	1145
Installing ESXi by Using a Script.....	1147
Install or Upgrade ESXi from a CD or DVD by Using a Script.....	1156
Install or Upgrade ESXi from a USB Stick by Using a Script.....	1157
Network Boot the ESXi Installer for a Scripted Installation or Upgrade.....	1157
Disk Device Names.....	1157
How to Boot an ESXi Host from a Network Device.....	1158
Network Booting the ESXi Installer.....	1158
Boot the ESXi Installer by Using PXE and TFTP.....	1161
Boot the ESXi Installer by Using iPXE and HTTP.....	1163
Boot the ESXi Installer by Using Native UEFI HTTP.....	1165
Sample DHCP Configurations.....	1166
How to Upgrade Hosts by Using ESXCLI Commands.....	1169
Upgrading Hosts by Using ESXCLI Commands.....	1170
See if ESXi Host Update Requires Maintenance Mode or Reboot.....	1172
Place a Host in Maintenance Mode.....	1173
Update a Host with Individual VIBs.....	1174
Upgrade or Update a Host with Image Profiles.....	1175
Update ESXi Hosts by Using Zip Files.....	1177

Remove VIBs from a Host.....	1178
Add Third-Party Extensions to Hosts with an ESXCLI Command.....	1181
Perform a Dry Run of an ESXCLI Installation or Upgrade.....	1181
List the Available VIBs and Profiles After Host Reboot.....	1181
Display the Image Profile and Acceptance Level of the Host.....	1182
After You Upgrade ESXi Hosts.....	1182
About ESXi Evaluation and Licensed Modes.....	1183
Licensing ESXi Hosts After Upgrade.....	1183
Run the Secure Boot Validation Script After ESXi Upgrade.....	1183
Configure Syslog on ESXi Hosts.....	1184
ESXi Syslog Options.....	1185
Fine-tune Syslog on ESXi Hosts.....	1189
Configure Log Filtering on ESXi Hosts.....	1195
Using vSphere Auto Deploy to Reprovision Hosts.....	1196
Introduction to vSphere Auto Deploy.....	1196
Install and Configure vSphere Auto Deploy.....	1198
vSphere Auto Deploy Preinstallation Checklist.....	1198
Prepare Your System for vSphere Auto Deploy.....	1199
Using vSphere Auto Deploy Cmdlets.....	1201
Set Up Bulk Licensing.....	1202
Reprovisioning Hosts.....	1203
Reprovision Hosts with Simple Reboot Operations.....	1203
Use PowerCLI To Reprovision a Host.....	1204
Write a Rule and Assign a Host Profile to Hosts.....	1205
Test and Repair Rule Compliance.....	1206
Collect Logs to Troubleshoot ESXi Hosts.....	1207
vCenter Server Installation and Setup.....	1208
Introduction to vSphere Installation and Setup.....	1208
How to Install and Set Up vSphere.....	1208
vCenter Server Components and Services.....	1209
What is vCenter Server Appliance.....	1211
Understanding vSphere Domains and Domain Names.....	1212
vCenter Enhanced Linked Mode.....	1212
vCenter Enhanced Linked Mode for vCenter Server Appliance.....	1212
Deploying the vCenter Server Appliance.....	1214
System Requirements for the vCenter Server Appliance.....	1215
Hardware Requirements for the vCenter Server Appliance.....	1215
Storage Requirements for the vCenter Server Appliance.....	1215
Software Requirements for the vCenter Server Appliance.....	1216
Required Ports for vCenter Server.....	1216

DNS Requirements for the vCenter Server Appliance.....	1217
vSphere Client Software Requirements.....	1217
Preparing for Deployment of the vCenter Server Appliance.....	1217
System Requirements for the vCenter Server Installer.....	1218
Download and Mount the vCenter Server Installer.....	1218
Synchronizing Clocks on the vSphere Network.....	1219
Synchronizing the Client and Server System Clock.....	1220
Prerequisites for Deploying the vCenter Server Appliance.....	1220
GUI Deployment of the vCenter Server Appliance.....	1221
Required Information for Deploying a vCenter Server Appliance.....	1222
Deploy the vCenter Server Appliance by Using the GUI.....	1224
CLI Deployment of the vCenter Server Appliance.....	1229
Prepare Your JSON Configuration File for CLI Deployment.....	1229
Syntax of the CLI Deployment Command.....	1238
Deploy a vCenter Server Appliance by Using the CLI.....	1239
Deploy Multiple vCenter Server Appliances Using the CLI.....	1240
File-Based Backup and Restore of vCenter Server.....	1241
Considerations and Limitations for File-Based Backup and Restore.....	1241
Schedule a File-Based Backup.....	1244
Manually Back Up vCenter Server by Using the vCenter Server Management Interface.....	1245
Restore vCenter Server from a File-Based Backup.....	1246
Stage 1 - Deploy a New Appliance.....	1248
Stage 2 - Transfer Data to the Newly Deployed Appliance.....	1251
Image-Based Backup and Restore of a vCenter Server Environment.....	1251
Considerations and Limitations for Image-Based Backup and Restore.....	1251
Restore a vCenter Server Image-based Environment.....	1254
Restore a vCenter Server Instance.....	1256
Restore a vCenter Enhanced Linked Mode Environment.....	1256
After You Deploy the vCenter Server Appliance.....	1256
Log in to Your vCenter Server.....	1257
Repoint vCenter Server to Another vCenter Server in a Different Domain.....	1257
Repoint a Single vCenter Server Node to an Existing Domain without a Replication Partner.....	1258
Repoint a vCenter Server Node to an Existing Domain with a Replication Partner.....	1260
Repoint a vCenter Server Node to a New Domain.....	1262
Syntax of the Domain Repoint Command.....	1264
Understanding Tagging and Authorization Conflicts.....	1265
vCenter Server Domain Repoint License Considerations.....	1268
Troubleshooting vCenter Server Installation or Deployment.....	1268
Collecting Logs for Troubleshooting a vCenter Server Installation or Upgrade.....	1268
Retrieve Installation Logs Manually.....	1268

Collect Deployment Log Files for the vCenter Server Appliance.....	1269
Export a vCenter Server Support Bundle for Troubleshooting.....	1269
vCenter Server Upgrade.....	1270
vSphere Upgrade Options.....	1270
How to Upgrade vSphere.....	1270
Overview of the vCenter Server Upgrade Process.....	1272
vCenter Server Upgrade Compatibility.....	1274
Differences Between vSphere Upgrades, Patches, Updates, and Migrations.....	1275
Component Behaviour Changes That Affect vCenter Server Upgrade.....	1276
Removal of Platform Services Controller.....	1276
Upgrade or Migration for vCenter Server Instances with an External Platform Services Controller.....	1277
Upgrading or Migrating to vSphere License Service.....	1277
Upgrading the Trust Authority vCenter Server.....	1278
Support for Upgrading a vCenter Server with Multi-Homing.....	1278
Support for Federal Information Processing Standard 140-2.....	1278
Support for Transport Security Layer 1.3.....	1278
Moving from a Deprecated to a Supported vCenter Server Deployment Topology Before Upgrade or Migration.....	1279
Example Upgrade Paths from vCenter Server6.7 to vCenter Server8.0.....	1280
Example Migration Paths from vCenter Server for Windows to vCenter Server8.0.....	1281
Upgrading the vCenter Server Appliance.....	1282
About the Upgrade Process of the vCenter Server Appliance.....	1284
System Requirements for the vCenter Server Appliance.....	1285
Hardware Requirements for the vCenter Server Appliance.....	1285
Storage Requirements for the vCenter Server Appliance.....	1285
Software Requirements for the vCenter Server Appliance.....	1286
Required Ports for vCenter Server.....	1286
DNS Requirements for the vCenter Server Appliance.....	1287
vSphere Client Software Requirements.....	1287
Preparing to Upgrade the vCenter Server Appliance.....	1287
System Requirements for the vCenter Server Installer.....	1288
Download and Mount the vCenter Server Installer.....	1288
Synchronizing Clocks on the vSphere Network.....	1289
Transferring Data from an Existing vCenter Server Appliance.....	1290
Prepare ESXi Hosts for vCenter Server Appliance Upgrade.....	1291
Determine the Oracle Database Size and the Storage Size for the New Appliance.....	1293
Download and Run VMware Migration Assistant on the Source Update Manager Machine.....	1295
Prerequisites for Upgrading the vCenter Server Appliance.....	1295
GUI Upgrade of the vCenter Server Appliance.....	1297
Required Information for Upgrading a vCenter Server Appliance 7.0 or 8.0.....	1298

Upgrade a vCenter Server Appliance 6.7 with an Embedded Platform Services Controller or 7.0 by Using the GUI.....	1302
Upgrade a vCenter Server Appliance 7.0 or 7.0 with an External Platform Services Controller Instance by Using the GUI.....	1308
Upgrading vCenter Server in High Availability Environments.....	1313
Prerequisites for Upgrading vCenter Server High Availability Environments.....	1313
Upgrade a vCenter Server Appliance 7.0 or 6.7 HA Cluster with an Embedded Platform Services Controller by Using the GUI.....	1314
Upgrade a vCenter Server Appliance 6.7 HA Cluster with an External Platform Services Controller by Using the GUI.....	1318
CLI Upgrade of the vCenter Server Appliance.....	1323
Prepare Your JSON Configuration File for CLI Upgrade.....	1323
Upgrade a vCenter Server Appliance by Using the CLI.....	1335
Syntax of the CLI Upgrade Command.....	1336
Reduced Downtime Upgrade.....	1337
About the Reduced Downtime Upgrade Process.....	1338
Configure Repository URL.....	1338
Download and Mount the ISO Image.....	1339
Upgrade the vCenter Lifecycle Plug-in.....	1340
Configure the Target vCenter Server Appliance.....	1340
Prepare Upgrade and Switch Over to the Target vCenter Server Appliance.....	1342
Migrating vCenter Server for Windows to vCenter Server Appliance.....	1342
How to Migrate from vCenter Server on Windows to vCenter Server Appliance.....	1342
Migration of Update Manager from Windows to a vCenter Server Appliance 8.0.....	1344
System Requirements for Migrating vCenter Server Deployments to vCenter Server Appliance Deployments...	1344
Pre-migration Checks.....	1345
Known Limitations.....	1346
Preparing for Migration.....	1346
Synchronize ESXi Clocks with a Network Time Server.....	1347
Preparing vCenter Server Databases for Migration.....	1347
Prepare Managed ESXi Hosts for Migration.....	1349
Preparing vCenter Server Certificates for Migration.....	1349
System Requirements for the vCenter Server Installer.....	1350
Determine the Oracle Database Size and the Storage Size for the New Appliance.....	1351
Determine the Microsoft SQL Server Database Size and the Storage Size for the New vCenter Server Appliance.....	1353
Download and Run VMware Migration Assistant on the Source Windows Machine.....	1354
Prerequisites for Migrating vCenter Server.....	1355
Required Information for Migrating vCenter Server from Windows to an Appliance.....	1356
GUI Migration of vCenter Server with an Embedded Platform Services Controller to an vCenter Server Appliance.....	1360

Deploy the OVA File for Migrating to the Target vCenter Server Appliance.....	1361
Set Up the Target vCenter Server Appliance.....	1365
GUI Migration of vCenter Server with an External Platform Services Controller to an Appliance.....	1365
Deploy the OVA File for the Target vCenter Server Appliance.....	1367
Set Up the Target vCenter Server Appliance.....	1370
CLI Migration of a vCenter Server Installation from Windows to an Appliance.....	1371
Prepare JSON Configuration Files for CLI Migration.....	1371
JSON Templates for CLI Migration of vCenter Server for Windows.....	1372
About Migrating a vCenter Server for Windows with an External Platform Services Controller Using the CLI.....	1373
Migration Configuration Parameters.....	1375
Run a Pre-Check Before a CLI Migration to vCenter Server Appliance.....	1383
Perform a CLI Migration of vCenter Server from Windows to an Appliance.....	1383
Syntax of the CLI Migrate Command.....	1384
After Upgrading or Migrating vCenter Server.....	1386
Verify Your vCenter Server Upgrade or Migration Is Successful.....	1387
Log in to Your vCenter Server.....	1387
Decommission the Platform Services Controller.....	1388
Decommission an External Platform Services Controller Appliance.....	1388
Decommission an External Platform Server Controller for Windows.....	1389
Identity Sources for vCenter Server with vCenter Single Sign-On.....	1389
Reregister Plug-In Solution in vCenter Server After Upgrade or Migration.....	1390
Roll Back a vCenter Server Upgrade or vCenter Server on Windows Migration.....	1390
Monitor and Manage Historical Data Migration.....	1390
Managing Software Updates, Upgrades, and Product Compatibility.....	1391
Monitor Interoperability for the Current vCenter Server Version.....	1392
Generate an Interoperability Report for Target vCenter Server Versions.....	1392
Generate Pre-Update Reports.....	1393
Patching and Updating vCenter Server8.0 Deployments.....	1394
Patching vCenter Server.....	1394
Patching vCenter Server Using the vCenter Server Management Interface.....	1394
Patching the vCenter Server Appliance by Using the Appliance Shell.....	1401
Patching a vCenter High Availability Environment.....	1405
Troubleshooting a vSphere Upgrade.....	1405
Collect Installation Logs for vCenter Server Appliance.....	1405
Errors and Warnings Returned by the Installation and Upgrade Precheck Script.....	1406
Upgrade Issues with vCenter Server Containing Host Profiles.....	1408
Roll Back a vCenter Server Instance on Windows When vCenter Server Upgrade Fails.....	1408
Collect Logs to Troubleshoot ESXi Hosts.....	1409
vSphere Authentication.....	1410

Getting Started with vSphere Certificate Management and Authentication.....	1411
Managing vCenter Server Certificates.....	1413
Manage vCenter Server Certificates Using the vSphere Client.....	1413
Manage vCenter Server Certificates Using CLIs.....	1414
Managing vCenter Server Authentication Services.....	1414
Manage vCenter Server Authentication Services Using the vSphere Client.....	1415
Manage vCenter Server Authentication Services Using Scripts.....	1415
Managing the vCenter Server.....	1416
Manage vCenter Server Using the Management Interface.....	1416
Manage vCenter Server Using the vCenter Server Shell.....	1416
Add a vCenter Server to an Active Directory Domain.....	1417
vSphere Security Certificates.....	1417
vSphere Certificate Requirements for Different Solution Paths.....	1418
vSphere Certificate Management.....	1422
Replacing vSphere Certificates.....	1423
Where vSphere Uses Certificates.....	1426
VMware Certificate Authority and VMware Core Identity Services.....	1428
VMware Endpoint Certificate Store.....	1428
Managing vSphere Certificate Revocation.....	1429
Replacing vSphere Certificates in Large Deployments.....	1430
Managing Certificates Using the vSphere Client.....	1431
Explore Certificate Stores Using the vSphere Client.....	1432
Set the Threshold for vCenter Certificate Expiration Warnings Using the vSphere Client.....	1432
Renew VMCA Certificates with New VMCA-Signed Certificates Using the vSphere Client.....	1432
Replace Certificates with Custom Certificates Using the vSphere Client.....	1433
Managing Certificates Using the vSphere Certificate Manager Utility.....	1436
Regenerate a New VMCA Root Certificate and Replace All Certificates Using the Certificate Manager.....	1437
Make VMCA an Intermediate Certificate Authority Using the Certificate Manager.....	1438
Replace All Certificates with a Custom Certificate Using the Certificate Manager.....	1442
Revert Last Performed Operation by Republishing Old Certificates Using the Certificate Manager.....	1445
Reset All Certificates Using the Certificate Manager.....	1445
Manual vSphere Certificate Replacement.....	1446
Guidelines for Stopping and Starting of vCenter Server Services.....	1446
Replace Existing VMCA-Signed Certificates with New VMCA-Signed Certificates Using the CLI.....	1446
Make VMCA Into an Intermediate Certificate Authority Using the CLI.....	1453
Replace Certificates with Custom Certificates Using the CLI.....	1460
vSphere Certificates and Services CLI Command Reference.....	1463
certool Initialization Commands Reference.....	1465
certool Management Commands Reference.....	1467
vecs-cli Command Reference.....	1470

dir-cli Command Reference.....	1474
vSphere Authentication with vCenter Single Sign-On.....	1480
How vCenter Single Sign-On Protects Your Environment.....	1481
vCenter Server Identity Provider Federation.....	1483
How vCenter Server Identity Provider Federation Works.....	1484
vCenter Server Identity Provider Federation Caveats and Interoperability.....	1486
vCenter Server Identity Provider Federation Life Cycle.....	1488
vCenter Server Identity Provider Federation and Enhanced Linked Mode.....	1489
Activation Process for External Identity Providers in Enhanced Linked Mode Configurations.....	1493
Configuring vCenter Server Identity Provider Federation.....	1494
vCenter Server Identity Provider Federation Configuration Process Flow.....	1494
Use the Trusted Root Certificates Store Instead of the JRE truststore.....	1498
Configure vCenter Server Identity Provider Federation for AD FS.....	1498
Configure vCenter Server Identity Provider Federation for Okta.....	1501
Configure vCenter Server Identity Provider Federation for Microsoft Entra ID.....	1504
Configuring vCenter Server Identity Provider for PingFederate.....	1507
Configure VMware Single Sign-On.....	1521
Managing VMware Identity Services.....	1522
vCenter Single Sign-On.....	1524
vCenter Single Sign-On Components.....	1524
Using vCenter Single Sign-On with vSphere.....	1525
Groups in the vCenter Single Sign-On Domain.....	1525
Configuring vCenter Single Sign-On Identity Sources.....	1527
Identity Sources for vCenter Server with vCenter Single Sign-On.....	1527
Set the Default Domain for vCenter Single Sign-On.....	1528
Add or Edit a vCenter Single Sign-On Identity Source.....	1528
Managing the vCenter Server Security Token Service.....	1532
Refresh a vCenter Server STS Certificate Using the vSphere Client.....	1533
Import and Replace a vCenter Server STS Certificate Using the vSphere Client.....	1534
Replace a vCenter Server STS Certificate Using the Command Line.....	1535
View the Active vCenter Server STS Signing Certificate Chain Using the vSphere Client.....	1536
Determine the Expiration Date of an LDAPS SSL Certificate Using the Command Line.....	1536
Managing vCenter Single Sign-On Policies.....	1537
Edit the vCenter Single Sign-On Password Policy.....	1537
Edit the vCenter Single Sign-On Lockout Policy.....	1538
Edit the vCenter Single Sign-On Token Policy.....	1539
Edit Password Expiration Notification for Active Directory (Integrated Windows Authentication) Users.....	1540
Managing vCenter Single Sign-On Users and Groups.....	1540
Add vCenter Single Sign-On Users.....	1541
Deactivate and Activate vCenter Single Sign-On Users.....	1541

Delete a vCenter Single Sign-On User.....	1542
Edit a vCenter Single Sign-On User.....	1542
Add a vCenter Single Sign-On Group.....	1543
Add Members to a vCenter Single Sign-On Group.....	1544
Remove Members from a vCenter Single Sign-On Group.....	1544
Change Your vCenter Single Sign-On Password.....	1545
Other vSphere Authentication Options.....	1545
Smart Card Authentication Login.....	1546
Configuring and Using Smart Card Authentication.....	1547
Set Up RSA SecurID Authentication.....	1553
Managing the Login Message to the vSphere Client Login Page.....	1554
Manage the Login Message to the vSphere Client Login Page.....	1554
vCenter Single Sign-On Security Best Practices.....	1555
Troubleshooting vCenter Server Authentication.....	1555
Determining the Cause of a Lookup Service Error.....	1555
Unable to Log In Using Active Directory Domain Authentication.....	1557
vCenter Server Login Fails Because the User Account Is Locked.....	1557
VMware Directory Service Replication Can Take a Long Time.....	1558
Export a vCenter Server Support Bundle.....	1558
vCenter Server Authentication Services Logs Reference.....	1558
Managing Host and Cluster Lifecycle.....	1560
Updated Information.....	1560
What is vSphere Lifecycle Manager.....	1560
The vSphere Lifecycle Manager User Interface in the vSphere Client.....	1562
Using vSphere Lifecycle Manager with VMware vSphere® Distributed Services Engine™.....	1564
Software Packaging Units That vSphere Lifecycle Manager Can Consume.....	1565
vSphere Lifecycle Manager Baselines and Images.....	1568
What Are the Requirements for Using vSphere Lifecycle Manager.....	1572
Required Privileges for Using vSphere Lifecycle Manager and vSphere Configuration Profiles.....	1573
The vSphere Lifecycle Manager Depot.....	1577
Browsing the vSphere Lifecycle Manager Depot.....	1580
Working With the vSphere Lifecycle Manager Depot.....	1582
Import Updates to the vSphere Lifecycle Manager Depot.....	1582
Import an ISO Image to the vSphere Lifecycle Manager Depot.....	1583
Delete an ISO Image from the vSphere Lifecycle Manager Depot.....	1583
Synchronize the vSphere Lifecycle Manager Depot.....	1584
Configuring the vSphere Lifecycle Manager Download Sources.....	1585
Configure vSphere Lifecycle Manager to Use the Internet as a Download Source.....	1586
Configure vSphere Lifecycle Manager to Use a Shared Repository as a Download Source.....	1586
Run the VMware vSphere vSphere Lifecycle Manager Update Download Task.....	1588

Add a New Download Source.....	1588
Add a New Download Source.....	1588
Modify a Download Source.....	1589
Configure the vSphere Lifecycle Manager Automatic Download Task.....	1589
Installing, Setting Up, and Using the Update Manager Download Service.....	1590
Installing UMDS.....	1591
Setting Up UMDS.....	1592
Using UMDS.....	1594
vSphere Lifecycle Manager Remediation Settings.....	1596
Configuring vSphere Lifecycle Manager for Fast Upgrades.....	1597
Configuring vSphere Lifecycle Manager for Live Patches.....	1600
How to Configure the vSphere Lifecycle Manager Remediation Settings.....	1600
Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image.....	1601
Override the Global vSphere Lifecycle Manager Remediation Settings for a Cluster or Host That You Manage with a Single Image.....	1604
Configure the vSphere Lifecycle Manager Remediation Settings for Hosts and Clusters That You Manage with Baselines.....	1607
Configure vSphere Lifecycle Manager Remediation Settings for vSAN Clusters that You Manage with vSphere Lifecycle Manager Baselines.....	1609
Creating and Managing vSphere Lifecycle Manager Clusters.....	1609
Create a Cluster That Uses a Single Image by Composing a New Image.....	1610
Create a Cluster That Uses a Single Image by Importing an Image from a Host.....	1611
Add Hosts to a Cluster that Uses a Single Image.....	1613
Managing Standalone ESXi Hosts with vSphere Lifecycle Manager Images.....	1616
Use the Current Image of a Host When Adding the Host to a Data Center or Folder.....	1616
Import an Image from a Reference Host When Adding a Host to a Data Center or Folder.....	1617
Set Up a New Image to a Host When Adding the Host to a Data Center or Folder.....	1618
How to Use vSphere Lifecycle Manager Images.....	1619
Convert a Cluster or a Host That Uses Baselines Into a Cluster or a Host That Uses vSphere Lifecycle Manager Images.....	1620
Cluster or Standalone Host Eligibility to Use vSphere Lifecycle Manager Images.....	1621
Option 1: Use an Image from a Host in the Cluster.....	1622
Option 2: Import an Image from a JSON File.....	1627
Option 3: Set Up a New Image Manually.....	1628
Working with Images.....	1630
Edit a vSphere Lifecycle Manager Image.....	1631
Export or Import a vSphere Lifecycle Manager Image.....	1633
Check the Compliance Against a Single Image.....	1635
Check the Compliance of a Cluster Against a vSphere Lifecycle Manager Image.....	1638
View Host Compliance Information.....	1639

Staging vSphere Lifecycle Manager Images to ESXi Hosts in a Cluster or Standalone Hosts.....	1640
Stage an Image to a Cluster.....	1641
Stage an Image to a Single Host Within a Cluster or a Standalone Host.....	1641
Run a Remediation Pre-Check for a Cluster, a Host Within a Cluster, or a Standalone Host.....	1642
Run a Remediation Pre-Check for a Single Host Within a Cluster or a Standalone Host.....	1642
Remediating a Cluster or a Standalone Host Against a Single Image.....	1643
Remediate a Cluster Against a Single Image.....	1644
Remediate a Single Host Within a Cluster or a Standalone Host Against an Image.....	1645
View Last Remediation or Remediation Pre-Check Results for a Cluster or a Standalone Host that Uses a Single Image.....	1646
Manage Depot Overrides for a Cluster or a Standalone Host.....	1646
vSphere Lifecycle Manager Recommended Images.....	1647
Check for Recommended Images.....	1649
View the Recommended or Not Recommended Images for a Cluster or a Standalone Host.....	1649
Use the vSphere Lifecycle Manager Recommended Image as the New Image for a Cluster or a Standalone Host.....	1650
Firmware Updates with vSphere Lifecycle Manager.....	1651
Deploying Hardware Support Managers.....	1652
Use an Image for Firmware Updates.....	1653
vSphere Lifecycle Manager Hardware Compatibility Checks for Clusters and Hosts.....	1654
Cluster-Level Hardware Compatibility Checks.....	1655
NIC validation.....	1658
Check the Hardware Compatibility of a Cluster.....	1660
Change the Compliance Status of a Disk Device Manually.....	1661
Host-Level Hardware Compatibility Checks.....	1662
Using vSphere Lifecycle Manager Baselines and Baseline Groups.....	1665
Types of vSphere Lifecycle Manager Baselines and Baseline Groups.....	1665
Create, Edit, or Delete vSphere Lifecycle Manager Baselines and Baseline Groups.....	1668
Create a Host Baseline Group.....	1671
Edit a vSphere Lifecycle Manager Baseline or Baseline Group.....	1671
Add or Remove a Single Update from a Custom Baseline.....	1672
Delete Baselines and Baseline Groups.....	1672
Attach Baselines and Baseline Groups to Objects.....	1672
Detach Baselines and Baseline Groups from Objects.....	1673
Checking Compliance Against vSphere Lifecycle Manager Baselines and Baseline Groups.....	1674
Initiate a Compliance Check for ESXi Hosts Manually.....	1674
Schedule Regular Compliance Checks for ESXi Hosts.....	1675
View Information About the Patches, Extensions, and ISO Images in a Baseline.....	1675
Host Upgrade Compliance Messages.....	1677
Host Upgrade Compliance Messages When Cisco Nexus 1000V Is Present.....	1679

Viewing Compliance Information About ESXi Hosts and Updates.....	1680
The vSphere Lifecycle Manager Compliance View.....	1681
Compliance Statuses of Updates.....	1682
Compliance Statuses of ESXi Hosts, Baselines, and Baseline Groups.....	1683
View Compliance Information About ESXi Hosts.....	1684
Stage Patches and Extensions to ESXi Hosts.....	1685
Remediating ESXi Hosts Against vSphere Lifecycle Manager Baselines and Baseline Groups.....	1686
Understanding the Remediation Operation.....	1688
Types of Host Remediation.....	1688
Remediating Hosts in a Cluster.....	1689
Remediating Hosts That Contain Third-Party Software.....	1690
Remediating ESXi6.7 or ESXi7.0 Hosts Against an ESXi8.0 Image.....	1691
Generate a Pre-Remediation Check Report.....	1691
Remediate ESXi Hosts Against a Single Baseline or Multiple Baselines.....	1694
Using vSphere Lifecycle Manager to Migrate an NSX Virtual Distributed Switch to a vSphere Distributed Switch.....	1696
Upgrading Virtual Machines with vSphere Lifecycle Manager.....	1697
Configure Virtual Machine Rollback Settings.....	1698
Checking the Status of Virtual Machines.....	1698
Check the Status of an Individual Virtual Machine.....	1699
Check the Status of the Virtual Machines in a Container Object.....	1699
The VMware Tools Status.....	1700
Upgrading Virtual Machines.....	1700
Upgrade the VM Hardware Compatibility of Virtual Machines.....	1701
Upgrade the VMware Tools Version of Virtual Machines.....	1702
Automatically Upgrade VMware Tools on Reboot.....	1703
vSphere Lifecycle Manager Images and Other VMware Products and Solutions.....	1704
vSAN Clusters and vSphere Lifecycle Manager.....	1705
Remediation Specifics of vSAN Clusters.....	1706
Updating Firmware in vSAN Clusters.....	1709
About Recommendation Baseline Groups.....	1709
vSphere Lifecycle Manager and vSphere Supervisor.....	1710
vSphere Lifecycle Manager and vSphere Supervisor with vSphere Networking.....	1710
vSphere Lifecycle Manager and vSphere Supervisor with NSX Networking.....	1711
vSphere Lifecycle Manager and VMware NSX®.....	1713
Using vSphere Lifecycle Manager Baselines to Upgrade ESXi Hosts in an Environment With VMware NSX® 3.0.....	1713
Using vSphere Lifecycle Manager Images in an Environment With NSX 3.1.....	1713
Using vSphere Lifecycle Manager Images for Standalone Hosts with NSX 4.2 and later.....	1715
Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level.....	1716

Enable vSphere Configuration Profiles During Cluster Creation.....	1719
Transition to Using vSphere Configuration Profiles.....	1720
Working with Configuration Documents and Configuration Schemas.....	1723
Export the Cluster Configuration or Cluster Configuration Schema.....	1724
Export the Draft Cluster Configuration or Draft Cluster Configuration Schema.....	1724
Extract The Configuration from a Reference Host.....	1725
View Host Settings in the Desired Configuration.....	1725
Create a Draft Configuration in vSphere Client.....	1728
Edit the Host Settings of a Draft Configuration in vSphere Client.....	1729
Import a Configuration Document or a Configuration from a Reference Host Into an Existing Draft Configuration.....	1733
Delete a Draft Configuration for a Cluster.....	1734
Check Host Compliance Against the Desired Configuration for the Cluster.....	1734
Set Up Alarms for Monitoring Configuration Drifts.....	1735
Run a Remediation Pre-Check or a Draft Pre-Check.....	1736
Remediate a Cluster Against the Desired or Draft Configuration.....	1737
Backup and Restore Scenarios When Using vSphere Lifecycle Manager.....	1738
vCenter Server Configuration.....	1740
How to Configure vCenter Server.....	1740
What Happened to the Platform Services Controller.....	1741
Using the vCenter Server Management Interface to Configure vCenter Server.....	1741
Log In to the vCenter Server Management Interface.....	1741
View vCenter ServerHealth Status.....	1742
Reboot or Shut Down vCenter Server.....	1742
Create a Support Bundle.....	1742
Monitor CPU and Memory Use.....	1743
Monitor Disk Use.....	1743
Monitor Network Use.....	1743
Monitor Database Use.....	1743
Enable or Disable SSH and Bash Shell Access.....	1744
Configure the DNS, IP Address, and Proxy Settings.....	1744
Reconfigure the Primary Network Identifier.....	1746
Edit the Firewall Settings.....	1747
Configure the System Time Zone and Time Synchronization Settings.....	1748
Start, Stop, and Restart Services.....	1748
Configure Update Settings.....	1748
Change the Password and Password Expiration Settings of the Root User.....	1749
Forward vCenter Server Log Files to Remote Syslog Server.....	1752
Configure and Schedule Backups.....	1752
Using the vSphere Client to Configure vCenter Server.....	1753

Configuring vCenter Server.....	1753
Configuring Statistics Settings.....	1753
Configure Runtime Settings for vCenter Server.....	1755
Configure User Directory Settings.....	1756
Configure Mail Sender Settings.....	1756
Configure SNMP Settings.....	1758
View Port Settings.....	1758
Configure Timeout Settings.....	1758
Configure Logging Options.....	1759
Configure Database Settings.....	1760
Verify SSL Certificates for Legacy Hosts.....	1761
Configure Advanced Settings.....	1761
Send a Message to Other Logged In Users.....	1762
Join or Leave an Active Directory Domain.....	1762
Add a User to the SystemConfiguration.BashShellAdministrators Group.....	1764
Reboot a Node.....	1764
View the Health Status of Nodes.....	1765
Export a Support Bundle.....	1765
Using the Appliance Shell to Configure vCenter Server.....	1766
Access the Appliance Shell.....	1766
Enable and Access the Bash Shell from the Appliance Shell.....	1766
Keyboard Shortcuts for Editing Commands.....	1766
Get Help on APIs and Utilities.....	1768
Plug-Ins in vCenter Server Shell.....	1768
Browse Log Files Using showlog Plug-In.....	1769
API Commands in Appliance Shell.....	1769
Configuring SNMP for vCenter Server.....	1774
Configure the SNMP Agent for Polling.....	1774
Configure vCenter Server for SNMP v1 and v2c.....	1774
Configure vCenter Server for SNMP v3.....	1775
Configure SNMP Agent to Filter Notifications.....	1778
Configure SNMP Management Client Software.....	1778
Reset SNMP Settings to Factory Defaults.....	1779
Configuring Time Synchronization Settings in vCenter Server.....	1779
Use VMware Tools Time Synchronization.....	1780
Add or Replace NTP Servers in the vCenter Server Configuration.....	1780
Synchronize the Time in vCenter Server with an NTP Server.....	1781
Managing Local User Accounts in vCenter Server.....	1781
User Roles in vCenter Server.....	1781
Get a List of the Local User Accounts in vCenter Server.....	1781

Create a Local User Account in vCenter Server.....	1782
Update Password of a Local User in vCenter Server.....	1782
Update a Local User Account in vCenter Server.....	1782
Delete a Local User Account in vCenter Server.....	1783
Monitor Health Status and Statistics in vCenter Server.....	1783
Using the vimtop Plug-In to Monitor the Resource Use of Services.....	1784
Monitor Services By Using vimtop in Interactive Mode.....	1784
Interactive Mode Command-Line Options.....	1784
Interactive Mode Single-Key Commands for vimtop.....	1784
Using the Direct Console User Interface to Configure vCenter Server.....	1786
Log In to the Direct Console User Interface.....	1786
Change the Password of the Root User.....	1786
Configure the Management Network of vCenter Server.....	1786
Restart the Management Network of vCenter Server.....	1787
Enable Access to the Bash Shell.....	1787
Access the Bash Shell for Troubleshooting.....	1788
Export a vCenter Server Support Bundle for Troubleshooting.....	1788
vCenter Server and Host Management.....	1789
Main vSphere Concepts and Features.....	1789
Basics of Virtualization.....	1789
Physical Topology of a vSphere Data Center.....	1789
Software Components of vSphere.....	1790
vSphere Cluster Services.....	1793
Monitoring vSphere Cluster Services.....	1793
Maintaining Health of vSphere Cluster Services.....	1794
Client Interfaces for vSphere.....	1795
Managed Inventory Objects in vSphere.....	1796
Optional vCenter Server Components.....	1797
vCenter Enhanced Linked Mode.....	1798
vCenter Server Plug-Ins.....	1798
Using the vSphere Client.....	1799
Use and Configure Your vSphere Client.....	1799
Log in to Your vCenter Server.....	1799
Manage the Login Message to the vSphere Client Login Page.....	1800
Configure the vSphere Client Timeout Value.....	1800
Refresh Data in Your vSphere Client.....	1801
Use the vSphere Client Inventory.....	1801
Use the Global Inventory Lists Navigator.....	1801
Inventory Keyboard Shortcuts in the vSphere Client.....	1802
Sort Your vSphere Client Inventory.....	1803

Use Quick Filters in vSphere Client Inventory.....	1803
Use Advanced Filters in vSphere Client Inventory.....	1803
Search Your vSphere Client Inventory.....	1804
Export vSphere Inventory Lists.....	1805
Manage vCenter Server Client Plug-Ins.....	1805
Start, Stop, and Restart Services.....	1807
Attach a File to a VMware Service Request.....	1808
Provide Feedback for vSphere.....	1808
Provide Feedback with the vSphere Client.....	1808
Use the vSphere Ideas Portal to Submit Suggestions for New Features and Feature Requests.....	1808
Configuring Customer Experience Improvement Program.....	1809
VMware Customer Experience Improvement Program.....	1809
Join or Leave Customer Experience Improvement Program in vSphere Client.....	1809
Organizing Your vSphere Inventory.....	1809
Create a vSphere Data Center.....	1810
Create a Folder.....	1811
Add a Host to a Folder or a Data Center.....	1811
How Do You Create and Configure Clusters in the vSphere Client.....	1812
Create a Cluster.....	1813
Extend and Configure a Cluster by Using the Quickstart Workflow.....	1814
Add a Host to a Cluster by Using the Quickstart Workflow.....	1815
Configure a Cluster by Using the Quickstart Workflow.....	1816
Extend a Cluster with Host Networking Configuration.....	1818
Extend and Configure a Cluster Manually.....	1819
vSphere Tags and Custom Attributes.....	1820
What Are vSphere Tags and How to Use Them.....	1821
Create and Edit a Tag Category.....	1821
Delete a Tag Category.....	1822
Create and Edit a Tag.....	1822
Delete a Tag.....	1823
Assign or Remove a Tag.....	1823
Add Permissions for Tags and Tag Categories.....	1824
vSphere Tagging Best Practices.....	1824
Custom Attributes.....	1825
Add and Edit Custom Attributes.....	1825
Working with vSphere Tasks.....	1827
View Tasks.....	1828
Cancel a Task.....	1828
Filter Tasks.....	1829
Export Tasks.....	1830

Export Tasks with Advanced Export.....	1830
Scheduling vSphere Tasks.....	1831
Create a Scheduled Task.....	1832
Change or Reschedule a Scheduled Task.....	1833
Remove a Scheduled Task.....	1833
vSphere Host Configuration.....	1834
Configure the Boot Device on an ESXi Host.....	1834
Configure Agent VM Settings.....	1834
Set Advanced Host Attributes.....	1835
Synchronizing Clocks on the vSphere Network.....	1835
Editing the Time Configuration Settings of Your ESXi Host.....	1836
Configure the Date and Time on Your ESXi Host Manually.....	1836
Use NTP Servers for Time and Date Synchronization of Your ESXi Host.....	1836
Use PTP for Time and Date Synchronization of Your ESXi Host.....	1837
Manage the NTP and PTP Service Status on Your ESXi Host.....	1838
Delete a Time Synchronization Service on Your ESXi Host.....	1839
Managing Hosts in vCenter Server.....	1839
Disconnecting and Reconnecting an ESXi Host.....	1839
Disconnect a Managed Host.....	1839
Reconnect a Managed Host.....	1840
Reconnect Hosts After Changes to the vCenter Server SSL Certificate.....	1840
Relocate a Host.....	1840
Remove a Host from vCenter Server.....	1840
Reboot or Shut Down an ESXi Host.....	1841
Verify SSL Certificates for Legacy Hosts.....	1842
Enabling SGX Remote Attestation on Multi-Socket ESXi Hosts.....	1842
View the SGX Registration Status of a Host.....	1843
Register a Multi-Socket ESXi Host with the Intel SGX Registration Server.....	1843
vSphere License Management.....	1844
vSphere Licensing and Subscription Terminology and Definitions.....	1844
VMware vSphere Foundation (VVF) Licensing.....	1845
Licensing and Subscription in vSphere.....	1846
Licensing for ESXi Hosts.....	1846
Licensing for vCenter Server.....	1848
Managing vSphere Licenses.....	1849
Create a New License.....	1849
Configuring License Settings for Assets in the vSphere Client.....	1850
Rename a License.....	1852
Remove a License.....	1852
Viewing and Exporting vSphere Licensing Information.....	1853

View Licensing Information About Your vSphere Environment.....	1853
View Available Licenses and Features About a vSphere Product.....	1854
View the vSphere Features That an Asset Can Use.....	1854
View the License Key of a License.....	1854
View the Licensed Features for an Asset.....	1855
Export Licensing Information.....	1855
vCenter Server Domain Repoint License Considerations.....	1855
vSphere Virtual Machines Migration.....	1856
Cold Migration in Your vSphere Environment.....	1857
Migrate a Powered Off or Suspended Virtual Machine.....	1858
Hot Migration.....	1861
Host Configuration for vSphere vMotion.....	1861
Virtual Machine Conditions and Limitations for vSphere vMotion.....	1864
Migrate a Virtual Machine to a New Compute Resource.....	1870
Migration with Storage vMotion.....	1871
Migrate Your Virtual Machine to a New Compute Resource and Storage.....	1873
vSphere vMotion in Environments Without Shared Storage.....	1876
Migrating Virtual Machines Between vCenter Server Systems.....	1877
Requirements for vMotion Between vCenter Server Instances.....	1877
Network Compatibility Checks During vMotion Between vCenter Server Instances.....	1878
MAC Address Management During Migration Between vCenter Server Systems.....	1878
Export or Clone a Virtual Machine with Advanced Cross vCenter vMotion.....	1879
Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion.....	1884
Migrating Virtual Machines Between Servers with DPU Devices.....	1888
Encrypted vSphere vMotion.....	1888
Activate Encrypted vMotion on a Virtual Machine.....	1890
How to Isolate Traffic for Migrating Your Virtual Machines.....	1891
Place vMotion Traffic on the vMotion TCP/IP Stack of a Host.....	1891
Place Traffic for Cold Migration on the Provisioning TCP/IP Stack.....	1892
vCenter Server Limits on Simultaneous Migrations.....	1894
vMotion Migration Compatibility Checks.....	1895
CPU Compatibility and vSphere Enhanced vMotion Compatibility.....	1896
CPU Compatibility Scenarios.....	1897
CPU Families and Feature Sets.....	1897
CPU Compatibility Masks.....	1898
Enhanced vMotion Compatibility.....	1898
Enhanced vMotion Compatibility Requirements for Hosts.....	1899
How Do You Activate vSphere EVC on an Existing Cluster.....	1899
Change the EVC Mode for a Cluster.....	1900
How to Determine the EVC Mode of Your Virtual Machine.....	1901

vMotion Migration and ESXi Hosts with AMD Processors Without 3DNow!.....	1903
How to View Feature Capabilities Details for Your EVC Cluster.....	1903
What Is vSphere Developer Center.....	1904
What Is vSphere API Explorer.....	1904
How Do You Retrieve APIs Using vSphere API Explorer.....	1904
What Is vSphere Code Capture.....	1904
How Do You Record Actions by Using vSphere Code Capture.....	1904
Automating Management Tasks by Using vRealize Orchestrator.....	1905
What Are VMware Aria Automation Orchestrator Workflows.....	1905
What Administration Tasks Can You Perform on vSphere Objects by Using VMware Aria Automation Orchestrator.....	1906
Managing Context Actions of Workflows on Your vSphere Inventory Objects.....	1906
Assign an Automation Orchestrator Group in the vSphere Client.....	1907
Associate Your Workflows with vSphere Inventory Object Types.....	1907
Edit the Context Actions of Your Workflows on vSphere Objects.....	1908
Export the Context Actions of Your Workflows on vSphere Objects.....	1908
Import Context Actions of Your Workflows on vSphere Objects.....	1908
Working with VMware Aria Automation Orchestrator Workflows in the vSphere Client.....	1909
Run Your Workflows on vSphere Inventory Objects.....	1909
Schedule Your Workflows on vSphere Inventory Objects.....	1910
Scheduling Workflows from the VMware Aria Automation Orchestrator view in the vSphere Client.....	1910
View Information About Your VMware Aria Automation Orchestrator Workflow Runs.....	1912
View Information About the Runs of a Specific Workflow.....	1913
View Workflows That Are Waiting for User Input.....	1913
Find a VMware Aria Automation Orchestrator Workflow.....	1913
Workflows for Managing Your vSphere Inventory Objects.....	1914
Working with Headless Systems.....	1914
How to Detect a Headless System.....	1914
Serial Mode Dynamic Switching.....	1914
What Serial Port Modes Does ESXi Support.....	1915
ESXi Keystrokes for Dynamic Serial Mode Switching.....	1915
Switch Serial Modes Through the ESXi CLI.....	1915
What Are the Keystrokes for Controlling the Serial DCUI.....	1916
Troubleshooting Overview.....	1916
Guidelines for Troubleshooting Your Implementation of vSphere.....	1916
How Do You Identify the Symptoms.....	1917
How Do You Define the Problem Space.....	1917
How Do You Test Possible Solutions.....	1917
Troubleshooting with vCenter Server Logs.....	1918
Troubleshooting vCenter Server and ESXi Host Certificates.....	1919

New vCenter Server Certificate Does Not Appear to Load.....	1919
vCenter Server Cannot Connect to Managed Hosts.....	1919
Troubleshooting ESXi Hosts.....	1919
Troubleshooting vSphere HA Host States.....	1919
Authentication Token Manipulation Error.....	1922
Unable to Download VIBs When Using vCenter Server Reverse Proxy.....	1923
Troubleshooting Licensing of ESXi Hosts and vCenter Server.....	1924
Troubleshooting ESXi Host Licensing.....	1924
Unable to Power On a Virtual Machine.....	1925
Unable to Configure or Use a Feature in Your vSphere Implementation.....	1925
vSphere Virtual Machine Administration.....	1926
Updated Information.....	1926
Introduction to VMware vSphere Virtual Machines.....	1926
Virtual Machine Files.....	1927
Virtual Machines and the Virtual Infrastructure.....	1927
Virtual Machine Lifecycle.....	1928
Virtual Machine Components.....	1928
Virtual Machine Hardware Available to vSphere Virtual Machines.....	1929
Virtual Machine Options.....	1931
ThevSphere Client.....	1932
Where to Go From Here.....	1932
Deploying Virtual Machines.....	1932
Create a Virtual Machine with the New Virtual Machine Wizard.....	1933
Clone a Virtual Machine to a Template.....	1935
Deploy a Virtual Machine from a Template.....	1937
Clone an Existing Virtual Machine.....	1939
Cloning a Virtual Machine with Instant Clone.....	1942
Clone a Template to a Template.....	1943
Convert a Template to a Virtual Machine.....	1944
Deploy and Export OVF and OVA Templates.....	1945
Browse VMware Virtual Appliance Marketplace.....	1949
Using Content Libraries.....	1949
Content Library Administrator Role, Permissions Hierarchy, and Inheritance.....	1950
Create and Edit a Content Library.....	1953
Configure Advanced Content Library Settings.....	1957
Create and Delete a Subscription for a Local Library.....	1958
Publish the Library Contents and a Single Template to a Subscriber.....	1961
Synchronize the Contents of a Subscribed Library.....	1961
How to Populate Libraries with Content.....	1963
Manage Content Library Items.....	1966

Creating Virtual Machines and vApps from Templates in a Content Library.....	1969
Deploy a Virtual Machine from an OVF Template in a Content Library.....	1969
Deploy a Virtual Machine from a VM Template in a Content Library.....	1971
Create New vApp From a Template in a Content Library.....	1973
Manage VM Templates in a Content Library.....	1974
Configuring Virtual Machine Hardware.....	1979
Virtual Machine Compatibility.....	1979
Set the Default Compatibility for Virtual Machine Creation.....	1981
Schedule a Compatibility Upgrade for a Single Virtual Machine.....	1982
Change the Default Virtual Machine Compatibility Setting.....	1982
Hardware Features Available with Virtual Machine Compatibility Settings.....	1983
Virtual CPU Configuration and Limitations.....	1991
Configure Virtual Machine CPU Resources.....	1992
Virtual Memory Configuration.....	1999
Change the Memory Configuration.....	1999
Allocate Memory Resources to a Virtual Machine in the VMware Host Client.....	2000
Manage Persistent Memory.....	2001
Change Memory Hot Add Settings.....	2004
Virtual Disk Configuration.....	2005
About Virtual Disk Provisioning Policies.....	2005
Large Capacity Virtual Disk Conditions and Limitations.....	2006
Change the Virtual Disk Configuration.....	2006
Use Disk Shares to Prioritize Virtual Machines.....	2007
Determine the Virtual Disk Format and Convert a Virtual Disk from Thin Provision to a Thick Provision Format.....	2008
How do I Add a Hard Disk to a Virtual Machine.....	2009
SCSI, SATA, and NVMe Storage Controller Conditions, Limitations, and Compatibility.....	2013
Virtual Machine Network Configuration.....	2018
Network Adapter Basics.....	2018
Network Adapters and Legacy Virtual Machines.....	2020
Change the Virtual Machine Network Adapter Configuration.....	2020
How to Add a Network Interface Controller to a Virtual Machine.....	2021
Virtual Machine Parallel and Serial Port Configuration.....	2022
Other Virtual Machine Device Configuration.....	2028
Change the CD/DVD Drive Configuration.....	2028
How do I Add or Modify a Virtual Machine CD or DVD Drive.....	2030
How do I Add a PCI Device to a Virtual Machine.....	2031
How do I Configure 3D Graphics.....	2035
How do I Add a Virtual Watchdog Timer Device to a Virtual Machine.....	2037
How do I Add a Precision Clock Device to a Virtual Machine.....	2038

Virtual Shared Graphics Acceleration.....	2039
Securing Virtual Machines with Intel Software Guard Extensions.....	2045
Enable vSGX on a Virtual Machine.....	2045
Remove vSGX from a Virtual Machine.....	2047
USB Configuration from an ESXi Host to a Virtual Machine.....	2047
How the USB Autoconnect Feature Works.....	2049
vSphere Features Available with USB Passthrough.....	2049
Conditions and Limitations When Configuring USB Devices for vMotion.....	2050
Avoiding Data Loss with USB Devices.....	2050
Connecting USB Devices to an ESXi Host.....	2050
Add USB Devices to an ESXi Host.....	2051
Add a USB Controller to a Virtual Machine.....	2052
Add USB Devices from an ESXi Host to a Virtual Machine.....	2053
Remove USB Devices That Are Connected Through an ESXi Host.....	2053
Remove USB Devices from an ESXi Host.....	2054
USB Configuration from a Client Computer to a Virtual Machine.....	2054
How to Connect USB Devices to a Client Computer.....	2055
Add USB Devices from a Client Computer to a Virtual Machine.....	2056
Remove USB Devices That Are Connected Through a Client Computer.....	2057
Remove a USB Controller from a Virtual Machine.....	2057
Add a Shared Smart Card Reader to Virtual Machines.....	2058
Securing Virtual Machines with Virtual Trusted Platform Module.....	2058
What Is a Virtual Trusted Platform Module.....	2059
Create a Virtual Machine with a Virtual Trusted Platform Module.....	2060
Add Virtual Trusted Platform Module to an Existing Virtual Machine.....	2060
Remove Virtual Trusted Platform Module from a Virtual Machine.....	2061
Identify Virtual Trusted Platform Module Enabled Virtual Machines.....	2062
Securing Virtual Machines with AMD Secure Encrypted Virtualization-Encrypted State.....	2062
vSphere and AMD Secure Encrypted Virtualization-Encrypted State.....	2062
Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the vSphere Client.....	2063
Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the vSphere Client.....	2064
Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the Command Line.....	2064
Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the Command Line.....	2065
Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the vSphere Client.....	2066
Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the Command Line.....	2066

Configuring Virtual Machine Options.....	2067
General Virtual Machine Options.....	2067
Change the Virtual Machine Name.....	2068
View the Virtual Machine Configuration and Working File Location.....	2068
Change the Configured Guest Operating System.....	2068
VMware Remote Console Options.....	2069
Change the Virtual Machine Console Options for Remote Users.....	2069
Configure Virtual Machine Encryption Options.....	2069
Encrypt an Existing Virtual Machine or Virtual Disk.....	2069
Decrypt an Encrypted Virtual Machine or Virtual Disk.....	2070
Clone an Encrypted Virtual Machine.....	2071
Virtual Machine Power Management Options.....	2072
Manage Power Management Settings for a Virtual Machine.....	2073
Configuring VMware Tools Options.....	2073
Configure the Virtual Machine Power States.....	2073
Synchronize the Time of a Virtual Machine Guest Operating System with the Host.....	2074
Virtualization Based Security.....	2075
Activate Virtualization-based Security on an Existing Virtual Machine.....	2075
Activate Virtualization-based Security on the Guest Operating System.....	2076
Deactivate Virtualization-based Security.....	2077
Identify VBS-Enabled Virtual Machines.....	2077
Configure Virtual Machine Boot Options.....	2077
Activate or Deactivate UEFI Secure Boot for a Virtual Machine.....	2077
Delay the Boot Sequence.....	2079
Configuring Virtual Machine Advanced Options.....	2079
Enable Virtual Machine Logging.....	2079
Configure Virtual Machine Debugging and Statistics.....	2080
Change the Swap File Location.....	2080
Configure the Virtual Machine Latency Sensitivity.....	2080
Set a Stun Time Limit for vGPU Virtual Machines.....	2081
Configure Fibre Channel NPIV Settings.....	2084
Configure the CPU Topology of a Virtual Machine.....	2085
Configure Virtual Machine Advanced File Parameters.....	2087
Managing Multi-Tiered Applications with vSphere vApp.....	2087
Create a vApp.....	2088
Perform vApp Power Operations.....	2088
Create or Add an Object to a vApp.....	2089
Clone a vApp.....	2090
Edit vApp Notes.....	2090
Configure vApp Properties.....	2090

Edit vApp Settings.....	2091
Configure vApp CPU and Memory Resources.....	2091
Configure vApp IP Allocation Policy.....	2092
Configure vApp Startup and Shutdown Options.....	2093
Configure vApp Product Properties.....	2093
View vApp License Agreements.....	2094
Add a Network Protocol Profile.....	2094
Assign a Port Group or Network to a Network Protocol Profile.....	2095
Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine or vApp.....	2096
Virtual Machine vApp Options.....	2097
Activate vApp Options for a Virtual Machine.....	2097
Edit Application Properties and OVF Deployment Options for a Virtual Machine.....	2097
OVF Authoring Options for a Virtual Machine.....	2098
Monitoring Solutions with the vCenter Solutions Manager.....	2101
View Solutions.....	2101
Managing Virtual Machines.....	2101
Install a Guest Operating System.....	2102
Customizing Guest Operating Systems.....	2103
Guest Operating System Customization Requirements.....	2103
Create a vCenter Server Application to Generate Computer Names and IP Addresses.....	2104
Customize Windows During Cloning or Deployment.....	2105
Customize Linux During Cloning or Deployment.....	2106
Apply a Customization Specification to an Existing Virtual Machine.....	2107
Create and Manage Customization Specifications for Windows and Linux.....	2107
Configure User Mappings on Guest Operating Systems.....	2115
Managing Power States of a Virtual Machine.....	2116
Edit Virtual Machine Startup and Shutdown Settings.....	2117
Using a Virtual Machine Console.....	2118
Install the VMware Remote Console Application.....	2119
Start the VMware Remote Console Application.....	2122
Open the Web Console.....	2122
Manage the VMware Remote Console Proxy Configuration.....	2123
Manage the VMware Remote Console Proxy Settings.....	2124
Answer Virtual Machine Questions.....	2125
Remove and Reregister VMs and VM Templates.....	2126
Manage Virtual Machine Templates.....	2127
Manage Virtual Machines With Snapshots.....	2128
Snapshot Files and Limitations.....	2129
Manage Snapshots.....	2131
Enhanced vMotion Compatibility as a Virtual Machine Attribute.....	2141

Virtual Machine Storage DRS Rules.....	2144
Distributing Content with GuestStore.....	2147
Sharing Data Between the vSphere Client and a Virtual Machine Guest Operating System with DataSets.....	2149
vSphere Virtual Machines Migration.....	2152
How do I Upgrade Virtual Machines.....	2153
Downtime for Upgrading Virtual Machines.....	2154
Upgrade the Compatibility of a Virtual Machine Manually.....	2155
Schedule a Compatibility Upgrade for a Virtual Machine.....	2156
Required vCenter Server Privileges for Common Tasks.....	2157
Troubleshooting Overview.....	2159
Guidelines for Troubleshooting Your Implementation of vSphere.....	2159
How Do You Identify the Symptoms.....	2160
How Do You Define the Problem Space.....	2160
How Do You Test Possible Solutions.....	2160
Troubleshooting with vCenter Server Logs.....	2161
Troubleshooting Virtual Machines.....	2162
Troubleshooting USB Passthrough Devices.....	2162
Error Message When You Try to Migrate Virtual Machine with USB Devices Attached.....	2162
Cannot Copy Data From an ESXi Host to a USB Device That Is Connected to the Host.....	2162
Recover Orphaned Virtual Machines.....	2163
vSphere Host Profiles.....	2164
What Are vSphere Host Profiles.....	2164
Reference Host Independence.....	2164
Reference Host Selection.....	2164
Host Profiles and vSphere Auto Deploy.....	2165
How Do You Manage vSphere Host Profiles in the vSphere Client.....	2166
How Do You Get Started with vSphere Host Profiles.....	2166
How Do You Create a vSphere Host Profile?.....	2166
How Do You Attach or Detach Hosts and Clusters to a vSphere Host Profile?.....	2167
How Do You Duplicate a vSphere Host Profile?.....	2168
How Do You Copy Settings from a Reference Host to a vSphere Host Profile?.....	2169
How Do You Import or Export a vSphere Host Profile?.....	2169
How Do You Copy Settings from a vSphere Host Profile?.....	2169
How Do You Manage Compliance Checks of Hosts and Clusters to a vSphere Host Profile.....	2170
How Do You Check Compliance of Hosts or Clusters to a vSphere Host Profile?.....	2170
How Do You Schedule a Compliance Check of Hosts or Clusters to a vSphere Host Profile?.....	2170
How Do You Remediate a Host to Comply with a vSphere Host Profile Configuration?.....	2171
How Do You Configure vSphere Host Profiles.....	2172
How Do You Manage Host Profile Policies and Policy Components.....	2172
Edit a Host Profile.....	2172

How Do You Edit a vSphere Host Profile Policy.....	2172
How Do You Deactivate a Host Profile Component or Subprofile?.....	2175
How Do You Configure a Security vSphere Host Profile.....	2175
How Do You Configure a Networking vSphere Host Profile.....	2176
How Do You Customize Hosts by Using vSphere Host Profiles.....	2177
How Do You Export Host Customizations?.....	2180
How Do You Edit Host Customizations?.....	2181
Answer File Field and vSphere Host Profile Extraction.....	2181
Recommended Upgrade Workflows for vSphere Host Profiles.....	2182
How Do You Upgrade Host Profiles and vCenter Server from 6.5 to 8.0.....	2183
How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi 6.5 Hosts Only?.....	2183
How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateless ESXi 6.5 Hosts Only?.....	2184
How Do You Upgrade vSphere Host Profiles and vCenter Server from 6.7 to 8.0.....	2185
How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi6.5 and 6.7 Hosts?.....	2185
How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi6.7 Hosts Only?.....	2186
How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateless ESXi6.7 Hosts Only?.....	2186
How Do You Upgrade vSphere Host Profiles and vCenter Server from 7.0 to 8.0.....	2187
How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi Hosts of version 6.7 and 7.0?.....	2187
How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi7.0 Hosts Only?.....	2188
How Do You Upgrade Host Profiles in a vCenter Server Instance With Stateless ESXi7.0 Hosts Only?.....	2188
Using vSphere Host Profiles to Migrate an NSX Virtual Distributed Switch to a vSphere Distributed Switch.....	2189
Troubleshooting vSphere Host Profiles.....	2190
Host Customization Data Is Missing.....	2190
Reference Host is Unavailable.....	2190
Compliance Failures with vSphere Storage Host Profiles.....	2191
Compliance Failures with Serial Attached SCSI (SAS) Devices.....	2191
Compliance failures with SAN Boot LUN Devices.....	2191
Edit Settings for vSphere Host Profiles is Failing.....	2192
vSphere Host Profile without an NFS Datastore.....	2192
Upgraded Host is Not Compliant with a Legacy Host Profile.....	2192
vSphere Networking.....	2194
Introduction to vSphere Networking.....	2194
How do I Set Up Networking with vSphere Standard Switches.....	2197
Create a vSphere Standard Switch.....	2199
Port Group Configuration for Virtual Machines.....	2200
Add a Virtual Machine Port Group.....	2201
Edit a Standard Switch Port Group.....	2202
Remove a Port Group from a vSphere Standard Switch.....	2202
vSphere Standard Switch Properties.....	2202

Change the Size of the MTU on a vSphere Standard Switch.....	2203
Change the Speed of a Physical Adapter.....	2203
Assign Physical Adapters to a vSphere Standard Switch.....	2204
Topology of a vSphere Standard Switch.....	2204
How do I Set Up Networking with vSphere Distributed Switches.....	2205
What is Network Offloads Compatibility.....	2209
Create a vSphere Distributed Switch.....	2211
Upgrade a vSphere Distributed Switch to a Later Version.....	2213
Edit General and Advanced vSphere Distributed Switch Settings.....	2214
Manage Networking on Multiple Hosts on a vSphere Distributed Switch.....	2215
Tasks for Managing Host Networking on a vSphere Distributed Switch.....	2216
Add Hosts to a vSphere Distributed Switch.....	2217
Configure Physical Network Adapters on a vSphere Distributed Switch.....	2220
Migrate VMkernel Adapters to a vSphere Distributed Switch.....	2221
Create a VMkernel Adapter on a vSphere Distributed Switch.....	2222
Migrate Virtual Machine Networking to the vSphere Distributed Switch.....	2224
Remove Hosts from a vSphere Distributed Switch.....	2225
Managing Networking on Host Proxy Switches.....	2225
Migrate Network Adapters on a Host to a vSphere Distributed Switch.....	2226
Migrate a VMkernel Adapter on a Host to a vSphere Standard Switch.....	2226
Assign a Physical NIC of a Host to a vSphere Distributed Switch.....	2227
Remove a Physical NIC from a vSphere Distributed Switch.....	2227
Removing NICs from Active Virtual Machines.....	2227
Distributed Port Groups.....	2228
Add a Distributed Port Group.....	2228
Edit General Distributed Port Group Settings.....	2236
Remove a Distributed Port Group.....	2236
Working with Distributed Ports.....	2237
Monitor the State of Distributed Ports in the vSphere Client.....	2237
Configure Distributed Port Settings in vSphere Client.....	2238
Configure Virtual Machine Networking on a vSphere Distributed Switch.....	2238
Migrate Virtual Machines to or from a vSphere Distributed Switch.....	2238
Connect an Individual Virtual Machine to a Distributed Port Group.....	2239
Topology of a vSphere Distributed Switch.....	2239
View the Topology of a vSphere Distributed Switch.....	2240
View the Topology of a Host Proxy Switch.....	2241
View the Topology of Network Offloads Switch.....	2241
How do I Set Up VMkernel Networking.....	2241
VMkernel Networking Layer.....	2241
View Information About VMkernel Adapters on a Host.....	2243

Create a VMkernel Adapter on a vSphere Standard Switch.....	2244
Create a VMkernel Adapter on a Host Associated with a vSphere Distributed Switch.....	2246
Edit a VMkernel Adapter Configuration.....	2248
Override the Default Gateway of a VMkernel Adapter.....	2250
Configure the VMkernel Adapter Gateway by Using esxcli Commands.....	2250
Configure the resolv.conf File by Using esxcli Commands.....	2251
Configure the DNS Hosts File by Using ESXCLI Commands.....	2252
View TCP/IP Stack Configuration on a Host.....	2253
Change the Configuration of a TCP/IP Stack on a Host.....	2254
Explicit Congestion Notification.....	2254
Create a Custom TCP/IP Stack.....	2255
Remove a VMkernel Adapter.....	2255
LACP Support on a vSphere Distributed Switch.....	2255
Configure LACP Teaming and Failover for Distributed Port Groups.....	2257
Configure a Link Aggregation Group to Handle the Traffic for Distributed Port Groups.....	2258
Edit a Link Aggregation Group.....	2261
Limitations of the LACP Support on a vSphere Distributed Switch.....	2262
Back Up and Restore Network Configurations.....	2262
Back Up and Restore a vSphere Distributed Switch Configuration.....	2262
Export vSphere Distributed Switch Configurations.....	2262
Import a vSphere Distributed Switch Configuration.....	2263
Restore a vSphere Distributed Switch Configuration.....	2263
Export, Import, and Restore vSphere Distributed Port Group Configurations.....	2264
Export vSphere Distributed Port Group Configurations.....	2264
Import a vSphere Distributed Port Group Configuration.....	2265
Restore a vSphere Distributed Port Group Configuration.....	2265
ESXi Configuration Manager Integration.....	2265
Export Host Configuration.....	2266
Import Host Configuration.....	2266
Rollback and Recovery of the Management Network.....	2267
vSphere Networking Rollback.....	2267
Disable Network Rollback.....	2268
Disable Network Rollback by Using the vCenter Server Configuration File.....	2268
Resolve Errors in the Management Network Configuration on a vSphere Distributed Switch.....	2268
vSphere Networking Policies.....	2269
Apply Networking Policies on a vSphere Standard or Distributed Switch.....	2269
Configure Overriding Networking Policies on Port Level.....	2270
What is Teaming and Failover Policy.....	2271
Load Balancing Algorithms Available for Virtual Switches.....	2272

Configure NIC Teaming, Failover, and Load Balancing on a vSphere Standard Switch or Standard Port Group.....	2275
Configure NIC Teaming, Failover, and Load Balancing on a Distributed Port Group or Distributed Port.....	2277
What is VLAN Policy.....	2279
Configure VLAN Tagging on a Distributed Port Group or Distributed Port.....	2279
Configure VLAN Tagging on an Uplink Port Group or Uplink Port.....	2280
What is Security Policy.....	2280
Configure the Security Policy for a vSphere Standard Switch or Standard Port Group.....	2281
Configure the Security Policy for a Distributed Port Group or Distributed Port.....	2282
What is Traffic Shaping Policy.....	2283
Configure Traffic Shaping for a vSphere Standard Switch or Standard Port Group.....	2284
Edit the Traffic Shaping Policy on a Distributed Port Group or Distributed Port.....	2285
What is Resource Allocation Policy.....	2286
Edit the Resource Allocation Policy on a Distributed Port Group.....	2286
What is Monitoring Policy.....	2286
Manage NetFlow Monitoring on a Distributed Port Group or Distributed Port.....	2286
What is Traffic Filtering and Marking Policy.....	2287
Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group.....	2287
Traffic Filtering and Marking on a Distributed Port or Uplink Port.....	2292
Qualifying Traffic for Filtering and Marking.....	2298
Manage Policies for Multiple Port Groups on a vSphere Distributed Switch.....	2300
Port Blocking Policies.....	2303
Edit the Port Blocking Policy for a Distributed Port Group.....	2304
Edit the Blocking Policy for a Distributed Port or Uplink Port.....	2304
What is MAC Learning Policy.....	2304
How to Use VLANs to Isolate Network Traffic.....	2305
VLAN Configuration.....	2305
Private VLANs.....	2306
Create a Private VLAN.....	2306
Remove a Primary Private VLAN.....	2306
Remove a Secondary Private VLAN.....	2307
Manage Network Resources.....	2307
DirectPath I/O.....	2307
Enable Passthrough for a Network Device on a Host.....	2307
Configure a PCI Device on a Virtual Machine.....	2308
Hot-add and Hot-remove support for VMDirectPath I/O Devices.....	2308
Enable Hot-Add and Hot-Remove of VMDirectPath I/O Devices.....	2310
Prepare and configure VM.....	2310
Verify If the Hot Plug Operation is Successful using vSphere Client.....	2311
Verify If the Hot Plug Operation is Successful using VM.....	2311

What is Single Root I/O Virtualization (SR-IOV).....	2312
SR-IOV Support.....	2312
SR-IOV Component Architecture and Interaction.....	2313
vSphere and Virtual Function Interaction.....	2315
DirectPath I/O vs SR-IOV.....	2316
Configure a Virtual Machine to Use SR-IOV.....	2316
Networking Options for the Traffic Related to an SR-IOV Enabled Virtual Machine.....	2317
Using an SR-IOV Physical Adapter to Handle Virtual Machine Traffic.....	2318
Enable SR-IOV by Using Host Profiles or an ESXCLI Command.....	2318
What is Remote Direct Memory Access for Virtual Machines.....	2320
PVRDMA Support.....	2320
Configure an ESXi Host for PVRDMA.....	2321
Assign a PVRDMA Adapter to a Virtual Machine.....	2322
Configure a Virtual Machine to Use PVRDMA Native Endpoints.....	2323
Configure a Virtual Machine to use PVRDMA Asynchronous Mode.....	2323
Network Requirements for RDMA over Converged Ethernet.....	2324
Configure Remote Direct Memory Access Network Adapters.....	2325
View RDMA Capable Network Adapter.....	2325
Configure Remote Direct Memory Access Network Adapters.....	2326
What are Jumbo Frames.....	2327
Enable Jumbo Frames on a vSphere Distributed Switch.....	2328
Enable Jumbo Frames on a vSphere Standard Switch.....	2328
Enable Jumbo Frames for a VMkernel Adapter.....	2328
Enable Jumbo Frame Support on a Virtual Machine.....	2329
What is TCP Segmentation Offload.....	2329
Manage Software TSO in the VMkernel.....	2329
How to determine whether TSO is Supported on the Physical Network Adapters on an ESXi Host.....	2330
Manage TSO on an ESXi Host.....	2330
How to find if TSO Is Enabled on an ESXi Host.....	2331
Manage TSO on a Linux Virtual Machine.....	2331
Manage TSO on a Windows Virtual Machine.....	2331
What is Large Receive Offload.....	2332
Manage Hardware LRO for All VMXNET3 Adapters on an ESXi Host.....	2332
Manage Software LRO for All VMXNET3 Adapters on an ESXi Host.....	2332
Determine Whether LRO Is Enabled for VMXNET3 Adapters on an ESXi Host.....	2333
Change the Size of the LRO Buffer for VMXNET 3 Adapters.....	2333
Activate or Deactivate LRO for All VMkernel Adapters on an ESXi Host.....	2333
Change the Size of the LRO Buffer for VMkernel Adapters.....	2334
Manage LRO on a VMXNET3 Adapter on a Linux Virtual Machine.....	2334
Manage LRO on a VMXNET3 Adapter on a Windows Virtual Machine.....	2334

Manage LRO Globally on a Windows Virtual Machine.....	2335
NetQueue and Networking Performance.....	2335
Activate NetQueue on a Host.....	2335
Deactivate NetQueue on a Host.....	2336
vSphere Network I/O Control.....	2336
What is vSphere Network I/O Control.....	2336
Enable Network I/O Control on a vSphere Distributed Switch.....	2337
Allocate Bandwidth for System Traffic.....	2337
Bandwidth Allocation Parameters for System Traffic.....	2337
Example of Bandwidth Reservation for System Traffic.....	2338
How to Allocate Bandwidth for System Traffic.....	2338
Allocate Bandwidth for Virtual Machine Traffic.....	2339
How to Allocate Bandwidth for Virtual Machines.....	2339
Bandwidth Allocation Parameters for Virtual Machine Traffic.....	2341
Admission Control for Virtual Machine Bandwidth.....	2342
Create a Network Resource Pool.....	2343
Add a Distributed Port Group to a Network Resource Pool.....	2343
Configure Bandwidth Allocation for a Virtual Machine.....	2344
Configure Bandwidth Allocation on Multiple Virtual Machines.....	2345
Modify the Quota of a Network Resource Pool.....	2345
Remove a Distributed Port Group from a Network Resource Pool.....	2346
Delete a Network Resource Pool.....	2346
Move a Physical Adapter Out of Network I/O Control.....	2346
MAC Address Management.....	2347
MAC Address Assignment from vCenter Server.....	2347
What is VMware OUI Allocation.....	2348
What is Prefix-Based MAC Address Allocation.....	2348
What is Range-Based MAC Address Allocation.....	2348
Assigning a MAC Address.....	2349
MAC Address Generation on ESXi Hosts.....	2351
How to set a Static MAC Address to a Virtual Machine.....	2351
VMware OUI in Static MAC Addresses.....	2351
Assign a Static MAC Address.....	2352
Assign a Static MAC Address in the Virtual Machine Configuration File.....	2352
Configure vSphere for IPv6.....	2353
vSphere IPv6 Connectivity.....	2353
Deploy vSphere on IPv6.....	2354
Enable IPv6 on a vSphere Installation.....	2354
Enable IPv6 on an Upgraded vSphere Environment.....	2355
Activate or Deactivate IPv6 Support on a Host.....	2357

Set Up IPv6 on an ESXi Host.....	2357
Set Up IPv6 on vCenter Server.....	2358
Monitor Network Connection and Traffic.....	2359
How to Capture Network Packets.....	2359
Capture and Trace Network Packets by Using the pktcap-uw Utility.....	2360
pktcap-uw Command Syntax for Capturing Packets.....	2361
pktcap-uw Command Syntax for Tracing Packets.....	2363
pktcap-uw Options for Output Control.....	2363
pktcap-uw Options for Filtering Packets.....	2364
Capturing Packets by Using the pktcap-uw Utility.....	2365
How to Trace Packets using the pktcap-uw Utility.....	2373
Configure the NetFlow Settings of a vSphere Distributed Switch.....	2374
What is Port Mirroring.....	2375
What is Port Mirroring Interoperability.....	2375
Create a Port Mirroring Session.....	2377
View Port Mirroring Session Details.....	2380
Edit Port Mirroring Session Details, Sources, and Destinations.....	2380
vSphere Distributed Switch Health Check.....	2381
Manage vSphere Distributed Switch Health Check.....	2382
View vSphere Distributed Switch Health Status.....	2382
Switch Discovery Protocol.....	2383
Enable Cisco Discovery Protocol on a vSphere Distributed Switch.....	2383
Enable Link Layer Discovery Protocol on a vSphere Distributed Switch.....	2384
View Switch Information.....	2384
View the Topology of an NSX Virtual Distributed Switch.....	2384
Configure Protocol Profiles for Virtual Machine Networking.....	2385
Add a Network Protocol Profile.....	2385
Select the Network Protocol Profile Name and Network.....	2387
Specify Network Protocol Profile IPv4 Configuration.....	2387
Specify Network Protocol Profile IPv6 Configuration.....	2387
Specify Network Protocol Profile DNS and Other Configuration.....	2388
Complete the Network Protocol Profile Creation.....	2388
Associate a Port Group with a Network Protocol Profile.....	2388
Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine or vApp.....	2389
What is Multicast Filtering.....	2390
Multicast Filtering Modes.....	2390
Enable Multicast Snooping on a vSphere Distributed Switch.....	2391
Edit the Query Time Interval for Multicast Snooping.....	2392
Edit the Number of Source IP Addresses for IGMP and MLD.....	2392
What is Stateless Network Deployment.....	2392

vSphere Networking Best Practices.....	2393
Troubleshoot vSphere Networking.....	2394
Guidelines for Troubleshooting Your Implementation of vSphere.....	2394
How Do You Identify the Symptoms.....	2394
How Do You Define the Problem Space.....	2395
How Do You Test Possible Solutions.....	2395
Troubleshooting with vCenter Server Logs.....	2395
Troubleshoot MAC Address Allocation.....	2396
Duplicate MAC Addresses of Virtual Machines on the Same Network.....	2397
Attempt to Power On a Virtual Machine Fails Due to a MAC Address Conflict.....	2398
Unable to Remove a Host from a vSphere Distributed Switch.....	2399
Hosts on a vSphere Distributed Switch Lose Connectivity to vCenter Server.....	2399
Alarm for Loss of Network Redundancy on a Host.....	2400
Virtual Machines Lose Connectivity After Changing the Uplink Failover Order of a Distributed Port Group.....	2401
Unable to Add a Physical Adapter to a vSphere Distributed Switch.....	2402
Troubleshoot SR-IOV Enabled Workloads.....	2402
SR-IOV Enabled Workload Cannot Communicate After You Change Its MAC Address.....	2402
Virtual Machine That Uses an SR-IOV Virtual Function Fails to Power On Because the Host Is Out of Interrupt Vectors.....	2403
A Virtual Machine that Runs a VPN Client Causes Denial of Service for Virtual Machines on the Host or Across a vSphere HA Cluster.....	2403
Low Throughput for UDP Workloads on Windows Virtual Machines.....	2405
VMs on the Same Distributed Port Group are Unable to Communicate.....	2407
Attempt to Power On a Migrated vApp Fails Because the Associated Protocol Profile Is Missing.....	2407
Networking Configuration Operation Is Rolled Back and a Host Is Disconnected from vCenter Server.....	2408
vSphere Storage.....	2410
Updated Information.....	2410
Introduction to Storage in vSphere Environment.....	2410
Traditional Storage Virtualization Models in vSphere Environment.....	2410
Software-Defined Storage Models in vSphere Environment.....	2412
vSphere Storage APIs.....	2412
Getting Started with a Traditional Storage Model in vSphere Environment.....	2413
What Types of Physical Storage Does ESXi Support.....	2413
Local Storage.....	2413
Networked Storage.....	2414
Comparing Types of Storage.....	2417
Target and Device Representations.....	2418
How Virtual Machines Access Storage.....	2419
View Storage Adapters Available on an ESXi Host.....	2420
Characteristics of ESXi Datastores.....	2421

Display Datastore Information.....	2423
Using Persistent Memory Devices with ESXi.....	2424
Monitor PMem Datastore Statistics.....	2426
Using ESXi with a SAN.....	2427
ESXi and SAN Use Cases.....	2427
Specifics of Using SAN Storage with ESXi.....	2428
Making LUN Decisions.....	2428
Selecting Virtual Machine Locations.....	2429
Third-Party Management Applications.....	2430
SAN Storage Backup Considerations.....	2430
Using ESXi with Fibre Channel SAN.....	2431
Setting Up ESXi Fibre Channel Environment.....	2433
ESXi Fibre Channel SAN Restrictions.....	2434
Setting LUN Allocations.....	2434
Setting Fibre Channel HBAs.....	2434
Configuring Fibre Channel over Ethernet.....	2434
Setup Workflow.....	2435
Bootting ESXi from Fibre Channel SAN.....	2435
Boot from SAN Benefits.....	2435
Requirements and Considerations when Bootting from Fibre Channel SAN.....	2436
Getting Ready to Boot ESXi from SAN.....	2436
Best Practices for Fibre Channel Storage with ESXi.....	2437
Preventing Fibre Channel SAN Problems.....	2437
Deactivate Automatic ESXi Host Registration.....	2438
Optimizing Fibre Channel SAN Storage Performance.....	2438
Using Fibre Channel NPIV with vSphere Virtual Machines.....	2439
How NPIV-Based LUN Access Works.....	2439
Requirements for Using NPIV.....	2440
NPIV Capabilities and Limitations.....	2440
Configure or Modify WWN Assignments.....	2440
Using ESXi with iSCSI SAN.....	2441
Bootting from iSCSI SAN with ESXi.....	2447
General Recommendations for Boot from iSCSI SAN.....	2448
Prepare the iSCSI SAN.....	2448
Configure Independent Hardware iSCSI Adapter for SAN Boot.....	2449
Best Practices for iSCSI Storage with ESXi.....	2450
Preventing iSCSI SAN Problems.....	2450
Optimizing iSCSI SAN Storage Performance.....	2451
Checking Ethernet Switch Statistics.....	2454
Configuring iSCSI and iSER Adapters and Storage with ESXi.....	2454

Set Up Independent Hardware iSCSI Adapters with ESXi.....	2456
View Independent Hardware iSCSI Adapters.....	2456
Edit Network Settings for Hardware iSCSI.....	2457
Configure Dependent Hardware iSCSI Adapters with ESXi.....	2458
Dependent Hardware iSCSI Considerations.....	2459
View Dependent Hardware iSCSI Adapters.....	2459
Determine Association Between iSCSI and Network Adapters.....	2460
Configure the Software iSCSI Adapter with ESXi.....	2460
Add or Remove the Software iSCSI Adapter.....	2461
Configure iSER with ESXi.....	2461
Install and View an RDMA Capable Network Adapter.....	2463
Enable the VMware iSER Adapter.....	2463
Modify General Properties for iSCSI or iSER Adapters on ESXi Hosts.....	2465
Setting Up Network for iSCSI and iSER with ESXi.....	2465
Multiple Network Adapters in iSCSI or iSER Configuration.....	2466
Best Practices for Configuring Networking with Software iSCSI.....	2468
Managing iSCSI Network.....	2471
iSCSI Network Troubleshooting.....	2472
Configure Port Binding for iSCSI or iSER on ESXi.....	2472
Create a Single VMkernel Adapter for iSCSI or iSER.....	2473
Bind iSCSI or iSER Adapters to VMkernel Adapters.....	2473
Review Port Binding Details on the ESXiHost.....	2474
Using Jumbo Frames with iSCSI and iSER with ESXi.....	2475
Enable Jumbo Frames for Networking.....	2475
Enable Jumbo Frames for Independent Hardware iSCSI.....	2476
Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host.....	2476
Remove Dynamic or Static iSCSI Targets.....	2477
Configuring CHAP Parameters for iSCSI or iSER Storage Adapters on ESXi Host.....	2477
Selecting CHAP Authentication Method.....	2477
Set Up CHAP for iSCSI or iSER Storage Adapter.....	2478
Set Up CHAP for Target.....	2479
Configuring Advanced Parameters for iSCSI on ESXi Host.....	2480
Configure Advanced Parameters for iSCSI on ESXi Host.....	2481
Managing iSCSI Session on ESXi Host.....	2482
Review iSCSI Sessions.....	2482
Add iSCSI Sessions.....	2483
Remove iSCSI Sessions.....	2483
Managing ESXi Storage Devices.....	2483
Viewing Storage Devices Available to ESXi Hosts.....	2483
Display Storage Devices for an ESXi Host.....	2484

Display Storage Devices for an Adapter.....	2485
Device Sector Formats.....	2485
ESXi Storage Device Names and Identifiers.....	2487
NVMe Devices with NGUID Device Identifiers.....	2488
Rename Storage Devices.....	2489
Rescan Operations for ESXi Storage.....	2489
Perform Storage Rescan.....	2490
Perform Adapter Rescan.....	2490
Change the Number of Scanned Storage Devices.....	2490
Identifying ESXi Storage Connectivity Problems.....	2491
Detecting PDL Conditions.....	2491
Performing Planned Storage Device Removal.....	2492
Detach Storage Devices.....	2492
Recovering from PDL Conditions.....	2493
Handling Transient APD Conditions.....	2494
Verify the Connection Status of a Storage Device on ESXi Host.....	2495
Use the Locator LED on ESXi Storage Devices.....	2495
Erase ESXi Storage Devices.....	2495
Change Perennial Reservation Settings on ESXiHosts.....	2496
Using Flash Devices with ESXi Hosts.....	2497
Best Practices for Flash Devices on ESXi Hosts.....	2497
Monitor Flash Devices.....	2498
Estimate Lifetime of Flash Devices.....	2498
Keeping Flash Disks VMFS-Free.....	2499
Marking Storage Devices on ESXi Hosts.....	2499
Mark Storage Devices as Flash.....	2499
Mark Storage Devices as Local.....	2500
About Virtual Flash Resource in vSphere Environment.....	2500
Considerations for Virtual Flash Resource.....	2500
Set Up Virtual Flash Resource.....	2501
Remove Virtual Flash Resource.....	2501
Set Alarm for Virtual Flash Use.....	2502
Configure Host Cache with VMFS Datastore in vSphere Environment.....	2502
About VMware NVMe Storage.....	2502
Requirements for VMware NVMe Storage.....	2507
Configuring Lossless Ethernet for NVMe over RDMA.....	2509
Configuring NVMe over RDMA (RoCE v2) on ESXi.....	2510
View RDMA Network Adapters.....	2510
Configure VMkernel Binding for the RDMA Adapter.....	2511
Add Software NVMe over RDMA or NVMe over TCP Adapters.....	2517

Add Controllers for NVMe over Fabrics.....	2518
Configuring NVMe over TCP on ESXi.....	2519
Configure VMkernel Binding for the NVMe over TCP Adapter.....	2519
Add Software NVMe over RDMA or NVMe over TCP Adapters.....	2524
Add Controllers for NVMe over Fabrics.....	2525
Remove Software NVMe Adapters from the ESXi Host.....	2526
Working with Datastores in vSphere Storage Environment.....	2526
vSphere VMFS Datastore Concepts and Operations.....	2527
Creating vSphere Datastores.....	2532
Create a vSphere VMFS Datastore.....	2532
Create an NFS Datastore in vSphere Environment.....	2533
Create a Virtual Volumes Datastore in vSphere Environment.....	2534
Basic Administrative Operations for vSphere Datastores.....	2536
Change Datastore Name.....	2536
Unmount Datastores.....	2536
Mount Datastores.....	2537
Remove VMFS Datastores.....	2537
Managing vSphere VMFS Datastores.....	2537
vSphere VMFS Locking Mechanisms.....	2538
vSphere VMFS Datastore Copies and Datastore Resignaturing.....	2541
Increase Capacity of vSphere VMFS Datastore.....	2542
Configure Settings for Clustered Virtual Disks on the vSphere VMFS6 Datastore.....	2543
Manage Filters for LUNs and vSphere VMFS Datastores.....	2544
Set Up Dynamic Disk Mirroring for Virtual Machines on vSphere VMFS.....	2545
Setting Up ESXi Core Dump on a VMFS Datastore.....	2546
Checking Metadata Consistency with vSphere On-disk Metadata Analyzer.....	2548
Configuring Pointer Block Cache for vSphere VMFS Datastores.....	2550
NFS Datastore Concepts and Operations in vSphere Environment.....	2552
Guidelines and Requirements for NFS Storage with ESXi.....	2554
Using Kerberos for NFS 4.1 with ESXi.....	2558
Set Up NFS Storage Environment with ESXi.....	2560
Configuring the NFS Datastore.....	2560
Collecting Statistical Information for NFS Storage on ESXi Host.....	2564
Using Datastore Browser in vSphere Environment.....	2565
Navigate to the Datastore Browser.....	2565
Upload Files or Folders to Datastores.....	2566
Download Files from Datastores.....	2566
Rename Datastore Files.....	2566
Move or Copy Datastore Folders or Files.....	2567
Inflate Thin Virtual Disks.....	2567

Understanding Multipathing and Failover in the ESXi Environment.....	2568
Using Pluggable Storage Architecture and Path Management with ESXi.....	2571
VMware Native Multipathing Plug-In.....	2573
VMware NMP Flow of I/O.....	2574
Display Multipathing Modules.....	2574
Display NMP Storage Devices.....	2574
VMware Path Selection Plug-Ins and Policies.....	2575
VMware SATPs.....	2576
Display SATPs for the Host.....	2576
VMware High Performance Plug-In and Path Selection Schemes.....	2577
Enable the High-Performance Plug-In and the Path Selection Schemes.....	2579
Set Latency Sensitive Threshold.....	2580
ESXi esxcli HPP Commands.....	2581
Viewing and Managing Storage Paths on ESXi Hosts.....	2582
View Storage Device Paths.....	2583
View Datastore Paths.....	2583
Change the Path Selection Policy.....	2585
Change Default Parameters for Latency Round Robin.....	2586
Deactivate Storage Paths.....	2587
Using Claim Rules to Control ESXi Multipathing Modules.....	2587
Multipathing Considerations.....	2587
List Multipathing Claim Rules for the Host.....	2588
Add Multipathing Claim Rules.....	2589
Delete Multipathing Claim Rules.....	2592
Mask Paths.....	2592
Unmask Paths.....	2593
Define NMP SATP Rules.....	2594
Scheduling Queues for Virtual Machine I/Os in vSphere Environment.....	2595
Edit Per File I/O Scheduling in the vSphere Client.....	2595
Use esxcli Commands to Edit Per File I/O Scheduling.....	2595
Raw Device Mapping in vSphere.....	2596
Introduction to Raw Device Mapping in vSphere.....	2596
Benefits of Raw Device Mapping.....	2597
RDM Considerations and Limitations.....	2599
Characteristics of Raw Device Mapping in vSphere.....	2599
RDM Virtual and Physical Compatibility Modes.....	2599
Dynamic Name Resolution.....	2599
Raw Device Mapping with Virtual Machine Clusters.....	2600
Comparing Available SCSI Device Access Modes.....	2600
Create vSphere Virtual Machines with RDMs.....	2601

Deactivate SCSI INQUIRY Cache for Improved Performance in vSphere.....	2602
Storage Policy Based Management in vSphere.....	2603
Default Storage Policies in vSphere.....	2604
Change the Default Storage Policy for a Datastore.....	2605
Workflow for vSphere VM Storage Policies.....	2605
Populating the VM Storage Policies Interface in vSphere.....	2606
Use Storage Providers to Populate the VM Storage Policies Interface.....	2606
Assign Tags to Datastores.....	2607
About Rules and Rule Sets in vSphere Storage Policies.....	2608
Creating and Managing vSphere VM Storage Policies.....	2610
Create a VM Storage Policy for VMFS Datastore.....	2610
Create a VM Storage Policy for Host-Based Data Services.....	2611
Create a VM Storage Policy for Virtual Volumes.....	2612
Create a VM Storage Policy for Tag-Based Placement.....	2614
Edit or Clone a VM Storage Policy.....	2614
About vSphere Storage Policy Components.....	2615
Create Storage Policy Components.....	2616
Edit or Clone Storage Policy Components.....	2617
vSphere Storage Policies and Virtual Machines.....	2617
Assign Storage Policies to Virtual Machines.....	2617
Change Storage Policy Assignment for Virtual Machine Files and Disks.....	2618
Check Compliance for a VM Storage Policy.....	2619
Find Compatible Storage Resource for Noncompliant Virtual Machine.....	2620
Reapply Virtual Machine Storage Policy.....	2620
Using Storage Providers in vSphere.....	2621
Managing Storage Providers.....	2623
Register Storage Providers.....	2623
View Storage Provider Information.....	2624
Re-authenticate Host VASA Client.....	2624
Manage Storage Providers.....	2625
Working with VMware vSphere Virtual Volumes.....	2627
vSphere Virtual Volumes Concepts and Architecture.....	2627
Virtual Volume Objects.....	2627
Virtual Volumes Storage Providers.....	2628
Virtual Volumes Storage Containers.....	2629
Static Protocol Endpoints.....	2629
Binding and Unbinding Virtual Volumes.....	2630
Virtual Volumes Datastores.....	2630
Virtual Volumes and VM Storage Policies.....	2631
Virtual Volumes and Storage Protocols.....	2631

Virtual Volumes Architecture.....	2633
vSphere Virtual Volumes Security.....	2635
vSphere Virtual Volumes Snapshots.....	2636
vSphere Virtual Volumes and Shared Virtual Disks.....	2637
Before You Enable vSphere Virtual Volumes.....	2637
Synchronize vSphere Storage Environment with a Network Time Server.....	2638
Configure vSphere Virtual Volumes.....	2638
Register Storage Providers for Virtual Volumes.....	2638
Create a Virtual Volumes Datastore in vSphere Environment.....	2642
Provision VMs on vSphere Virtual Volumes Datastores.....	2644
Review and Manage Protocol Endpoints in vSphere.....	2644
Review and Manage Static Protocol Endpoints.....	2644
Change the Path Selection Policy for a SCSI Protocol Endpoint.....	2647
Review NVMe Protocol Endpoints.....	2648
Manage Storage Providers for vSphere Virtual Volumes.....	2650
Re-authenticate Host VASA Client.....	2651
NVMe and Virtual Volumes in vSphere.....	2652
Configure NVMe with Virtual Volumes.....	2654
vSphere Virtual Volumes and Replication.....	2660
Requirements for Replication with Virtual Volumes.....	2660
Virtual Volumes and Replication Groups.....	2661
Virtual Volumes and Fault Domains.....	2661
Virtual Volumes Replication Workflow.....	2663
Replication Guidelines and Considerations.....	2663
Best Practices for Working with vSphere Virtual Volumes.....	2664
Guidelines and Limitations when Using Virtual Volumes.....	2664
Best Practices for Storage Container Provisioning.....	2665
Best Practices for Virtual Volumes Performance.....	2666
Using Stretched Storage Clustering with Virtual Volumes.....	2667
Using the esxcli storage vvol Command to Troubleshoot vSphere Virtual Volumes.....	2667
Virtual Volumes and esxcli Commands.....	2667
Collecting Statistical Information for Virtual Volumes.....	2668
Virtual Volumes Datastore Is Inaccessible.....	2669
Failures When Migrating VMs or Deploying VM OVFes to Virtual Volumes Datastores.....	2669
Filtering Virtual Machine I/O in vSphere.....	2669
Deploying and Managing I/O Filters in vSphere.....	2671
Configure I/O Filters in the vSphere Environment.....	2672
Handling I/O Filter Installation Failures.....	2673
Uninstall I/O Filters from a Cluster.....	2673
Upgrade I/O Filters in a Cluster.....	2674

Activate I/O Filter Data Services on Virtual Disks in vSphere.....	2674
Assign the I/O Filter Policy to Virtual Machines.....	2674
Migrating Virtual Machines with I/O Filters.....	2675
Storage Hardware Acceleration in vSphere.....	2676
vSphere Hardware Acceleration for Block Storage Devices.....	2677
Deactivate Hardware Acceleration for Block Storage Devices.....	2678
Managing Hardware Acceleration on Block Storage Devices.....	2678
Configure XCOPY Parameters.....	2681
vSphere Hardware Acceleration on NAS Devices.....	2683
Enable NAS Native Snapshots on Virtual Machines.....	2684
Storage Provisioning and Space Reclamation in vSphere.....	2684
Virtual Disk Thin Provisioning with vSphere Storage.....	2685
About Virtual Disk Provisioning Policies.....	2685
Create Thin Provisioned Virtual Disks.....	2686
View Virtual Machine Storage Resources.....	2687
Determine the Disk Format of a Virtual Machine.....	2687
Inflate Thin Virtual Disks.....	2687
Handling Datastore Over-Subscription.....	2688
ESXi and Array Thin Provisioning.....	2688
Monitoring Space Use.....	2688
Identify Thin-Provisioned Storage Devices.....	2689
Storage Space Reclamation in vSphere.....	2689
Space Reclamation on vSphere VMFS Datastores.....	2690
Reclaim Space on the vSphere Virtual Volumes Datastores.....	2696
Space Reclamation Requests from Guest Operating Systems.....	2697
Using vmkfstools in vSphere.....	2698
File System Options of vSphere vmkfstools Command.....	2699
Listing Attributes of a VMFS Datastore.....	2699
Creating a VMFS Datastore or a Scratch Partition.....	2700
Adding an Extent to a VMFS Datastore.....	2700
Expanding a VMFS Datastore.....	2701
Virtual Disk Options.....	2701
Supported Disk Formats.....	2701
Creating a Virtual Disk.....	2702
Initializing a Virtual Disk.....	2702
Inflating a Thin Virtual Disk.....	2702
Converting a Zeroedthick Virtual Disk to an Eagerzeroedthick Disk.....	2703
Removing Zeroed Blocks.....	2703
Deleting a Virtual Disk.....	2703
Renaming a Virtual Disk.....	2703

Cloning or Converting a Virtual Disk or RDM.....	2703
Extending a Virtual Disk.....	2704
Upgrading Virtual Disks.....	2704
Creating a Virtual Compatibility Mode Raw Device Mapping.....	2705
Creating a Physical Compatibility Mode Raw Device Mapping.....	2705
Listing Attributes of an RDM.....	2705
Displaying Virtual Disk Geometry.....	2706
Checking and Repairing Virtual Disks.....	2706
Checking Disk Chain for Consistency.....	2706
Storage Device Options of vSphere vmkfstools Command.....	2706
Managing SCSI Reservations of LUNs.....	2706
Breaking Device Locks.....	2707
Getting Started with Cloud Native Storage in vSphere.....	2707
Cloud Native Storage Concepts and Terminology in vSphere.....	2707
Cloud Native Storage Components.....	2710
Using vSAN File Service to Provision vSphere File Volumes.....	2712
vSphere Cloud Native Storage Users.....	2714
Cloud Native Storage for vSphere Administrators.....	2715
Requirements of vSphere Cloud Native Storage.....	2715
vSphere Cloud Native Storage Roles and Privileges.....	2717
Create a Storage Policy for Kubernetes in vSphere.....	2718
Configure Kubernetes Cluster in vSphere Virtual Machines.....	2720
Monitor Container Volumes Across Kubernetes Clusters in vSphere.....	2720
Use Encryption with vSphere Cloud Native Storage.....	2721
Migrating Container Volumes in vSphere.....	2722
vSphere Security.....	2726
Updated Information.....	2727
Security in the vSphere Environment.....	2728
Securing the ESXi Hypervisor.....	2728
Securing vCenter Server Systems and Associated Services.....	2729
Securing Virtual Machines.....	2730
Securing the Virtual Networking Layer.....	2731
Securing Passwords in Your vSphere Environment.....	2733
vCenter Server and ESXi Security Best Practices and Resources.....	2734
vSphere Permissions and User Management Tasks.....	2735
Understanding Authorization in vSphere.....	2735
Hierarchical Inheritance of Permissions in vSphere.....	2738
How Do Multiple Permission Settings Work in vSphere.....	2740
Example 1: Permission Inheritance from Multiple Groups.....	2741
Example 2: Child Permissions Overriding Parent Permissions.....	2741

Example 3: User Role Overriding Group Role.....	2742
Managing Permissions for vCenter Server Components.....	2743
Add a Permission to an Inventory Object.....	2743
Change or Remove Permissions on an Inventory Object.....	2744
Change vCenter Server User Validation Settings.....	2745
Using vCenter Server Global Permissions.....	2745
Add a Global Permission.....	2746
vCenter Server Permissions on Tag Objects.....	2746
Using vCenter Server Roles to Assign Privileges.....	2748
Create a vCenter Server Custom Role.....	2750
Using Privilege Recorder.....	2751
Enable the Privilege Recorder.....	2751
Best Practices for vCenter Server Roles and Permissions.....	2751
Required vCenter Server Privileges for Common Tasks.....	2752
Securing ESXi Hosts.....	2755
General ESXi Security Recommendations.....	2755
ESXi Advanced System Settings.....	2756
Configure ESXi Hosts with Host Profiles.....	2759
Use Scripts to Manage ESXi Host Configuration Settings.....	2760
ESXi Passwords and Account Lockout.....	2761
ESXi Cryptographic Key Generation.....	2762
SSH Security in ESXi.....	2763
PCI and PCIe Devices and ESXi.....	2765
Deactivate the vSphere Managed Object Browser.....	2765
ESXi Networking Security Recommendations.....	2765
Modifying ESXi Web Proxy Settings.....	2766
vSphere Auto Deploy Security Considerations.....	2766
Control Access for CIM-Based Hardware Monitoring Tools.....	2767
vSphere Distributed Services Engine Security Best Practices.....	2768
Controlling ESXi Entropy.....	2768
Managing Certificates for ESXi Hosts.....	2769
ESXi Host Upgrades and Certificates.....	2771
ESXi Certificate Mode Switch Workflows.....	2772
ESXi Certificate Default Settings.....	2773
View Certificate Expiration Information for ESXi Hosts.....	2775
Renew or Refresh ESXi Certificates.....	2776
Change the ESXi Certificate Mode.....	2777
Replacing the Default ESXi Certificate with a Custom Certificate.....	2777
Make Auto Deploy a Subordinate Certificate Authority.....	2782
Use Custom Certificates with Auto Deploy.....	2784

Restore ESXi Certificate and Key Files When Certificate Replacement Fails.....	2787
Customizing ESXi Host Security.....	2787
Configuring the ESXi Firewall.....	2787
Activate or Deactivate an ESXi Service.....	2790
Configuring and Managing Lockdown Mode on ESXi Hosts.....	2792
Using vSphere Installation Bundles to Perform Secure Updates.....	2796
Manage the Acceptance Levels of ESXi Hosts and vSphere Installation Bundles.....	2797
Assigning Privileges for ESXi Hosts.....	2798
Using Active Directory to Manage ESXi Users.....	2800
Configure an ESXi Host to Use Active Directory.....	2800
Add an ESXi Host to a Directory Service Domain.....	2801
View Directory Service Settings for an ESXi Host.....	2802
Using vSphere Authentication Proxy.....	2802
Start the vSphere Authentication Proxy Service.....	2803
Add a Domain to vSphere Authentication Proxy Using the vSphere Client.....	2804
Add a Domain to vSphere Authentication Proxy Using the camconfig Command.....	2804
Use vSphere Authentication Proxy to Add a Host to a Domain.....	2804
Activate Client Authentication for vSphere Authentication Proxy.....	2805
Import the vSphere Authentication Proxy Certificate to ESXi Host.....	2806
Generate a New Certificate for vSphere Authentication Proxy.....	2806
Set Up vSphere Authentication Proxy to Use Custom Certificates.....	2807
Configuring and Managing Smart Card Authentication for ESXi.....	2808
Activate Smart Card Authentication.....	2808
Deactivate Smart Card Authentication.....	2809
Authenticating With User Name and Password in Case of Connectivity Problems.....	2809
Using Smart Card Authentication in Lockdown Mode.....	2809
Using the ESXi Shell.....	2809
Set Idle Timeout for the ESXi Shell Using the vSphere Client.....	2810
Set Availability Timeout for the ESXi Shell Using the vSphere Client.....	2811
Set Availability Timeout or Idle Timeout for the ESXi Shell Using the DCUI.....	2811
Activate Access to the ESXi Shell Using the vSphere Client.....	2812
Activate Access to the ESXi Shell Using the DCUI.....	2812
Log in to the ESXi Shell for Troubleshooting.....	2813
UEFI Secure Boot for ESXi Hosts.....	2813
Securing ESXi Hosts with Trusted Platform Module.....	2814
View ESXi Host Attestation Status.....	2815
Troubleshoot ESXi Host Attestation Problems.....	2816
ESXi Log Files.....	2816
Configure Syslog on ESXi Hosts.....	2816
ESXi Syslog Options.....	2816

ESXi Log File Locations.....	2820
Securing Fault Tolerance Logging Traffic.....	2821
Activate Fault Tolerance Encryption.....	2822
Managing ESXi Audit Records.....	2822
Securing the ESXi Configuration.....	2823
Manage a Secure ESXi Configuration.....	2825
Deactivate the execInstalledOnly Internal Runtime Option.....	2830
Securing vCenter Server Systems.....	2831
Best Practices for vCenter Server Access Control.....	2831
Set the vCenter Server Password Policy.....	2832
Removing Expired or Revoked Certificates and Logs from Failed Installations.....	2832
Limiting vCenter Server Network Connectivity.....	2833
Evaluate the Use of Linux Clients with CLIs and SDKs.....	2833
Examine vSphere Client Plug-Ins.....	2833
vCenter Server Security Best Practices.....	2834
vCenter Password Requirements and Lockout Behavior.....	2834
Verify Thumbprints for Legacy ESXi Hosts.....	2836
Required Ports for vCenter Server.....	2836
Securing Virtual Machines.....	2836
Activate or Deactivate UEFI Secure Boot for a Virtual Machine.....	2837
Virtual Machine Security Best Practices.....	2838
General Virtual Machine Protection.....	2838
Use Templates to Deploy Virtual Machines.....	2839
Minimize Use of the Virtual Machine Console.....	2839
Prevent Virtual Machines from Taking Over Resources.....	2839
Deactivate Unnecessary Functions Inside Virtual Machines.....	2840
Securing Virtual Machines with Intel Software Guard Extensions.....	2843
Getting Started with vSGX.....	2844
Enable vSGX on a Virtual Machine.....	2845
Enable vSGX on an Existing Virtual Machine.....	2845
Remove vSGX from a Virtual Machine.....	2846
Securing Virtual Machines with AMD Secure Encrypted Virtualization-Encrypted State.....	2846
vSphere and AMD Secure Encrypted Virtualization-Encrypted State.....	2846
Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the vSphere Client.....	2847
Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the Command Line.....	2848
Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the vSphere Client.....	2849
Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the Command Line.....	2849

Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the vSphere Client.....	2850
Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the Command Line.....	2850
Virtual Machine Encryption.....	2851
Comparison of vSphere Key Providers.....	2851
How vSphere Virtual Machine Encryption Protects Your Environment.....	2853
vSphere Virtual Machine Encryption Components.....	2857
Encryption Process Flow.....	2859
Virtual Disk Encryption.....	2861
Virtual Machine Encryption Errors.....	2862
Prerequisites and Required Privileges for Virtual Machine Encryption Tasks.....	2862
Encrypted vSphere vMotion.....	2864
Virtual Machine Encryption Best Practices.....	2866
Virtual Machine Encryption Caveats.....	2869
Virtual Machine Encryption Interoperability.....	2869
vSphere Key Persistence on ESXi Hosts.....	2871
Configuring and Managing a Standard Key Provider.....	2872
What Is a Standard Key Provider.....	2873
Setting Up the Standard Key Provider.....	2873
Add a Standard Key Provider Using the vSphere Client.....	2873
Establish a Standard Key Provider Trusted Connection by Exchanging Certificates.....	2874
Finish the Trust Setup for a Standard Key Provider.....	2877
Set Up Separate Key Providers for Different Users.....	2878
Delete a Standard Key Provider.....	2878
Configuring and Managing vSphere Native Key Provider.....	2879
vSphere Native Key Provider Overview.....	2879
vSphere Native Key Provider Process Flows.....	2881
Configure a vSphere Native Key Provider.....	2882
Back up a vSphere Native Key Provider.....	2883
Recovering a vSphere Native Key Provider.....	2884
Restore a vSphere Native Key Provider Using the vSphere Client.....	2884
Update a vSphere Native Key Provider.....	2884
Delete a vSphere Native Key Provider.....	2885
vSphere Trust Authority.....	2886
vSphere Trust Authority Concepts and Features.....	2886
How vSphere Trust Authority Protects Your Environment.....	2886
vSphere Trust Authority Trusted Infrastructure.....	2889
vSphere Trust Authority Process Flows.....	2891
vSphere Trust Authority Topology.....	2895

Prerequisites and Required Privileges for vSphere Trust Authority.....	2895
vSphere Trust Authority Best Practices, Caveats, and Interoperability.....	2897
vSphere Trust Authority Life Cycle.....	2898
Configuring vSphere Trust Authority.....	2900
Set up Your Workstation to Configure vSphere Trust Authority.....	2902
Enable the Trust Authority Administrator.....	2902
Enable the Trust Authority State.....	2903
Collect Information About ESXi Hosts and vCenter Server to Be Trusted.....	2904
Import the Trusted Host Information to the Trust Authority Cluster.....	2911
Create the Key Provider on the Trust Authority Cluster.....	2913
Export the Trust Authority Cluster Information.....	2920
Import the Trust Authority Cluster Information to the Trusted Hosts.....	2921
Configure the Trusted Key Provider for Trusted Hosts Using the vSphere Client.....	2925
Configure the Trusted Key Provider for Trusted Hosts Using the Command Line.....	2925
Managing vSphere Trust Authority in Your vSphere Environment.....	2926
Start, Stop, and Restart vSphere Trust Authority Services.....	2927
View the Trust Authority Hosts.....	2927
View the vSphere Trust Authority Cluster State.....	2927
Restart the Trusted Host Service.....	2927
Adding and Removing vSphere Trust Authority Hosts.....	2928
Add a Host to a Trusted Cluster Using the vSphere Client.....	2928
Add a Host to a Trusted Cluster Using the Command Line.....	2928
Decommission Trusted Hosts from a Trusted Cluster.....	2929
Backing Up the vSphere Trust Authority Configuration.....	2930
Change the Primary Key of a Trusted Key Provider.....	2930
Trusted Host Attestation Reporting.....	2931
Checking and Remediating Trusted Cluster Health.....	2933
Using Encryption in Your vSphere Environment.....	2935
Create an Encryption Storage Policy.....	2935
Activate Host Encryption Mode Explicitly.....	2936
Deactivate Host Encryption Mode Using the API.....	2936
Create an Encrypted Virtual Machine.....	2937
Clone an Encrypted Virtual Machine.....	2938
Encrypt an Existing Virtual Machine or Virtual Disk.....	2941
Decrypt an Encrypted Virtual Machine or Virtual Disk.....	2942
Change the Encryption Policy for Virtual Disks.....	2943
Resolve Missing Encryption Key Issues.....	2943
Unlock Locked Virtual Machines.....	2945
Resolve ESXi Host Encryption Mode Issues.....	2945
Re-Activate ESXi Host Encryption Mode.....	2945

Set Key Server Certificate Expiration Threshold.....	2946
vSphere Virtual Machine Encryption and Core Dumps.....	2946
Collect a vm-support Package for an ESXi Host That Uses Encryption.....	2947
Decrypt or Re-Encrypt an Encrypted Core Dump.....	2948
Activate and Deactivate Key Persistence on an ESXi Host.....	2949
Rekey an Encrypted Virtual Machine Using the vSphere Client.....	2949
Rekey an Encrypted Virtual Machine Using the CLI.....	2950
Set the Default Key Provider Using the vSphere Client.....	2951
Set the Default Key Provider Using the Command Line.....	2951
Securing Virtual Machines with Virtual Trusted Platform Module.....	2952
What Is a Virtual Trusted Platform Module.....	2952
Create a Virtual Machine with a Virtual Trusted Platform Module.....	2953
Add Virtual Trusted Platform Module to an Existing Virtual Machine.....	2954
Remove Virtual Trusted Platform Module from a Virtual Machine.....	2954
Identify Virtual Trusted Platform Module Enabled Virtual Machines.....	2955
View Virtual Trusted Platform Module Device Certificates.....	2955
Export and Replace Virtual Trusted Platform Module Device Certificates.....	2956
Securing Windows Guest Operating Systems with Virtualization-based Security.....	2956
vSphere Virtualization-based Security Best Practices.....	2957
Activate Virtualization-based Security on a Virtual Machine.....	2958
Activate Virtualization-based Security on an Existing Virtual Machine.....	2959
Activate Virtualization-based Security on the Guest Operating System.....	2959
Deactivate Virtualization-based Security.....	2960
Identify VBS-Enabled Virtual Machines.....	2960
Securing vSphere Networking.....	2960
Securing the Network with Firewalls.....	2961
Firewalls for Configurations with vCenter Server.....	2962
Connecting to vCenter Server Through a Firewall.....	2962
Connecting ESXi Hosts Through Firewalls.....	2963
Firewalls for Configurations Without vCenter Server.....	2963
Connecting to the Virtual Machine Console Through a Firewall.....	2963
Secure the Physical Switch on ESXi Hosts.....	2964
Securing Standard Switch Ports with Security Policies.....	2964
Securing vSphere Standard Switches.....	2965
MAC Address Changes.....	2966
Forged Transmits.....	2966
Promiscuous Mode Operation.....	2966
Standard Switch Protection and VLANs.....	2967
Secure vSphere Distributed Switches and Distributed Port Groups.....	2968
Securing Virtual Machines with VLANs.....	2969

Security Considerations for VLANs.....	2970
Secure VLANs.....	2971
Creating Multiple Networks Within a Single ESXi Host.....	2971
Using Internet Protocol Security on ESXi Hosts.....	2973
List Available Security Associations on ESXi Hosts.....	2973
Add an IPsec Security Association to an ESXi Host.....	2973
Remove an IPsec Security Association from an ESXi Host.....	2974
List Available IPsec Security Policies on an ESXi Host.....	2974
Create an IPsec Security Policy on an ESXi Host.....	2975
Remove an IPsec Security Policy from an ESXi Host.....	2976
Ensure Proper SNMP Configuration on ESXi Hosts.....	2976
vSphere Networking Security Best Practices.....	2976
General vSphere Networking Security Recommendations.....	2976
Labeling vSphere Networking Components.....	2977
Document and Check the vSphere VLAN Environment.....	2978
Adopting Network Isolation Practices in vSphere.....	2978
Use Virtual Switches with the vSphere Network Appliance API Only If Required.....	2979
Best Practices Involving Multiple vSphere Components.....	2980
Synchronizing Clocks on the vSphere Network.....	2980
Synchronize ESXi Clocks with a Network Time Server.....	2981
Configuring Time Synchronization Settings in vCenter Server.....	2981
Storage Security Best Practices.....	2983
Securing iSCSI Storage.....	2983
Masking and Zoning SAN Resources.....	2984
Using Kerberos for NFS 4.1.....	2984
Verify That Sending Host Performance Data to Guests Is Deactivated.....	2985
Setting Timeouts for the ESXi Shell and the vSphere Client.....	2986
vSphere TLS Configuration.....	2986
Managing vSphere TLS.....	2989
View the TLS Profile of an ESXi Host Using the vSphere Client.....	2990
View the TLS Profile of an ESXi Host Using the CLI.....	2990
Change the TLS Profile of an ESXi Host Using the vSphere Client.....	2990
Change the TLS Profile of an ESXi Host Using the CLI.....	2991
Edit the Parameters in the MANUAL TLS Profile Using the CLI.....	2991
Manage the TLS Profile of a vCenter Server Host.....	2993
Defined Privileges.....	2993
Alarms Privileges.....	2994
Auto Deploy and Image Profile Privileges.....	2994
Certificates Privileges.....	2995
Certificate Authority Privileges.....	2995

Certificate Management Privileges.....	2996
Cns Privileges.....	2996
Compute Policy Privileges.....	2997
Content Library Privileges.....	2997
Cryptographic Operations Privileges.....	3000
dvPort Group Privileges.....	3003
Distributed Switch Privileges.....	3003
Datacenter Privileges.....	3004
Datastore Privileges.....	3005
Datastore Cluster Privileges.....	3006
ESX Agent Manager Privileges.....	3007
Extension Privileges.....	3007
External Stats Provider Privileges.....	3008
Folder Privileges.....	3008
Global Privileges.....	3008
Interact with the Guest Data Publisher Privileges.....	3010
Hybrid Linked Mode Privileges.....	3010
Health Update Provider Privileges.....	3010
Host CIM Privileges.....	3010
Host Configuration Privileges.....	3011
Host Entropy Pool Privileges.....	3012
Host Intel Software Guard Extensions Privileges.....	3013
Host Inventory Privileges.....	3013
Host Local Operations Privileges.....	3014
Host Statistics Privileges.....	3014
Host Trusted Platform Module Privileges.....	3015
Host vSphere Replication Privileges.....	3015
Host Profile Privileges.....	3015
vCenter Server Profiles Privileges.....	3016
vSphere Namespaces Privileges.....	3016
Network Privileges.....	3017
NSX Privileges.....	3017
VMware Observability Privileges.....	3018
OvfManager Privileges.....	3018
Interact with Partner Rest Daemons Privileges.....	3018
Performance Privileges.....	3018
Plug-in Privileges.....	3019
Replication as a Service Privileges.....	3019
Permissions Privileges.....	3019
VM Storage Policies Privileges.....	3020

Resource Privileges.....	3020
Scheduled Task Privileges.....	3021
Sessions Privileges.....	3022
Storage Views Privileges.....	3022
Supervisor Services Privileges.....	3023
Tasks Privileges.....	3023
Tenant Management Privileges.....	3024
Transfer Service Privileges.....	3024
VcTrusts/VcIdentity Privileges.....	3024
Trusted Infrastructure Administrator Privileges.....	3025
vApp Privileges.....	3026
VcIdentityProviders Privileges.....	3027
VMware vSphere Lifecycle Manager Configuration Privileges.....	3028
VMware vSphere Lifecycle Manager Manager Desired Configuration Management Privileges.....	3028
VMware vSphere Lifecycle Manager ESXi Health Perspectives Privileges.....	3029
VMware vSphere Lifecycle Manager Depots Privileges.....	3029
VMware vSphere Lifecycle Manager General Privileges.....	3030
VMware vSphere Lifecycle Manager Hardware Compatibility Privileges.....	3030
VMware vSphere Lifecycle Manager Image Privileges.....	3031
VMware vSphere Lifecycle Manager Image Remediation Privileges.....	3031
VMware vSphere Lifecycle Manager Settings Privileges.....	3032
VMware vSphere Lifecycle Manager Manage Baseline Privileges.....	3032
VMware vSphere Lifecycle Manager Manage Patches and Upgrades Privileges.....	3033
VMware vSphere Lifecycle Manager Upload File Privileges.....	3033
Virtual Machine Change Configuration Privileges.....	3034
Virtual Machine Guest Operations Privileges.....	3036
Virtual Machine Interaction Privileges.....	3037
Virtual Machine Edit Inventory Privileges.....	3040
Virtual Machine Provisioning Privileges.....	3041
Virtual Machine Service Configuration Privileges.....	3042
Virtual Machine Snapshot Management Privileges.....	3043
Virtual Machine vSphere Replication Privileges.....	3043
Virtual Machine Classes Privileges.....	3044
vSAN Privileges.....	3044
vSAN Stats Privileges.....	3044
vSphere Zones Privileges.....	3045
vService Privileges.....	3045
vSphere Tagging Privileges.....	3046
vSphere Client Privileges.....	3046
vSphere Data Protection Privileges.....	3046

vSphere Stats Privileges.....	3047
vSphere Hardening and Compliance.....	3047
Security Versus Compliance in the vSphere Environment.....	3047
vSphere Security Controls Reference.....	3049
vSphere System Design Security Controls Reference.....	3051
vSphere Hardware Security Controls Reference.....	3055
ESXi Security Controls Reference.....	3059
vCenter Server Security Controls Reference.....	3093
Virtual Machine Security Controls Reference.....	3111
Guest Operating System Security Controls.....	3119
vSAN Security Controls Reference.....	3126
About the National Institute of Standards and Technology.....	3128
About DISA STIGs.....	3129
About NERC CIP.....	3129
About VMware Security Development Lifecycle.....	3129
Audit Logging in vSphere.....	3130
Single Sign-On Audit Events.....	3130
Understanding Security and Compliance Next Steps.....	3131
vCenter Server and FIPS.....	3131
FIPS Modules Used in ESXi.....	3132
Activate and Deactivate FIPS on the vCenter Server Appliance.....	3132
Considerations When Using FIPS.....	3133
vSphere Resource Management.....	3134
Configure Your Resource Allocation Settings in vSphere.....	3134
CPU Virtualization in vSphere.....	3138
Administering CPU Resources in vSphere.....	3139
Hyperthreading with vSphere.....	3140
Using CPU Affinity with vSphere.....	3142
Host Power Management Policies in ESXi.....	3143
Memory Virtualization with vSphere.....	3145
Administer Your Memory Resources with vSphere.....	3148
Using Swap Files with vSphere.....	3155
Using Persistent Memory with vSphere.....	3158
Configure Virtual Graphics on vSphere.....	3162
Manage Storage I/O Resources with vSphere.....	3165
Managing Resource Pools with vSphere.....	3170
vSphere Cluster Services.....	3177
Embedded vCLS.....	3177
vSphere DRS and vCLS VMs.....	3179
Datastore selection for External vCLS.....	3180

External vCLS Datastore Placement.....	3180
Monitoring vSphere Cluster Services.....	3180
Maintaining Health of vSphere Cluster Services.....	3181
Putting a Cluster in Retreat Mode.....	3182
Retrieving Password for External vCLS.....	3183
vCLS VM Anti-Affinity Policies.....	3183
Create or Delete a vCLS VM Anti-Affinity Policy.....	3183
Create a vSphere DRS Cluster.....	3184
Admission Control and Initial Placement.....	3184
Single Virtual Machine Power On.....	3185
Group Power-on.....	3185
Virtual Machine Migration.....	3185
DRS Migration Threshold.....	3186
Migration Recommendations.....	3187
DRS Cluster Requirements.....	3187
Shared Storage Requirements.....	3187
Shared VMFS Volume Requirements.....	3187
Processor Compatibility Requirements.....	3188
vMotion Requirements for DRS Clusters.....	3188
Configuring DRS with Virtual Flash.....	3189
Create a Cluster.....	3189
Edit Cluster Settings.....	3190
Set a Custom Automation Level for a Virtual Machine.....	3191
Deactivate DRS.....	3192
Restore a Resource Pool Tree.....	3192
DRS Awareness of vSAN Stretched Cluster.....	3193
DRS Placement of vGPUs.....	3193
DRS Overhead Memory Management for VMs.....	3194
vSphere DRS Maintenance Mode Functionality with ROBO Enterprise License.....	3194
Use DRS Clusters to Manage Resources with vSphere.....	3196
Adding Hosts to a Cluster.....	3196
Adding Virtual Machines to a Cluster.....	3197
Removing Virtual Machines from a Cluster.....	3198
Removing a Host from a Cluster.....	3198
Place a Host in Maintenance Mode.....	3199
Remove a Host from a Cluster.....	3199
Using Standby Mode.....	3200
Configuring VM Best Effort Restart Policy.....	3200
DRS Cluster Validity.....	3201
Valid DRS Clusters.....	3201

Overcommitted DRS Clusters.....	3203
Invalid DRS Clusters.....	3204
Managing Power Resources.....	3205
Configure IPMI or iLO Settings for vSphere DPM.....	3206
Test Wake-on-LAN for vSphere DPM.....	3206
Activating vSphere DPM for a DRS Cluster.....	3207
Monitoring vSphere DPM.....	3208
VM Static Power.....	3209
Using Affinity Rules with vSphere DRS.....	3209
Create a VM-Host Group.....	3211
Create a Virtual Machine Group.....	3212
Create a VM-VM Affinity Rule.....	3212
VM-VM Affinity Rule Conflicts.....	3212
Create a VM-Host Affinity Rule.....	3212
Using VM-Host Affinity Rules.....	3213
Using Affinity Rules without vSphere DRS.....	3214
Creating a Datastore Cluster in vSphere.....	3215
Initial Placement and Ongoing Balancing.....	3215
Storage Migration Recommendations.....	3216
Create a Datastore Cluster.....	3216
Activate and Deactivate Storage DRS.....	3216
Set the Automation Level for Datastore Clusters.....	3217
Setting the Aggressiveness Level for Storage DRS.....	3217
Datastore Cluster Requirements.....	3218
Adding and Removing Datastores from a Datastore Cluster.....	3219
Use Datastore Clusters to Manage Storage Resources in vSphere.....	3219
Using Storage DRS Maintenance Mode.....	3219
Applying Storage DRS Recommendations.....	3220
Change Storage DRS Automation Level for a Virtual Machine.....	3221
Set Up Off-Hours Scheduling for Storage DRS.....	3222
Storage DRS Anti-Affinity Rules.....	3222
Clear Storage DRS Statistics.....	3225
Storage vMotion Compatibility with Datastore Clusters.....	3226
Using NUMA Systems with ESXi.....	3226
What is NUMA?.....	3226
Challenges for Operating Systems.....	3227
How ESXi NUMA Scheduling Works.....	3227
VMware NUMA Optimization Algorithms and Settings.....	3228
Home Nodes and Initial Placement.....	3228
Dynamic Load Balancing and Page Migration.....	3228

Transparent Page Sharing Optimized for NUMA.....	3229
Resource Management in NUMA Architectures.....	3229
Using Virtual NUMA.....	3229
Virtual Topology in ESXi 8.0.....	3230
Virtual NUMA Controls.....	3230
Specifying NUMA Controls.....	3231
Associate Virtual Machines with Specific Processors.....	3232
Associate Memory Allocations with Specific NUMA Nodes Using Memory Affinity.....	3233
Associate Virtual Machines with Specified NUMA Nodes.....	3234
Advanced Attributes in vSphere.....	3234
Set Advanced Host Attributes.....	3235
Advanced Memory Attributes.....	3235
Advanced NUMA Attributes.....	3236
Set Advanced Virtual Machine Attributes.....	3237
Advanced Virtual Machine Attributes.....	3237
Advanced Virtual NUMA Attributes.....	3238
Latency Sensitivity.....	3239
Virtual Hyperthreading support for VMs.....	3240
vHT Full CPU Reservation.....	3240
Activate vHT for VMs.....	3241
Reliable Memory.....	3241
Backing Guest vRAM with 1GB Pages.....	3241
Fault Definitions.....	3242
Virtual Machine is Pinned.....	3242
Virtual Machine not Compatible with any Host.....	3242
VM/VM DRS Rule Violated when Moving to another Host.....	3242
Host Incompatible with Virtual Machine.....	3242
Host Has Virtual Machine That Violates VM/VM DRS Rules.....	3242
Host has Insufficient Capacity for Virtual Machine.....	3243
Host in Incorrect State.....	3243
Host Has Insufficient Number of Physical CPUs for Virtual Machine.....	3243
Host has Insufficient Capacity for Each Virtual Machine CPU.....	3243
The Virtual Machine Is in vMotion.....	3243
No Active Host in Cluster.....	3243
Insufficient Resources.....	3243
Insufficient Resources to Satisfy Configured Failover Level for HA.....	3243
No Compatible Hard Affinity Host.....	3243
No Compatible Soft Affinity Host.....	3243
Soft Rule Violation Correction Disallowed.....	3244
Soft Rule Violation Correction Impact.....	3244

DRS Troubleshooting Information.....	3244
Cluster Problems.....	3244
Load Imbalance on Cluster.....	3244
Cluster is Yellow.....	3244
Cluster is Red Because of Inconsistent Resource Pool.....	3245
Cluster Is Red Because Failover Capacity Is Violated.....	3245
No Hosts are Powered Off When Total Cluster Load is Low.....	3245
Hosts Are Powered-off When Total Cluster Load Is High.....	3245
DRS Seldom or Never Performs vMotion Migrations.....	3246
Host Problems.....	3246
DRS Recommends Host Be Powered on to Increase Capacity When Total Cluster Load Is Low.....	3246
Total Cluster Load Is High.....	3246
Total Cluster Load Is Low.....	3247
DRS Does Not Evacuate a Host Requested to Enter Maintenance or Standby Mode.....	3247
DRS Does Not Move Any Virtual Machines onto a Host.....	3247
DRS Does Not Move Any Virtual Machines from a Host.....	3248
VM migrations initiated by DRS due to incompatibility between the VMs and their host.....	3248
Virtual Machine Problems.....	3249
Insufficient CPU or Memory Resources.....	3249
VM/VM DRS Rule or VM/Host DRS Rule Violated.....	3250
Virtual Machine Power On Operation Fails.....	3250
DRS Does Not Move the Virtual Machine.....	3251
vSphere Availability.....	3252
You Can Minimize Downtime with vSphere.....	3252
Reduce Planned Downtime with vSphere.....	3252
Prevent Unplanned Downtime with vSphere.....	3252
vSphere HA Provides Rapid Recovery from Outages.....	3253
vSphere Fault Tolerance Provides Continuous Availability.....	3254
Protecting vCenter Server with vCenter High Availability.....	3254
Protecting vCenter Server with VMware Service Lifecycle Manager.....	3254
Creating and Using vSphere HA Clusters.....	3254
How vSphere HA Works.....	3255
Primary and Secondary Hosts.....	3255
Host Failure Types.....	3256
Determining Responses to Host Issues.....	3256
VM and Application Monitoring.....	3258
VM Component Protection.....	3259
Network Partitions.....	3260
Datastore Heartbeating.....	3260
vSphere HA Security.....	3261

vSphere HA Admission Control.....	3262
Cluster Resources Percentage Admission Control.....	3262
Slot Policy Admission Control.....	3264
Dedicated Failover Hosts Admission Control.....	3266
vSphere HA Interoperability.....	3267
Using vSphere HA with vSAN.....	3267
Using vSphere HA and DRS Together.....	3268
Other vSphere HA Interoperability Issues.....	3269
Creating a vSphere HA Cluster.....	3270
vSphere HA Checklist.....	3270
Create a vSphere HA Cluster in the vSphere Client.....	3271
Configuring vSphere Availability Settings.....	3271
Configuring Responses to Failures.....	3272
Configure Proactive HA.....	3274
Configure Admission Control.....	3275
Configure Heartbeat Datastores.....	3276
Set Advanced Options.....	3276
Best Practices for VMware vSphere® High Availability Clusters.....	3279
Best Practices for Networking.....	3279
Best Practices for Interoperability.....	3281
Best Practices for Cluster Monitoring.....	3282
Change in behavior for HA VIBs.....	3282
Providing Fault Tolerance for Virtual Machines.....	3282
How Fault Tolerance Works.....	3282
Fault Tolerance Use Cases.....	3283
Fault Tolerance Requirements, Limits, and Licensing.....	3283
Fault Tolerance Interoperability.....	3284
vSphere Features not Supported with Fault Tolerance.....	3284
Features and Devices Incompatible with Fault Tolerance.....	3285
Using Fault Tolerance with DRS.....	3285
Preparing Your Cluster and Hosts for Fault Tolerance.....	3286
Fault Tolerance Checklist.....	3286
Configure Networking for Host Machines.....	3287
Using Fault Tolerance.....	3287
Validation Checks for Turning On Fault Tolerance.....	3288
Turn On Fault Tolerance.....	3288
Turn Off Fault Tolerance.....	3289
Suspend Fault Tolerance.....	3290
Migrate Secondary.....	3290
Test Failover.....	3290

Test Restart Secondary.....	3290
Upgrade Hosts Used for Fault Tolerance.....	3291
Activate Fault Tolerance Encryption.....	3291
Best Practices for Fault Tolerance.....	3292
Activate Metro Cluster Fault Tolerance.....	3293
Legacy Fault Tolerance.....	3294
Troubleshooting Fault Tolerant Virtual Machines.....	3294
Hardware Virtualization Not Enabled.....	3294
Compatible Hosts Not Available for Secondary VM.....	3295
Secondary VM on Overcommitted Host Degrades Performance of Primary VM.....	3295
Increased Network Latency Observed in FT Virtual Machines.....	3295
Some Hosts Are Overloaded with FT Virtual Machines.....	3296
Losing Access to FT Metadata Datastore.....	3296
Turning On vSphere FT for Powered-On VM Fails.....	3296
FT Virtual Machines not Placed or Evacuated by vSphere DRS.....	3296
Fault Tolerant Virtual Machine Failovers.....	3297
vCenter High Availability.....	3298
Plan the vCenter HA Deployment.....	3298
vCenter Architecture Overview.....	3298
vCenter HA Hardware and Software Requirements.....	3299
vCenter HA Configuration Workflow.....	3299
Configure the Network.....	3300
Configure vCenter HA With the vSphere Client.....	3301
Manage the vCenter HA Configuration.....	3302
Set Up SNMP Traps.....	3302
Set Up Your Environment to Use Custom Certificates.....	3303
Manage vCenter HA SSH Keys.....	3303
Initiate a vCenter HA Failover.....	3304
Edit the vCenter HA Cluster Configuration.....	3304
Perform Backup and Restore Operations.....	3305
Remove a vCenter HA Configuration.....	3306
Reboot All vCenter HA Nodes.....	3306
Change the Server Environment.....	3306
Collecting Support Bundles for a vCenter HA Node.....	3306
Troubleshoot Your vCenter HA Environment.....	3307
vCenter HA Clone Operation Fails During Deployment.....	3307
Redeploy the Passive or Witness node.....	3307
vCenter HA Deployment Fails with an Error.....	3308
Troubleshooting a Degraded vCenter HA Cluster.....	3308
Recovering from Isolated vCenter HA Nodes.....	3309

Resolving Failover Failures.....	3309
VMware vCenter® HA Alarms and Events.....	3310
Patching a vCenter High Availability Environment.....	3311
Reduced Downtime Upgrade for vCenter HA.....	3311
vSphere Monitoring and Performance.....	3312
Monitoring Inventory Objects with Performance Charts.....	3312
Performance Chart Types.....	3313
Data Counters.....	3313
Metric Groups in vSphere.....	3315
Data Collection Intervals.....	3316
Data Collection Levels.....	3317
View Performance Charts.....	3318
Performance Charts Options Available Under the View Menu.....	3318
Overview Performance Charts.....	3319
Clusters.....	3319
Data centers.....	3329
Datastores and Datastore Clusters.....	3332
Hosts.....	3341
Resource Pools.....	3370
vApps.....	3376
Virtual Machines.....	3380
Working with Advanced and Custom Charts.....	3406
View Advanced Performance Charts in the vSphere Client.....	3407
Change Advanced Chart Settings.....	3407
Create a Custom Advanced Chart.....	3408
Delete a Custom Advanced Chart View.....	3408
Save Chart Data to a File.....	3409
Troubleshoot and Enhance Performance.....	3409
Solutions for Consistently High CPU Usage.....	3409
Solutions for Memory Performance Problems.....	3410
Solutions for Storage Performance Problems.....	3410
Solutions for Disk Performance Problems.....	3411
Solutions for Poor Network Performance.....	3412
Empty Performance Charts.....	3413
Solutions for Performance Issues With Memory Mode.....	3414
Monitoring Guest Operating System Performance.....	3414
Enable Statistics Collection for Guest Operating System Performance Analysis.....	3415
View Performance Statistics for Windows Guest Operating Systems.....	3415
Monitoring Host Health Status.....	3415
Monitor Hardware Health Status in the vSphere Client.....	3416

Monitoring and Diagnostics of vSphere Health.....	3416
Skyline Health.....	3416
Check health of the system.....	3416
VMware Skyline Health Diagnostics Tool.....	3417
Register VMware Skyline Health Diagnostics plugin from vSphere Client.....	3417
Use VMware Skyline Health Diagnostics Plugin from vSphere Client.....	3418
Monitoring Events, Alarms, and Automated Actions.....	3418
View Events in vSphere Client.....	3420
Export Events in vSphere Client.....	3420
View System Event Log.....	3420
Export System Event Log Data.....	3421
Consolidate Identical Events.....	3421
Configure Events Burst Filter.....	3422
Streaming Events to a Remote Syslog Server.....	3423
Forward vCenter Server Log Files to Remote Syslog Server.....	3424
Configure Streaming of Events to a Remote Syslog Server.....	3424
Retention of Events in the vCenter Server Database.....	3425
Configure Database Settings.....	3425
View Triggered Alarms.....	3426
Live Refresh of Recent Tasks and Alarms.....	3426
Set an Alarm in the vSphere Client.....	3426
Create or Edit Alarms.....	3426
Specify Alarm Name, Description, and Target.....	3427
Specify Alarm Rules.....	3427
Specify Alarm Reset Rules.....	3429
Review and Enable Alarm.....	3430
Disable Alarms on Specific Objects.....	3431
Acknowledge Triggered Alarms.....	3431
Reset Triggered Event Alarms.....	3431
Preconfigured vSphere Alarms.....	3431
View Alarm Definitions for Preconfigured vSphere Alarms.....	3436
Monitoring Solutions with the vCenter Solutions Manager.....	3436
View Solutions.....	3436
Monitoring the Health of Services and Nodes.....	3437
View the Health Status of Nodes.....	3437
View the Health Status of Services.....	3437
Performance Monitoring Utilities: resxtp and esxtp.....	3438
Using the esxtp Utility.....	3438
Using the resxtp Utility.....	3439
Using esxtp or resxtp in Interactive Mode.....	3440

Interactive Mode Command-Line Options.....	3440
CPU Panel.....	3442
CPU Power Panel.....	3445
Memory Panel.....	3446
Storage Adapter Panel.....	3449
Storage Device Panel.....	3450
Virtual Machine Storage Panel.....	3452
Network Panel.....	3453
Interrupt Panel.....	3454
Using Batch Mode.....	3455
Prepare for Batch Mode.....	3455
Use esxtop or resxtop in Batch Mode.....	3455
Batch Mode Command-Line Options.....	3455
Using Replay Mode.....	3456
Prepare for Replay Mode.....	3456
Use esxtop in Replay Mode.....	3456
Replay Mode Command-Line Options.....	3457
Using the vimtop Plug-In to Monitor the Resource Use of Services.....	3457
Monitor Services By Using vimtop in Interactive Mode.....	3457
Interactive Mode Command-Line Options.....	3457
Interactive Mode Single-Key Commands for vimtop.....	3458
Monitoring Networked Devices with SNMP and vSphere.....	3459
Using SNMP Traps with vCenter Server.....	3459
Configure SNMP Settings for vCenter Server.....	3459
Configure SNMP for ESXi.....	3460
Configure the SNMP Agent for Polling.....	3460
Configure ESXi for SNMPv1 and SNMPv2c.....	3461
Configure ESXi for SNMP v3.....	3462
Configure the Source of Hardware Events Received by the SNMP Agent.....	3466
Configure the SNMP Agent to Filter Notifications.....	3466
Configure SNMP Management Client Software.....	3466
SNMP Diagnostics.....	3467
Monitor Guest Operating Systems with SNMP.....	3467
VMware MIB Files.....	3467
SNMPv2 Diagnostic Counters.....	3469
System Log Files.....	3470
View System Logs on an ESXi Host.....	3470
Configure Syslog on ESXi Hosts.....	3470
System Logs.....	3470
ESXi System Logs.....	3471

vSphere Client Logs.....	3471
Export System Log Files.....	3471
ESXi Log Files.....	3472
Upload Logs Package to a VMware Service Request.....	3472
Manage the Amount of Virtual Machine Log Data.....	3473
Handle Virtual Machine Log Throttling.....	3473
Collecting Log Files.....	3474
Set Verbose Logging.....	3474
Collect vSphere Log Files.....	3474
Collect ESXi Log Files.....	3475
ESXi Log File Locations.....	3475
Configure Log Filtering on ESXi Hosts.....	3476
Turn Off Compression for vpxd Log Files.....	3477
ESXi VMkernel Files.....	3477
vSphere Single Host Management - VMware Host Client.....	3478
What is VMware Host Client.....	3478
Use the VMware Host Client.....	3478
Start the VMware Host Client and Log In.....	3479
Log Out of the VMware Host Client.....	3479
How to Customize Your VMware Host Client UI Theme.....	3479
Configure a Login Banner for the UI Login Screen of the VMware Host Client.....	3480
Leave and Rejoin the Customer Experience Improvement Program in the VMware Host Client.....	3487
Host Management with the VMware Host Client.....	3487
Manage System Settings in the VMware Host Client.....	3487
Manage Advanced Settings in the VMware Host Client.....	3487
Create an Initial Welcome Message for the Direct Console User Interface and the VMware Host Client.....	3487
Configure the VMware Host Client User Interface Session Timeout.....	3488
Configure the SOAP Session Timeout in the VMware Host Client.....	3489
Configure the Passwords and Account Lockout Policy in the VMware Host Client.....	3489
Configure Syslog in the VMware Host Client.....	3493
Configure Advanced TLS/SSL Key Options.....	3493
Configure Userworld Memory Zeroing.....	3494
Change Autostart Configuration in the VMware Host Client.....	3495
Edit the Time Configuration of an ESXi Host in the VMware Host Client.....	3496
Manage Hardware for an ESXi Host by Using the VMware Host Client.....	3497
Host Power Management Policies in ESXi.....	3497
Change Power Management Policies in the VMware Host Client.....	3497
Change the Hardware Label in the VMware Host Client.....	3498
Licensing for ESXi Hosts.....	3498
View Licensing Information About the VMware Host Client Environment.....	3500

Assign a License Key to an ESXi Host in the VMware Host Client.....	3500
Remove a License from an ESXi Host in the VMware Host Client.....	3500
Manage Services in the VMware Host Client.....	3501
Managing Security and Users for an ESXi Host by Using the VMware Host Client.....	3501
Manage Host Authentication by Using the VMware Host Client.....	3501
Manage Host Certificates by Using the VMware Host Client.....	3502
Manage Users with the VMware Host Client.....	3503
Manage ESXi Roles in the VMware Host Client.....	3505
Managing Hosts in vCenter Server.....	3506
Update Your VMware Host Client Environment to the Latest Version.....	3506
Unable to Connect from the VMware Host Client to an ESXi Host After Upgrading to a Newer Version of ESXi.....	3507
Switch to the vSphere Client.....	3508
Disconnect an ESXi Host from vCenter Server by Using the VMware Host Client.....	3508
Reboot or Shut Down an ESXi Host in the VMware Host Client.....	3508
Using the ESXi Shell.....	3508
Enable the Secure Shell (SSH) in the VMware Host Client.....	3509
Enable the ESXi Console Shell in the VMware Host Client.....	3509
Create a Timeout for ESXi Shell Availability in the VMware Host Client.....	3509
Create a Timeout for Idle ESXi Shell Sessions in the VMware Host Client.....	3510
Place an ESXi Host in Maintenance Mode in the VMware Host Client.....	3510
Manage Permissions in the VMware Host Client.....	3511
Permission Validation.....	3511
Assign Permissions to a User for an ESXi Host in the VMware Host Client.....	3511
Remove Permissions for a User in the VMware Host Client.....	3512
Assign User Permissions for a Virtual Machine in the VMware Host Client.....	3512
Remove Permissions for a Virtual Machine in the VMware Host Client.....	3512
Generate a Support Bundle in the VMware Host Client.....	3512
Lockdown Mode in the VMware Host Client.....	3513
Put an ESXi Host in Normal Lockdown Mode by Using the VMware Host Client.....	3514
Put an ESXi Host in Strict Lockdown Mode by Using the VMware Host Client.....	3514
Exit Lockdown Mode by Using the VMware Host Client.....	3514
Specify Lockdown Mode Exception Users in the VMware Host Client.....	3515
Manage CPU Resources by Using the VMware Host Client.....	3515
View Processor Information by Using the VMware Host Client.....	3515
Assign a Virtual Machine to a Specific Processor in the VMware Host Client.....	3515
Monitor an ESXi Host in the VMware Host Client.....	3516
View Charts in the VMware Host Client.....	3516
Monitor Hardware Health Status in the VMware Host Client.....	3516
View Events in the VMware Host Client.....	3516

View Tasks in the VMware Host Client.....	3517
View System Logs in the VMware Host Client.....	3517
View Notifications in the VMware Host Client.....	3518
Virtual Machine Management with the VMware Host Client.....	3518
Create a Virtual Machine in the VMware Host Client.....	3518
Deploy a Virtual Machine from an OVF or OVA File in the VMware Host Client.....	3521
OVF and OVA Limitations for the VMware Host Client.....	3521
Deploy a Virtual Machine from an OVF or OVA File in the VMware Host Client.....	3522
Register an Existing Virtual Machine in the VMware Host Client.....	3523
Use Consoles in the VMware Host Client.....	3523
Install the VMware Remote Console Application in the VMware Host Client.....	3524
Launch Remote Console for a Virtual Machine in the VMware Host Client.....	3524
Open a Virtual Machine Console in the VMware Host Client.....	3525
Manage a Guest Operating System in the VMware Host Client.....	3525
Shut Down or Restart a Guest Operating System by Using the VMware Host Client.....	3525
Change the Guest Operating System in the VMware Host Client.....	3525
Introduction to VMware Tools.....	3526
Configuring a Virtual Machine in the VMware Host Client.....	3527
Check the Virtual Machine Compatibility in the VMware Host Client.....	3528
Change the Name of a Virtual Machine in the VMware Host Client.....	3528
View the Location of the Virtual Machine Configuration File in the VMware Host Client.....	3528
Configure the Virtual Machine Power States in the VMware Host Client.....	3528
Edit the Configuration File Parameters in the VMware Host Client.....	3530
Configure Autostart Options for a Virtual Machine in the VMware Host Client.....	3530
Upgrade Virtual Machine Compatibility by Using the VMware Host Client.....	3531
Virtual CPU Configuration and Limitations.....	3531
Virtual Memory Configuration.....	3535
Network Virtual Machine Configuration.....	3538
Virtual Disk Configuration.....	3541
Virtual Machine Controller Configuration in the VMware Host Client.....	3547
Other Virtual Machine Device Configuration in the VMware Host Client.....	3551
Securing Virtual Machines with Intel Software Guard Extensions.....	3557
Manage Virtual Machines in the VMware Host Client.....	3558
Access a Virtual Machine in the VMware Host Client.....	3559
Power States of a Virtual Machine in the VMware Host Client.....	3559
Use Virtual Machine Column Configuration in the VMware Host Client.....	3560
Remove Virtual Machines from a Host in the VMware Host Client.....	3560
Remove Virtual Machines from a Datastore in the VMware Host Client.....	3560
Register a Virtual Machine in the VMware Host Client.....	3561
Manage Virtual Machines With Snapshots.....	3561

Monitor a Virtual Machine in the VMware Host Client.....	3573
View Virtual Machine Performance Charts in the VMware Host Client.....	3573
View Virtual Machine Events in the VMware Host Client.....	3573
View Virtual Machine Tasks in the VMware Host Client.....	3574
View Virtual Machine Log Browser in the VMware Host Client.....	3574
View Virtual Machine Notifications in the VMware Host Client.....	3574
Configure Virtual Machine Hardware in the VMware Host Client.....	3575
Virtual CPU Configuration and Limitations.....	3575
Virtual CPU Limitations.....	3576
Configuring Multicore Virtual CPUs.....	3577
Change the Number of Virtual CPUs.....	3577
Allocate CPU Resources in the VMware Host Client.....	3578
Virtual Memory Configuration.....	3578
Change the Memory Configuration.....	3579
Allocate Memory Resources to a Virtual Machine in the VMware Host Client.....	3580
Change Memory Hot Add Settings.....	3580
Add an NVDIMM device to a VM in the VMware Host Client.....	3581
Network Virtual Machine Configuration.....	3581
Network Adapter Basics.....	3582
Network Adapters and Legacy Virtual Machines.....	3583
Change the Configuration of the Virtual Network Adapter in the VMware Host Client.....	3583
Add a Network Adapter to a Virtual Machine in the VMware Host Client.....	3584
Virtual Disk Configuration.....	3584
About Virtual Disk Provisioning Policies.....	3585
Change the Virtual Disk Configuration in the VMware Host Client.....	3585
Add a New Standard Hard Disk to a Virtual Machine in the VMware Host Client.....	3586
Add an Existing Hard Disk to a Virtual Machine in the VMware Host Client.....	3588
Add a Persistent Memory Disk in the Host Client.....	3588
Use Disk Shares to Prioritize Virtual Machines in the VMware Host Client.....	3589
Virtual Machine Controller Configuration in the VMware Host Client.....	3590
Add a USB Controller to a Virtual Machine.....	3590
Add SCSI Controllers in the VMware Host Client.....	3591
Change the SCSI Bus Sharing Configuration in the VMware Host Client.....	3591
Change the SCSI Controller Type in the VMware Host Client.....	3592
About VMware Paravirtual SCSI Controllers.....	3592
Add a Paravirtual SCSI Controller in the VMware Host Client.....	3593
Add a SATA Controller to a Virtual Machine in the VMware Host Client.....	3593
Add an NVMe Controller in the VMware Host Client.....	3594
Other Virtual Machine Device Configuration in the VMware Host Client.....	3594
Add a CD or DVD Drive to a Virtual Machine in the VMware Host Client.....	3594

Add a Floppy Drive to a Virtual Machine in the VMware Host Client.....	3595
Add a USB Device to a Virtual Machine in the VMware Host Client.....	3595
Add a Sound Controller to a Virtual Machine in the VMware Host Client.....	3596
Parallel and Serial Port Configuration in the VMware Host Client.....	3596
How do I Add a Virtual Watchdog Timer Device to a Virtual Machine.....	3598
Add a Precision Clock Device to a Virtual Machine in the VMware Host Client.....	3599
Add a PCI Device to a Virtual Machine in the VMware Host Client.....	3599
Securing Virtual Machines in the VMware Host Client.....	3600
Activate vSGX on a Virtual Machine in the VMware Host Client.....	3600
Deactivate vSGX on a Virtual Machine in the VMware Host Client.....	3601
Remove a vTPM device from a VM in the VMware Host Client.....	3601
Activate or Deactivate Virtualization-based Security on an Existing VM in the VMware Host Client.....	3602
Managing Storage in the VMware Host Client.....	3603
Datastores in the VMware Host Client.....	3603
View Datastore Information in the VMware Host Client.....	3603
Create a VMFS Datastore in the VMware Host Client.....	3603
Increasing VMFS Datastore Capacity.....	3604
Mounting a Network File System Datastore in the VMware Host Client.....	3605
Unmount a Datastore in the VMware Host Client.....	3606
Using Datastore File Browser in the VMware Host Client.....	3607
Rename a Datastore in the VMware Host Client.....	3609
Delete a VMFS Datastore in the VMware Host Client.....	3609
Storage Thin Provisioning in the VMware Host Client.....	3610
Managing Storage Adapters in the VMware Host Client.....	3611
View Storage Adapters in the VMware Host Client.....	3611
Configuring Software iSCSI Adapters in the VMware Host Client.....	3611
Managing Storage Devices in the VMware Host Client.....	3619
View Storage Devices in the VMware Host Client.....	3620
Clear a Device Partition Table in the VMware Host Client.....	3620
Edit Individual Device Partitions in the VMware Host Client.....	3620
Manage Persistent Memory.....	3620
Modes of Consumption of the Persistent Memory Resources of the Host.....	3620
Structure of the PMem Datastore.....	3622
Monitoring Storage in the VMware Host Client.....	3623
Monitor Datastores in the VMware Host Client.....	3624
Monitoring vSAN in the VMware Host Client.....	3624
Performing Storage Refresh and Rescan Operations in the VMware Host Client.....	3628
Perform Adapter Rescan in the VMware Host Client.....	3628
Perform Device Rescan in the VMware Host Client.....	3628
Change the Number of Scanned Storage Devices in the VMware Host Client.....	3628

Networking in the VMware Host Client.....	3629
Manage Port Groups in the VMware Host Client.....	3629
View Port Group Information in the VMware Host Client.....	3629
Add a Virtual Switch Port Group in the VMware Host Client.....	3629
Edit Port Group Settings in the VMware Host Client.....	3630
Remove a Virtual Switch Port Group in the VMware Host Client.....	3633
Manage Virtual Switches in the VMware Host Client.....	3633
View Virtual Switch Information in the VMware Host Client.....	3633
Add a Standard Virtual Switch in the VMware Host Client.....	3633
Remove a Standard Virtual Switch in the VMware Host Client.....	3635
Add a Physical Uplink to a Virtual Switch in the VMware Host Client.....	3635
Edit Virtual Switch Settings in the VMware Host Client.....	3635
Manage Physical Network Adapters in the VMware Host Client.....	3637
View Physical Network Adapter Information in the VMware Host Client.....	3638
Edit Physical NICs in the VMware Host Client.....	3638
Manage VMkernel Network Adapters in the VMware Host Client.....	3638
View VMkernel Network Adapter Information in the VMware Host Client.....	3638
Add a VMkernel Network Adapter in the VMware Host Client.....	3638
Edit VMkernel Network Adapter Settings in the VMware Host Client.....	3639
Remove a VMkernel Network Adapter in the VMware Host Client.....	3640
View TCP/IP Stack Configuration on a Host in the VMware Host Client.....	3640
Change the Configuration of a TCP/IP Stack on a Host in the VMware Host Client.....	3640
Configure ESXi Firewall in the VMware Host Client.....	3641
Manage ESXi Firewall Settings by Using the VMware Host Client.....	3641
Add Allowed IP Addresses for an ESXi Host by Using the VMware Host Client.....	3642
Monitor Networking Events and Tasks in the VMware Host Client.....	3642
Monitor Port Groups in the VMware Host Client.....	3642
Monitor Virtual Switches in the VMware Host Client.....	3642
Monitor Physical Network Adapters in the VMware Host Client.....	3643
Monitor VMkernel Network Adapters in the VMware Host Client.....	3643
Monitor TCP/IP Stacks in the VMware Host Client.....	3643
Setup for Windows Server Failover Clustering on VMware vSphere.....	3644
Getting Started with WSFC on VMware® vSphere®	3644
Clustering Configuration on VMware® vSphere®.....	3644
Hosting virtual nodes of a WSFC on a Single Host.....	3645
Clustering Virtual Machines Across Physical Hosts.....	3645
Clustering Physical Machines with Virtual Machines.....	3646
Hardware and Software Requirements for WSFC on VMware® vSphere®.....	3647
What Shared Storage Configurations are Supported with WSFC?.....	3648
ESXi Supports Round Robin Path Selection Policy (PSP_RR) for disk resources used by WSFC.....	3649

ESXi Supports iSCSI for WSFC.....	3649
ESXi Supports FCoE for WSFC.....	3650
ESXi Supports vMotion for WSFC.....	3650
VMware vSphere® Virtual Volumes™ Support for WSFC.....	3651
You Can Hot Extend a Shared vVol Disk.....	3651
vSphere WSFC Setup Limitations.....	3653
WSFC and Booting from a SAN.....	3654
Cluster Virtual Machines Across Physical ESXi Hosts.....	3654
Clustered VMDK support for WSFC.....	3655
Activating Clustered VMDK Support.....	3655
Configuration Maximums for Clustered VMDK Support.....	3655
Recommendations for using Clustered VMDKs with WSFC.....	3656
Requirements for using Clustered VMDKs with WSFC.....	3656
Limitations of Clustered VMDK support for WSFC.....	3656
Create the First Node for WSFC Clusters Across Physical Hosts.....	3657
Create Additional Nodes for a WSFC Across Physical Hosts.....	3657
Add Hard Disks to the First VM, first node of a WSFC, for Clusters Across Physical Hosts with pRDMs.....	3658
Add Hard Disks to the First Node of a WSFC Across Physical Hosts with vVols.....	3659
Add Hard Disks to First Node for Cluster Across Physical Hosts with Clustered VMDKs on VMFS.....	3660
Add Hard Disks to Additional Nodes for Clusters Across Physical Hosts.....	3662
Use WSFC in an vSphere HA and vSphere DRS Environment.....	3665
Enable vSphere HA and vSphere DRS in ESXi.....	3665
Create VM-VM Affinity Rules for WSFC Virtual Machines.....	3665
Set DRS Automation Level for WSFC Virtual Machines.....	3666
Using vSphere DRS Groups and VM-Host Affinity Rules with WSFC Virtual Machines.....	3666
Create a Virtual Machine DRS Group (WSFC).....	3667
Create a Host DRS Group (WSFC).....	3667
Set up VM-Host Affinity Rules for DRS Groups (WSFC).....	3668
vSphere WSFC Setup Checklist.....	3668
Cluster Virtual Machines on One Physical Host.....	3671
Create the First Node of a WSFC on One Physical ESXi Host.....	3671
Create Additional Nodes for Clusters on One Physical Host.....	3671
Add Hard Disks to the First Node for Clusters on One Physical Host.....	3672
Add Hard Disks to Additional Nodes for Clusters on One Physical Host.....	3672
Cluster Physical and Virtual Machines.....	3673
Create the First Node for a Cluster of Physical and Virtual Machines.....	3673
Create the Second Node for a Cluster of Physical and Virtual Machines.....	3673
Add Hard Disks to the Second Node for a Cluster of Physical and Virtual Machines.....	3674
Install Microsoft Cluster Service.....	3675
Create Additional Physical-Virtual Pairs.....	3675

VMware Aria Operations Plugin in vCenter Server..... 3676

Introduction..... 3676

Installing VMware Aria Operations.....3676

 Install a New Instance of VMware Aria Operations in vCenter Server.....3676

Using the VMware Aria Operations Plugin in vCenter Server..... 3677

 Monitor Cluster Capacity.....3677

 View and Analyse Cluster Usage..... 3677

 Access VMware Aria Operations..... 3678

Documentation Legal Notice..... 3679

Release Notes

Release notes include product enhancements and notices, bug fixes, and resolved issues.

[VMware vSphere 8.0 Release Notes](#)

[VMware vCenter Server Photon OS Security Patches](#)

[VMware Host Client 2.5.0 Release Notes](#)

[ESXi Update and Patch Release Notes](#)

[vCenter Server Update and Patch Release Notes](#)

[VMware Host Client Release Notes](#)

VMware vSphere 8.0 Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [General Availability](#)
- [Internationalization](#)
- [Compatibility](#)
- [Before You Begin](#)
- [Installation and Upgrades for This Release](#)
- [Open Source Components for vSphere 8.0](#)
- [Product Support Notices](#)
- [Known Issues](#)

Introduction

VMware vSphere 8.0 | 11 OCT 2022
ESXi 8.0 | 11 OCT 2022 | Build 20513097
vCenter Server 8.0 | 11 OCT 2022 | Build 20519528
Check for additions and updates to these release notes.

What's New

This release of VMware vSphere 8.0 includes VMware ESXi 8.0 and VMware vCenter Server 8.0. Read about the new and enhanced features in this release in the [vSphere 8 Technical Overview Blog](#).

The vSphere 8.0 release notes do not include the following release notes:

- [VMware vSphere With Tanzu](#)
- [Tanzu Kubernetes releases](#)
- [VMware vSAN](#)
- [VMware Host Client](#)

General Availability

vSphere 8.0 is designated General Availability (GA). Custom ISO images that use ESXi 8.0 GA as a base image and include OEM firmware and drivers are available.

For more information, read the [vSphere 8 General Availability](#) blog.

IMPORTANT: Do not upgrade to this release if you use the Tanzu Kubernetes Grid service (TKG guest clusters) on vSphere along with the NSX Advanced Load Balancer (formerly known as Avi Networks) and have multiple Service Engine Groups configured. Upgrading to VMware vSphere with Tanzu 8.0 for such environments might result in failure to create new Tanzu Kubernetes guest clusters or cause upgrades of existing Supervisor Clusters to fail.

Internationalization

VMware vSphere 8.0 is available in the following languages:

- English
- Italian
- French
- German
- Spanish
- Japanese
- Korean
- Simplified Chinese
- Traditional Chinese

Components of vSphere 8.0, including vCenter Server, ESXi, the vSphere Client, and the VMware Host Client, do not accept non-ASCII input.

Compatibility

Virtual Machine Compatibility for ESXi

Virtual machines that are compatible with ESX 3.x and later (hardware version 4) are supported with ESXi 8.0. Virtual machines that are compatible with ESX 2.x and later (hardware version 3) are not supported. To use such virtual machines on ESXi 8.0, upgrade the virtual machine compatibility. See the [ESXi Upgrade](#) documentation.

Guest Operating System Compatibility for ESXi

To determine which guest operating systems are compatible with vSphere 8.0, use the ESXi 8.0 information in the [VMware Compatibility Guide](#).

The following guest operating system releases are deprecated or terminated in this release. Future vSphere releases will not support these guest operating systems:

- Windows Vista, Windows 2003 / R2, Windows XP: Deprecated
- Oracle Linux 5.x: Deprecated
- Oracle Linux 4.9: Terminated
- CentOS 5.x: Deprecated
- Asianux 3.0: Deprecated
- SUSE Linux Enterprise Server 9 SP4: Terminated
- SUSE Linux Enterprise Server 10 SP4: Deprecated
- SUSE Linux Enterprise Desktop 12: Deprecated
- Ubuntu releases 12.04, 18.10, 19.04 and 19.10: Terminated
- Debian 7.x and 8.x: Deprecated
- Debian 6.0: Terminated

- Photon OS 1.0: Terminated
- Flatcar Container Linux non-LTS releases: Terminated
- All OS X and macOS releases: Terminated
- FreeBSD 9.x and 10.x: Deprecated
- FreeBSD 7.x and 8.x: Terminated
- Solaris 10.x: Deprecated
- All eComStation releases: Terminated
- All SCO releases: Terminated
- All CoreOS releases: Terminated

Device Compatibility for ESXi

To determine which devices are compatible with ESXi 8.0, use the ESXi 8.0 information in the [VMware Compatibility Guide](#).

Hardware Compatibility for ESXi

To view a list of processors, storage devices, SAN arrays, and I/O devices that are compatible with vSphere 8.0, use the ESXi 8.0 information in the [VMware Compatibility Guide](#).

ESXi and vCenter Server Version Compatibility

The [VMware Product Interoperability Matrix](#) provides details about the compatibility of current and earlier versions of VMware vSphere components, including ESXi, VMware vCenter Server, and optional VMware products. Check the [VMware Product Interoperability Matrix](#) also for information about supported management and backup agents before you install ESXi or vCenter Server.

The vSphere Lifecycle Manager and vSphere Client are packaged with vCenter Server.

Before You Begin

vSphere 8.0 requires one CPU license for up to 32 physical cores. If a CPU has more than 32 cores, additional CPU licenses are required as announced in [Update to VMware's per-CPU Pricing Model](#). Prior to upgrading ESXi hosts, you can determine the number of licenses required using the license counting tool described in [Counting CPU licenses needed under new VMware licensing policy](#).

Installation and Upgrades for This Release

For download instructions, see [Download Broadcom products and software](#).

VMware Tools Bundling Changes in ESXi 8.0

- The following VMware Tools ISO images are bundled with ESXi 8.0:
 - **windows.iso**: VMware Tools 12.0.6 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - **linux.iso**: VMware Tools 10.3.24 ISO image for Linux OS with `glibc` 2.11 or later.
- The following VMware Tools ISO images are available for download:
- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
 - VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with `glibc` version earlier than 2.5.
 - **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
 - **darwin.iso**: Supports Mac OS X versions 10.11 and later.
- Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:
- [VMware Tools 12.0.6 Release Notes](#)
 - [Earlier versions of VMware Tools](#)
 - [What Every vSphere Admin Must Know About VMware Tools](#)
 - [VMware Tools for hosts provisioned with Auto Deploy](#)
 - [Updating VMware Tools](#)

Installation Notes

Read the [ESXi Installation and Setup](#) and the [vCenter Server Installation and Setup](#) documentation for guidance about installing and configuring ESXi and vCenter Server.

Although the installations are straightforward, several subsequent configuration steps are essential. Read the following documentation:

"License Management" in the [vCenter Server and Host Management](#) documentation

"Networking" in the [vSphere Networking](#) documentation

"Security" in the [vSphere Security](#) documentation for information on firewall ports

VMware's Configuration Maximums tool helps you plan your vSphere deployments. Use this tool to view the limits for virtual machines, ESXi, vCenter Server, vSAN, networking, and others. You can also compare limits for two or more product releases. The [VMware Configuration Maximums](#) tool is best viewed on larger format devices such as desktops and laptops.

- **Deprecation of USB or SD card devices for full ESXi installation:** Starting with ESXi 8.0, legacy SD and USB devices are supported with limitations and certification of new platforms with SD cards is not supported. SD and USB devices are supported for boot bank partitions. You can find a list of validated devices on partnerweb.vmware.com. The use of SD and USB devices for storing ESX-OSData partitions is being deprecated and the best practice is to provide a separate persistent local device with a minimum of 32 GB to store the ESX-OSData volume. For more details, see VMware knowledge base article [85685](#).
- **ESXi boot memory requirements increased:** The minimum memory requirements for ESXi to boot have increased from 4GB to 8GB. The minimum amount of memory needed to run VMs remains unchanged at 8GB.
- **The 'reboot -f' option is not supported for ESXi installation with a DPU:** Although ESXi supports the `-f force` reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.
- You cannot install ESXi on DPUs by using Auto Deploy. You can use the interactive and scripted methods for fresh installations of ESXi on DPU.

Upgrades and Installations Disallowed for Unsupported CPUs

- For ESXi hosts using Broadcom `bnxtnet` NIC drivers, make sure the NIC firmware is a compatible version such as 222.1.68.0 or higher, before you install or upgrade to ESXi 8.0. If you do not use a compatible firmware version, as specified in the [VMware Compatibility Guide](#), or as recommend by the OEM, you might see issues such as a drop in performance, firmware failure or ESXi host failure.
- vSphere 8.0 no longer supports CPUs which have been marked as End of Support or End of Life from hardware vendors. For more details, see VMware knowledge base article [82794](#).
- Deprecation of legacy BIOS:** In vSphere 8.0, booting ESXi hosts with the Unified Extensible Firmware Interface (UEFI) is strongly recommended. Some ESXi 8.0 hosts might not successfully boot in legacy BIOS mode. If this change affects your vSphere system, see VMware knowledge base articles [84233](#) and [89682](#) for details and action plans.

Upgrade notes

- Best practice for vSphere system upgrades is that the vCenter version is always greater than or equal to the ESXi version to ensure that you can use all new capabilities introduced with the latest vSphere release. For more information about vSphere build numbers, see [Build numbers and versions of VMware vCenter Server](#) and [Build numbers and versions of VMware ESXi/ESX](#). For vSphere back-in-time release upgrade restrictions, see VMware knowledge base article [67077](#).
- Before upgrading vCenters in Enhanced Link Mode (ELM), best practice is to create powered-off concurrent snapshots for all vCenter Servers and Platform Services Controllers within the Single Sign-On domain to prevent replication sync issues. If a rollback to snapshot is necessary, all PSC and vCenter Servers can be reverted to the previous state.
- vSphere 8.0 is not compatible with VMware NSX for vSphere (NSX-V). Upgrade paths to vSphere 8.0 from systems featuring NSX for vSphere are not supported. For more information, see the [NSX Migration Guide](#).
- Direct upgrade from ESXi 6.5 to 8.0 is not supported, because VMKAPI version 2.4, introduced in ESXi 6.5, is removed from ESXi 8.0. If you have ESXi 6.5 VIBs that depend on VMKAPI version 2.4 in a 6.7.x or 7.x environment, such VIBs prevent upgrade from 6.7.x or 7.x to 8.0. You must use only VIBs of version 6.7.x or later in the image that you use for upgrade to ESXi 8.0. To provide newer versions of the VIBs, you can recertify with the corresponding development kit for versions of VMKAPI later than 2.4. For more information, see VMware knowledge base article [88646](#).
- If you proceed to upgrade to ESXi 8.0 from a host with a device supported by the `nmlx4_en` driver, or a device removed in the `lpfc` driver, the following non-recoverable consequences can occur: lose access to storage or datastores, lose network access, or lose previous configuration on the host. Before upgrading to ESXi 8.0, you should replace devices previously supported by the `nmlx4_en` driver or devices that have been removed in the 8.0 `lpfc` driver. For the full list of devices no longer supported in ESXi 8.0, see VMware knowledge base article [88172](#).
- ESXi 8.0 installation and upgrade workflows block VIBs that do not have SHA256 checksum in their metadata, such as VIBs of ESXi version earlier than 6.7. You must replace such VIBs with later versions: 6.7.x, 7.x and 8.0.
- ESXi upgrade on DPUs is not supported by the interactive or scripted method. Instead, you can use vSphere Lifecycle Manager or ESXCLI. For example, an `esxcli software *` command run on an ESXi host automatically triggers the same operation on the DPU if such is present on the host.

Open Source Components for vSphere 8.0

After you log in to your [Customer Connect](#) account, the copyright statements and licenses applicable to the open source software components distributed in vSphere 8.0 are available at https://customerconnect.vmware.com/en/downloads/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/8_0#open_source. On the **Open Source** tab, you can also download the source files for any GPL, LGPL, or other similar licenses that require the source code or modifications to source code to be made available for the most recent available release of vSphere.

Product Support Notices

- **Data Processing Unit (DPU) Support:** vSphere 8.0 adds support for NVIDIA and AMD Pensando SmartNIC devices, also called DPUs, in ESXi. For more information, see [Introducing VMware vSphere® Distributed Services Engine™ and Networking Acceleration by Using DPUs](#).
- **Marvell qlnativefc driver update:** With vSphere 8.0, the qlnativefc driver can generate and provide information about Fibre Channel Port Identifier (FPIN) reports.
- **TPM provision policy:** Starting with ESX 8.0, you can use the TPM provision policy where virtual TPM devices can be automatically replaced during clone or deployment operations. For more information, see [Windows 11 Support on vSphere](#).
- **Deprecation of N-Port ID Virtualization (NPIV):** NPIV, which is an ANSI T11 standard that defines how to register a single Fibre Channel HBA port with the fabric by using several worldwide port names (WWPNs), will be deprecated in a future release of vSphere due to many existing alternatives.
- **Deprecation of Integrated Windows Authentication (IWA):** Due to performance issues with IWA-based authentication, vSphere 8.0 deprecates the use of IWA. You can use AD over LDAP or ADFS.
- **Deprecation of Common Information Model (CIM) and Service Location Protocol (SLP):** Support for CIM and SLP is deprecated in ESXi 8.0 due to security issues and will be removed in a future release. As an alternative, consider using the Daemon Software Development Kit (DSDK) for solutions that rely on CIM, such as the CIM Provider Development Kit (CIMPDK) and the vSphere APIs for I/O Filtering (VAIO) Development Kit. No CIMPDK is released for vSphere 8.0, but CIM Providers for ESXi 7.x. continue to work on ESXi 8.0 to support a smooth upgrade process.
- **Smart Card mutual authentication moves to port 3128:** vCenter Server 8.0 moves the Smart Card mutual authentication to port 3128 and requires a restart of the Security Token Service (STS) during configuration. For more information, see [Configure the Reverse Proxy to Request Client Certificates](#).
- VMware inbox qedentv NIC driver from Marvell does not support NVMe/TCP in Enhanced Networking Stack (ENS) mode in ESXi 8.0.
- **Deprecation of vSphere Lifecycle Manager baselines:** Managing clusters with vSphere Lifecycle Manager baselines and baseline groups (legacy vSphere Update Manager workflows) is supported in vSphere 8.0, but support will drop in a future vSphere release. Instead of baselines and baseline groups, you can use vSphere Lifecycle Manager images to perform tasks on a cluster level such as install a desired ESXi version on all hosts in a cluster, install and update third-party software, update, and upgrade ESXi or firmware, generate recommendations, and use a recommended image for your cluster. For more details, see the blog [Introducing vSphere Lifecycle Management \(vLCM\)](#).
- **Predefined vSphere Lifecycle Manager baselines restriction to VMware content only:** As of vSphere 8.0, vSphere Lifecycle Manager predefined baselines are restricted only to VMware-provided content. Third-party content such as async drivers and tools no longer can be part of predefined baselines. If you need to add third-party content to update baselines, you must manually create custom baselines.
- **Deprecation of Patch Manager APIs:** With vSphere 8.0, Patch Manager APIs are deprecated. Patch Manager APIs are supported in vSphere 8.0, but support will discontinue in a future release of vSphere. Instead of Patch Manager APIs, you can use the latest vSphere APIs, documented in the [vSphere API automation reference guide](#).
- **Deprecation of local plug-ins:** Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. VMware plans to discontinue support for local plug-ins in a future vSphere release. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).
- **Removal of support for 32-bit userworlds:** ESXi 8.0 does not support 32-bit user worlds and you need to recompile any solution using the 32-bit subsystem.
- **Removal of VMKAPI v2.4:** vSphere 8.0 does not support VMKAPI v2.4 and you must use newer versions of the API and recertify with the latest development kit.
- **Removal of nmlx4_en driver:** ESXi 8.0 removes the nmlx4_en driver and all devices for this driver are not supported.
- **Removal of Trusted Platform Module (TPM) 1.2:** VMware discontinues support of TPM 1.2 and associated features such as TPM 1.2 with TXT. To get full use of vSphere features, you can use TPM 2.0 instead of TPM 1.2.

- **Removal of insecure ciphers:** X.509 certificates that specify a SHA-1 signature algorithm, or other weak signature algorithms, are no longer supported in vSphere 8.0. Prechecks prevent upgrade to vCenter Server 8.0 and ESXi 8.0 if certificates with a weak signature algorithm are in use. For remediation steps, see VMware knowledge base article [89424](#).
- **Removal of Software FCoE Adapters:** In vSphere 8.0, the option to configure software FCoE adapters that use the native FCoE stack in ESXi is removed and not supported. The change does not impact hardware FCoE adapters and drivers.
- **Discontinuation of support for I/O devices used by nmlx4_en and lpfc drivers:** VMware intends to discontinue support for I/O devices that reach EOL, such as devices previously supported by the `lpfc` driver and all devices for the `nmlx4_en` driver. For more information, see VMware knowledge base article [88172](#).
- **You cannot register a read-only VMDK as FCD:** In vSphere 8.0, any Virtual Storage Lifecycle Management API invoked on read only vmdk registered as FCD is not supported.
- **Discontinuation of support for Guest OS:** vSphere 8.0 drops support for the following guest operating systems:
 - eComStation
 - SCO Openserver
 - SCO Unixware
 - Oracle Linux 4.x
 - SLES9 SP4
 - Ubuntu 12.04 LTS
 - Debian 6.0
 - FreeBSD 7.x
 - FreeBSD 8.x
- **OpenSSL 3.0 support:** ESXi 8.0 supports OpenSSL 3.0 and drops support for TLS 1.0 and TLS 1.1. vCenter does not support OpenSSL 3.0. In vCenter, TLS 1.2 is enabled by default and TLS 1.0, and TLS 1.1 are disabled by default, but you can temporarily enable them.
- **Hardware timestamp-based Precision Time Protocol (PTP) certification for NICs:** vSphere 8.0 adds certification for NICs that support hardware timestamp-based PTP as part of the I/O Vendor Partner (IOVP) certification program.
- **NVMe over Fabrics (NVMe-oF) support for vSphere Virtual Volumes:** vSphere 8.0 adds NVMe-oF support for vSphere Virtual Volumes as part of the IOVP NVMe-FC certification program.
- **NVMeoF-RDMA scale enhancements:** With NVMeoF, you can scale NVMe namespaces and paths to 256 and 4,000 respectively in vSphere 8.0.
- **Advanced NVMe-oF Discovery Service Support:** vSphere 8.0 adds dynamic discovery of devices for compliant NVMe Discovery Services and Storage Arrays.
- **Syslog enhancements:** vSphere 8.0 unifies the format of logs by the ESXi syslog daemon across all of vSphere and VCF products and if your system consumes syslog log files directly from ESXi, you need to update your solutions. For more information, see [Configuring System Logging](#). You can optionally configure transmissions of syslog messages from ESXi in compliance with RFC 5424 or frame messages. For more information, see [Protocols, Formats and Framing of ESXi Syslog Messages](#). You can set all syslog controlling parameters by using either the vSphere Client or the VMware Host Client and do not need to use SSH or ESXCLI.
- **Phasing out LSI SAS controllers:** vSphere 8.0 can automatically and safely replace the LSI SAS controller for VMs on Windows 10 and later or Windows Server 2016 and later with the native VMware PVSCI controller, because the LSI SAS driver for Windows has reached end of life. For VMs on versions earlier than Windows 10 and Windows Server 2016, you can manually replace the LSI SAS controller with PVSCI, SATA, or a NVMe controller.
- **Compliance of the default RSA key length with U.S. Department of Defense Impact Level 6 (IL-6):** To conform with IL6 standards, the default RSA key length generated for vCenter Server certificates increases to 3072 bits in vSphere 8.0 from 2048 bits.
- **Administer user access to ESXi Shell:** Starting with vSphere 8.0, users with the Administrator role can remove or grant ESXi Shell access to user accounts. For more information, see [How Do You Configure a Security vSphere Host Profile](#).

- **Discontinuation of support for Apple Mac platforms:** ESXi 8.0 does not support Apple MacPro and Apple MacMini platforms, and macOS as a guest operating system. For more information, see VMware knowledge base article [88698](#).
- **Virtual hardware version 20:** ESXi 8.0 introduces virtual hardware version 20 to enable support for virtual machines with higher resource maximums, and:
 - Virtual NUMA Topology
 - Enhanced Direct Path I/O
 - Virtual Hyperthreading
 - vMotion App Notification
 - VM DataSets
 - OpenGL 4.3
 - UEFI 2.7A
- **Increased scalability with vSphere Lifecycle Manager:** With vSphere 8.0, scalability for operations with vSphere Lifecycle Manager images increases to 1,000 ESXi hosts from 280.
- **Support for UEFI 2.7A:** vSphere 8.0 complies with UEFI specification version 2.7A to support some Microsoft Windows 11 features.
- **vSphere Configuration Profiles:** vSphere 8.0 launches vSphere Configuration Profiles in tech preview. This capability allows you to manage ESXi cluster configurations by specifying a desired host configuration at the cluster level, automates the scanning of ESXi hosts for compliance to the specified Desired Configuration and remediates any host that is not compliant. The tech preview launch is applicable only to customers that use standard switches and does not support VMware vSphere Distributed Switch (VDS) and requires that you use vSphere Lifecycle Manager images to manage your cluster lifecycle.
- **Removal of RDMA over Converged Ethernet (RoCE) v1:** Starting with vSphere 8.0, VMware does not support the network protocol RoCE v1. You can use RoCEv2. Make sure you migrate your paravirtualized remote direct memory access (PVRDMA) network adapters for virtual machines and guest operating systems to an adapter that supports RoCEv2.

Known Issues

Installation, Upgrade, and Migration Issues

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliance$` shell. If you use VAMI, refresh your browser.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is deactivated, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

Miscellaneous Issues

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reinstall ESXi. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In the vSphere API Explorer, VMware Datacenter CLI (DCLI) and PowerCLI, you see an API option "contentinternal" that is not functional

You see an API option **contentinternal** in the metadata of either the vSphere API Explorer, DCLI and PowerCLI. For example, when you open `https://<your vCenter IP>/ui/app/devcenter/api-explorer`, you see the option in the **select API** drop-down menu. This option is not functional.

Workaround: Ignore the **contentinternal** API option and do not use it.

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client.

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: None.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example `vmnic6` and `vmnic7`) and one management controller (for example `vmnic8`):

```
:~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
:~] esxcli -nics -ls
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers `vmnic6` and `vmnic7` register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -
n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller `vmnic8` registers first with the vSphere Distributed Switch, the high-speed ethernet controllers `vmnic6` or `vmnic7` uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:

```
:~] esxcli -module -u
ionic_en :~] esxcli -module ionic_en :~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/
deviceInternal bind.
```

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcli -advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcli -advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough (SR-IOV or UPT) to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

Networking Issues

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

You see link flapping on NICs that use the `ntg3` driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Storage Issues

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vCenter Server and vSphere Client Issues

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade.`

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [90259](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files.`

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Guest OS Issues

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

VMware vCenter Server Photon OS Security Patches

This document contains the following sections

- [Introduction](#)
- [Installation Steps](#)
- [vCenter Server 8.0 Photon OS Security Patches](#)

Introduction

This document tracks the release of 8.x patches to the Photon Operating System bundled in VMware vCenter Server.

You can download the deliverables from the [Broadcom Support Portal](#). For download instructions, see [Download Broadcom products and software](#).

Installation Steps

To apply the Photon OS security patches to the vCenter Server Appliance, you can use one of the methods.

- Deploy a new vCenter Server by using either the GUI or the CLI installer.
For information about doing a fresh install of the vCenter Server Appliance, see [Deploying the vCenter Server Appliance](#).
- Upgrade to the version of the vCenter Server Appliance containing the latest Photon OS security patches by using either the GUI or the CLI installer.
For information about upgrading the vCenter Server Appliance, see [Upgrading the vCenter Server Appliance](#).
- Patch the appliance either by using the appliance shell or the Appliance Management Interface.
IMPORTANT: You can update the vCenter Server Appliance with Photon OS patches released within one and the same Update release.
For information on patching the vCenter Server Appliance, see [Patching and Updating vCenter Server 8.0 Deployments](#).
- Perform a file-based backup and restore where in the restore process you deploy a new appliance containing the latest Photon OS security patches..
For information performing a file-based backup and restore of the vCenter Server Appliance, see [Restore vCenter Server from a File-Based Backup](#).
- Migrate a vCenter Server on Windows instance to a version of the vCenter Server Appliance containing the latest Photon OS security patches.
For information about performing a migration of vCenter Server on Windows to vCenter Server Appliance, see [Migrating vCenter Server for Windows to vCenter Server Appliance](#).

vCenter Server 8.0 Photon OS Security Patches

vCenter Server 8.0 Update 3g

vCenter Server 8.0 Update 3g

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
July 29, 2025	24853646	vCenter Server 8.0 Update 3g (Security fixes for Photon OS)	apache-tomcat9	9.0.104-1.ph4	CVE-2025-24813, CVE-2025-31650, CVE-2025-31651
			expat	2.7.1-1.ph4	CVE-2024-8176
			freetype2	2.13.3-1.ph4	CVE-2025-27363
			gettext	0.21-6.ph4	CVE-2024-56171, CVE-2025-24928, CVE-2025-27113
			libarchive	3.4.3-11.ph4	CVE-2025-25724
			libtasn1	4.14-3.ph4	CVE-2024-12133
			libxml2	2.9.12-16.ph4	CVE-2024-56171, CVE-2025-24928, CVE-2025-27113, CVE-2025-32415

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
			linux	5.10.238-2.ph4	CVE-2021-47200, CVE-2021-47489, CVE-2022-21546, CVE-2022-49651, CVE-2024-57795, CVE-2023-52531, CVE-2023-52621, CVE-2024-49991, CVE-2024-50067, CVE-2023-52752, CVE-2024-26739, CVE-2024-26928, CVE-2024-35863, CVE-2024-35864, CVE-2024-35867, CVE-2024-35868, CVE-2024-27056, CVE-2024-27415, CVE-2024-35839, CVE-2024-56604, CVE-2024-56658, CVE-2024-57979, CVE-2025-21726, CVE-2025-21727, CVE-2025-21731, CVE-2025-21756, CVE-2025-21760, CVE-2025-21761, CVE-2025-21762, CVE-2025-21763, CVE-2025-21764, CVE-2025-21791, CVE-2025-21796, CVE-2025-21858, CVE-2024-8805, CVE-2025-21703, CVE-2025-21759, CVE-2025-21863, CVE-2025-21999, CVE-2025-22035, CVE-2025-22056
			openssh	8.9p1-10.ph4	CVE-2025-32728
			openssl	3.0.16-1.ph4	CVE-2024-9143
			perl-Data-Validate-I P	0.30-1.ph4	CVE-2021-29662
			postgresql13	13.20-1.ph4	CVE-2025-1094
			postgresql14	14.17-1.ph4	CVE-2025-1094
			python3-certifi	2023.11.17-2.ph4	CVE-2024-39689
			python3-idna	2.10-3.ph4	CVE-2024-3651
			ruby	2.7.4-17.ph4	CVE-2025-27219, CVE-2025-27220

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
			rubygem-nokogiri	1.13.6-1.ph4	CVE-2018-25032, CVE-2021-30560, CVE-2022-24836, CVE-2022-29181
			runc	1.1.14-1.ph4	CVE-2024-45310
			xerces-c	3.2.5-1.ph4	CVE-2024-23807
			vim	9.1.1441-1.ph4	CVE-2025-26603, CVE-2025-22134, CVE-2025-29768, CVE-2025-24014, CVE-2025-1215

vCenter Server 8.0 Update 3e

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
April 10, 2025	24674346	vCenter Server 8.0 Update 3e (Security fixes for Photon OS)	apache-tomcat9	9.0.98-1.ph4	CVE-2024-52316 CVE-2024-34750 CVE-2024-50379
			apr	1.7.2-3.ph4	CVE-2023-49582
			bash	5.0-5.ph4	CVE-2019-18276
			bindutils	9.18.27-1.ph4	CVE-2021-25216 CVE-2020-8625 CVE-2022-38177 CVE-2023-5517 CVE-2022-38178 CVE-2023-50387 CVE-2023-6516 CVE-2023-4408 CVE-2021-25215 CVE-2022-3094 CVE-2023-3341 CVE-2023-2828 CVE-2024-1975 CVE-2024-0760 CVE-2024-1737 CVE-2024-4076
			curl	8.7.1-4.ph4	CVE-2022-32207 CVE-2023-38545 CVE-2023-27534 CVE-2023-27533 CVE-2022-22576 CVE-2022-27775 CVE-2021-22946 CVE-2021-22926 CVE-2022-27782 CVE-2022-27781 CVE-2024-6197 CVE-2024-7264
			dbus	1.13.18-5.ph4	CVE-2023-34969

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
			docker	24.0.9-1.ph4	CVE-2024-24557
			expat	2.4.9-4.ph4	CVE-2024-45491 CVE-2022-25315 CVE-2022-23852 CVE-2024-45492 CVE-2022-25235 CVE-2022-25236 CVE-2022-22824 CVE-2022-22823 CVE-2022-22822 CVE-2021-45960 CVE-2022-22827 CVE-2022-22826 CVE-2022-40674 CVE-2021-46143 CVE-2022-43680 CVE-2022-23990 CVE-2024-45490 CVE-2022-25314 CVE-2023-52425
			file	5.39-3.ph4	CVE-2023-5388 CVE-2022-48554 CVE-2023-52434 CVE-2024-26584
			gdb	13.2-3.ph4	CVE-2023-1579 CVE-2023-25584
			glib	2.68.4-4.ph4	CVE-2023-32643 CVE-2023-32636 CVE-2023-29499 CVE-2024-34397
			gnutls	3.7.10-3.ph4	CVE-2002-2439 CVE-2024-0567 CVE-2024-0553
			krb5	1.17-12.ph4	CVE-2024-37371 CVE-2022-42898 CVE-2024-37370 CVE-2020-28196 CVE-2019-14844 CVE-2021-36222
			libssh2	1.11.0-2.ph4	CVE-2024-34161 CVE-2024-31079 CVE-2023-48795 CVE-2024-35200 CVE-2024-32760
			linux	5.10.233-2.ph4	CVE-2024-39484 CVE-2024-38538 CVE-2024-44987 CVE-2024-46722 CVE-2024-42284

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
					CVE-2024-53141 CVE-2024-50154 CVE-2024-42228 CVE-2024-44998 CVE-2024-43882 CVE-2024-38577 CVE-2024-38381 CVE-2024-56615 CVE-2024-46759 CVE-2024-43839 CVE-2024-46744 CVE-2024-46865 CVE-2024-42301 CVE-2024-46854 CVE-2024-50055 CVE-2024-56600 CVE-2024-42271 CVE-2024-46725 CVE-2024-46800 CVE-2024-56672 CVE-2024-44986 CVE-2024-56704 CVE-2024-49967 CVE-2024-46738 CVE-2024-50143 CVE-2024-49960 CVE-2024-41014 CVE-2024-50083 CVE-2024-41011 CVE-2024-50301 CVE-2024-46756 CVE-2024-56662 CVE-2024-57896 CVE-2024-41071 CVE-2024-44999 CVE-2024-46743 CVE-2024-50125 CVE-2024-46758 CVE-2024-50262 CVE-2024-50228 CVE-2024-50264 CVE-2024-38588 CVE-2024-42285 CVE-2024-53096 CVE-2024-41073 CVE-2024-56759 CVE-2024-56606 CVE-2024-46747 CVE-2024-44934 CVE-2024-50036 CVE-2024-50282 CVE-2024-46782 CVE-2024-46740 CVE-2024-49983 CVE-2024-46724 CVE-2024-50121

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
					CVE-2024-46798 CVE-2024-50115 CVE-2024-50131 CVE-2024-45026 CVE-2024-56605 CVE-2024-46723 CVE-2024-46757 CVE-2024-56601 CVE-2024-53142 CVE-2024-46731 CVE-2024-42302 CVE-2024-53103 CVE-2024-53099 CVE-2024-56631 CVE-2024-46821 CVE-2024-56603
			nghttp2	1.57.0-2.ph4	CVE-2024-28182
			nss	3.72-6.ph4	CVE-2021-43527 CVE-2023-5388 CVE-2024-0743
			openjre8	1.8.0.442-1.ph4	CVE-2012-0881 CVE-2013-4002
			openssl	3.0.15-1.ph4	CVE-2022-2068 CVE-2022-1292 CVE-2023-0464 CVE-2023-0215 CVE-2021-23840 CVE-2022-0778 CVE-2021-3712 CVE-2023-0286 CVE-2023-4807 CVE-2023-5363
			postgresql13	13.18-1.ph4	CVE-2023-39417 CVE-2024-0985 CVE-2024-7348
			postgresql14	14.15-1.ph4	CVE-2024-7348 CVE-2023-5869
			python3 python3-curses python3-devel python3-libs python3-macros python3-xml	3.10.11-12.ph4	CVE-2023-24329 CVE-2024-6232 CVE-2024-7592 CVE-2024-23342 CVE-2024-6923 CVE-2023-6597
			python3-Twisted	22.10.0-2.ph4	CVE-2020-10108 CVE-2022-39348 CVE-2022-24801 CVE-2022-21712
			python3-cryptography	41.0.7-2.ph4	CVE-2023-50782
			python3-pycryptodome	3.20.0-1.ph4	CVE-2023-52323
			python3-setuptools	65.5.1-6.ph4	CVE-2024-6345

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
			python3-urllib3	1.26.19-1.ph4	CVE-2021-33503
			rsync	3.4.1-1.ph4	CVE-2024-12084 CVE-2024-12085 CVE-2024-12086 CVE-2024-12087 CVE-2024-12747
			ruby	2.7.4-14.ph4	CVE-2020-36327 CVE-2021-43809 CVE-2021-3 CVE-2024-41946 CVE-2024-49761 CVE-2024-41123 CVE-2021-33621 CVE-2021-41819 CVE-2022-28739 CVE-2021-41817
			rubygem-nokogiri	1.12.5-5.ph4	CVE-2021-30560CV E-2018-25032
			rubygem-activesupp ort	7.1.3.4-1.ph4	CVE-2024-28103
			sysstat	12.7.2-3.ph4	CVE-2023-33204
			tcpdump	4.99.4-2.ph4	CVE-2024-2397
			vim	9.1.0724-2.ph4	CVE-2024-22667 CVE-2023-4733 CVE-2023-4734 CVE-2023-4752 CVE-2023-4735 CVE-2023-4736 CVE-2023-4750 CVE-2023-4738 CVE-2023-4781 CVE-2023-5535 CVE-2023-5344

vCenter Server 8.0 Update 3b

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
17 September 2024	24262322	vCenter Server 8.0 Update 3b (Security fixes for Photon OS)	ruby	2.7.4-12	CVE-2024-27282
			zlib	1.2.11-6.ph4	CVE-2023-45853
			util-linux	2.37.4-2.ph4	CVE-2024-28085
			openssl	3.0.14-2.ph4	CVE-2023-50782
					CVE-2024-2511
					CVE-2024-4741
			httpd	2.4.62-1.ph4	CVE-2024-5535
					CVE-2024-27316 CVE-2024-39573

			glibc	2.32-19.ph4	CVE-2023-4911 CVE-2024-2961
			expat	2.4.9-2.ph4	CVE-2023-52425 & CVE-2023-52426 CVE-2024-28757
			curl	8.7.1-1.ph4	CVE-2024-2004,CV E-2024-2398 CVE-2023-46218,CV E-2023-46219
			linux	5.10.219-3.ph4	CVE-2024-23307 and CVE-2024-22099 CVE-2024-26584
					CVE-2023-52447/20 23-52458/2023-524 82
					CVE-2024-26583, CVE-2024-26585, and CVE-2024-26589
					CVE-2023-52585
					CVE-2024-26642, CVE-2023-52620
					CVE-2024-26643
					CVE-2023-1192
					CVE-2024-36901
					CVE-2022-48666
			libxml2	2.9.12-14.ph4	CVE-2024-34459

vCenter Server 8.0 Update 2b

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
29 February 2024	23319993	vCenter Server 8.0 Update 2b (Security fixes for Photon OS)	apache-tomcat	8.5.93-2.ph4	CVE-2023-34981 CVE-2023-28709 CVE-2023-28708
			bindutils	9.16.42-4.ph4	CVE-2023-3341 CVE-2023-2829 CVE-2023-2828 CVE-2023-2911
			c-ares	1.19.1-1.ph4	CVE-2023-31130 CVE-2023-31147 CVE-2023-32067 CVE-2023-31124
			curl	8.1.2-6.ph4	CVE-2023-38545 CVE-2023-38546

		CVE-2023-38039
glib	2.68.4-1.ph4	CVE-2023-32643
grub2	2.06-11.ph4	CVE-2023-4692 CVE-2023-4693 CVE-2022-28736 CVE-2021-3697 CVE-2021-3695 CVE-2022-28734
httpd	2.4.58-2.ph4	CVE-2023-31122
libcap	2.43-3.ph4	CVE-2023-2602 CVE-2023-2603
libssh2	1.11.0-1.ph4	CVE-2020-22218
libuv	1.45.0-1.ph4	CVE-2024-24806
libxml2	2.9.12-12.ph4	CVE-2023-45322 CVE-2023-39615
linux	5.10.201-1.ph4	CVE-2023-4244 CVE-2023-42754 CVE-2023-42756 CVE-2023-42755 CVE-2023-42753 CVE-2023-22995 CVE-2023-2176 CVE-2023-2007 CVE-2023-4147 CVE-2023-4128 CVE-2023-0597
nss	3.72-3.ph4	CVE-2022-36320
open-vm-tools	12.3.5-1.ph4	CVE-2023-34058 CVE-2023-34059 CVE-2023-20900 CVE-2023-20867
openldap	2.4.57-4.ph4	CVE-2023-2953
openssh	8.9p1-4.ph4	CVE-2023-28531 CVE-2023-38408
openssl	3.0.9-8.ph4	CVE-2023-5678 CVE-2023-5363 CVE-2023-4807 CVE-2023-3817 CVE-2023-3446 CVE-2023-2975 CVE-2023-0464 CVE-2023-0465
perl	5.30.1-6.ph4	CVE-2023-31486
postgresql14	14.10-1.ph4	CVE-2023-2454 CVE-2023-2455
runc	1.1.4-11.ph4	CVE-2023-25809 CVE-2023-27561
sqlite	3.38.5-3.ph4	CVE-2023-36191

			tcpdump	4.99.4-1.ph4	CVE-2020-8036
			vim	9.0.2068-1.ph4	CVE-2023-46246 CVE-2023-3896 CVE-2023-2610 CVE-2023-2609 CVE-2023-2426
			xinetd	2.3.15-10.ph4	CVE-2013-4342
			samba-client	4.18.5-2.ph4	CVE-2018-10919 CVE-2023-0225 CVE-2023-0614 CVE-2023-0922

vCenter Server 8.0 Update 1c

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
27 July 2023	22088981	vCenter Server 8.0 Update 1c (Security fixes for Photon OS)	apache-tomcat	8.5.86-1.ph3	CVE-2022-42252
					CVE-2022-45143
			bindutils	9.16.38-1.ph3	CVE-2022-3736
			containerd	1.6.8-5.ph3	CVE-2023-25173
			curl	7.86.0-5.ph3	CVE-2022-43551
					CVE-2023-23914
					CVE-2023-23916
					CVE-2023-27535
					CVE-2023-27536
			dmidecode	3.5-1.ph3	CVE-2023-30630
			dnsmasq	2.85-4.ph3	CVE-2022-0934
			e2fsprogs	1.46.5-2.ph3	CVE-2022-1304
			expat	2.2.9-11.ph3	CVE-2022-43680
			gnutls	3.6.16-4.ph3	CVE-2023-0361
			grub2	2.06-4.ph3	CVE-2022-2601
					CVE-2022-28733
			httpd	2.4.56-1.ph3	CVE-2006-20001
					CVE-2022-36760
					CVE-2023-25690
					CVE-2023-27522
			krb5	1.17-4.ph3	CVE-2022-42898
			libarchive	3.3.3-9.ph3	CVE-2022-36227
			libxml2	2.9.11-10.ph3	CVE-2022-40303
					CVE-2022-40304
			linux	4.19.283-1.ph3	CVE-2020-28374
					CVE-2022-3643
					CVE-2022-42896

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
					CVE-2022-4378
					CVE-2023-0461
					CVE-2023-1281
					CVE-2023-26545
					CVE-2023-2002
					CVE-2023-32233
					CVE-2023-34256
					CVE-2023-23559
			nss	3.44-10.ph3	CVE-2020-25648
					CVE-2022-36320
			openssl	1.0.2zg-1.ph3	CVE-2022-4450
					CVE-2023-0215
			pkg-config	0.29.2-4.ph3	CVE-2021-3800
			python3	3.7.5-26.ph3	CVE-2020-10735
					CVE-2022-37454
					CVE-2022-45061
			sudo	1.9.5-5.ph3	CVE-2023-22809
			tar	1.30-6.ph3	CVE-2022-48303
			vim	8.2.5169-9.ph3	CVE-2022-47024
					CVE-2023-0049
					CVE-2023-0051
					CVE-2023-0054
					CVE-2023-0433
					CVE-2023-2426

vCenter Server 8.0b

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
14 February 2023	21216066	vCenter Server 8.0b (Security fixes for Photon OS)	apache-ant	1.10.12-1.ph3	CVE-2020-11979
					CVE-2017-1000487
					CVE-2022-24839
			apache-tomcat	8.5.78-2.ph3	CVE-2022-29885
			c-ares	1.16.1-2.ph3	CVE-2020-8277
			cifs-utils	6.8-4.ph3	CVE-2022-27239
			containerd	1.4.13-5.ph3	CVE-2022-23648
			curl	7.86.0-2.ph3	CVE-2022-32207
					CVE-2022-22576
					CVE-2022-27782
					CVE-2022-27781

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
					CVE-2022-27775
					CVE-2021-22946
					CVE-2021-22926
					CVE-2020-8286
					CVE-2020-8285
					CVE-2022-42915
					CVE-2022-42916
			glibc	2.28-22.ph3	CVE-2022-23219
					CVE-2022-23218
					CVE-2021-33574
					CVE-2021-35942
					CVE-2021-3999
					CVE-2021-3998
					CVE-2021-43396
					CVE-2021-38604
					CVE-2021-3326
			libarchive	3.3.3-8.ph3	CVE-2021-31566
					CVE-2021-23177
			libtasn1	4.14-2.ph3	CVE-2021-46848
			libxml2	2.9.11-8.ph3	CVE-2022-2309
			libxslt	1.1.34-1.ph3	CVE-2021-30560
			linux	4.19.264-6.ph3	CVE-2022-2977
					CVE-2022-39842
					CVE-2022-2938
					CVE-2022-20368
					CVE-2022-2327
					CVE-2022-34918
					CVE-2022-32250
					CVE-2022-1786
					CVE-2022-29581
					CVE-2022-30594
					CVE-2022-1652
					CVE-2022-27666
					CVE-2022-1011
					CVE-2022-24958
					CVE-2021-4197
					CVE-2021-4037
					CVE-2022-36946
			nghttp2	1.41.0-2.ph3	CVE-2015-8659
					CVE-2020-11080

Release Date	Build Number	Patch Name	Affect Package	New Package Version	CVEs Addressed
			openldap	2.4.57-3.ph3	CVE-2022-29155
			openssl	1.0.2ze-3.ph3	CVE-2022-2068
					CVE-2022-1292
					CVE-2022-0778
					CVE-2021-23840
			pcre	8.44-2.ph3	CVE-2019-20838
			postgresql	10.22-1.ph3	CVE-2022-2625
			ruby	2.5.8-4.ph3	CVE-2020-10663
					CVE-2020-25613
			runc	1.1.3-3.ph3	CVE-2022-29162
			sysstat	12.7.1-1.ph3	CVE-2022-39377

VMware Host Client 2.5.0 Release Notes

This document contains the following sections

- [Introduction](#)
- [What's in the Release Notes](#)
- [Resolved Issues](#)
- [Known Issues](#)

Introduction

VMware vSphere 8.0 | 11 OCT 2022
 VMware Host Client 2.5.0 | 11 OCT 2022 | Build 20374953
 Check for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

VMware Host Client Overview

The VMware Host Client is an HTML5-based client that is used to connect to and manage single ESXi hosts. It can be used to perform administrative tasks to manage host resources such as virtual machines, networking, and storage. The VMware Host Client can also be helpful for troubleshooting individual virtual machines or hosts if vCenter Server and the vSphere Client are unavailable.

For more information, see the [vSphere 8.0 Documentation](#).

What's New in This Release

- Support for virtual machine hardware version 20
- Improvements of the user interface to comply with the accessibility standards
- Ability to change the user interface theme
- User interface localization available for Italian language
- Customizable login banner screen of the ESXi host

Previous Releases

- VMware Host Client 1.37.0
- VMware Host Client 1.34.8
- VMware Host Client 1.34.4
- VMware Host Client 1.34.0
- VMware Host Client 1.33.3
- VMware Host Client 1.30.0
- VMware Host Client 1.25.0
- VMware Host Client 1.21
- VMware Host Client 1.8
- VMware Host Client 1.4
- VMware Host Client 1.2
- VMware Host Client 1.0

Resolved Issues

The VMware Host Client supports direct SSH to the ESXi host within a browser

When you log in to the VMware Host Client with Google Chrome, you can download and install the **Secure Shell** plug-in. As a result, when you navigate to **Actions > SSH Console**, the terminal window appears and you have direct SSH access to the ESXi host.

This issue was resolved in this release.

Known Issues

The **Current reading** parameter of the Fan Redundancy health sensor displays **Redundancy lost**

Although the system fan health sensor is in green status, the **Current reading** parameter displays **Redundancy lost**. When you run the ESXCLI command `esxcli hardware ipmi sdr list`, the **Computed reading** returns a **Fully Redundant** value. As both **Current reading** and **Computed reading** parameters are the same, the **Redundancy lost** value is false.

Perform the following steps to reproduce this issue:

1. Install an ESXi 8.0 guest operating system.
2. Log in to the VMware Host Client.
3. Click **Monitor** in the VMware Host Client inventory.
4. On the **Hardware tab**, click **System sensors**.
5. Search for Fan sensors.

Result: **Current reading** column for **System Board 1 Fan Redundancy 0** displays **Redundancy lost**.

Workaround: None

You might not change the Load balancing setting of NIC teaming

The Load balancing setting of NIC teaming might not change when you select the **Inherit from vSwitch** option.

Perform the following steps to reproduce the issue:

1. Open Google Chrome and log in to the VMware Host Client.
2. Click **Networking** in the VMware Host Client inventory and click **Management Network**.
3. Click **Edit settings**.

Result: The **Edit port group - Management Network** window appears.

4. Expand **NIC teaming**.
 - a. From the **Load balancing** drop-down menu, select **Inherit from vSwitch** and click **Save**.
5. Complete steps 1-4.
6. The **Load balancing** option remains unchanged.

Workaround:

Override the **Policy** with the same settings as vSwitch, without using **Inherit from vSwitch**.

Do not edit **configstore**.

ESXi Update and Patch Release Notes

VMware ESXi 8.0 Update 3g Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 3g | 29 JUL 2025 | ISO Build 24859861
Check for additions and updates to these release notes.

What's New

- This release contains bugfixes as listed in the Resolved Issues section.

Patches Contained in This Release

VMware ESXi 8.0 Update 3g

Build Details

VMware vSphere Hypervisor (ESXi ISO) image

Download Filename:	VMware-VMvisor-Installer-8.0U3g-24859861.x86_64.iso
Build:	24859861
Download Size:	629 MB
SHA256 checksum:	d8baeb93aa81d8243c841f451a42571ed70cf3a48c2169585e6a476ecfd00113
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

VMware vSphere Hypervisor (ESXi) Offline Bundle

Download Filename:	VMware-ESXi-8.0U3g-24859861-depot.zip
---------------------------	---------------------------------------

Build:	24859861
Download Size:	1020.3 MB
SHA256 checksum:	c8c66780ffef1faaa24f21c9bef994c11b82c493d365d23621a8805ad3d9b9e5
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.3-0.79.24859861	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.79.24859861	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.79.24859861	Bugfix	Critical
Intel NVME Driver with VMD Technology	Intel-Volume-Mgmt-Device_3.0.0.1010-13vmw.803.0.79.24859861	Bugfix	Critical
VMware NVMe over TCP Driver	VMware-NVMeoF-TCP_1.0.1.36-1vmw.803.0.79.24859861	Bugfix	Critical
MMC/SD HCI Driver	VMware-SDHCI-Driver_1.0.3-7vmw.803.0.79.24859861	Bugfix	Critical
ESXi component with core ESXi VIBs	ESXi_8.0.3-0.76.24853260	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.76.24853260	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.76.24853260	Security	Critical
ESXi Tools Component	VMware-VM-Tools_12.5.3.24819442-24853260	Security	Critical

Rollup Bulletins

These rollup bulletins contain the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U3g-24859861	Bugfix	Critical	Security and Bugfix image
ESXi80U3sg-24853260	Security	Critical	Security only image

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U3g-24859861-standard
ESXi-8.0U3g-24859861-no-tools
ESXi-8.0U3sg-24853260-standard

ESXi-8.0U3sg-24853260-no-tools

ESXi Images

Name and Version	Release Date	Category	Detail
ESXi 8.0 U3g - 24859861	29 JUL 2025	General	Security and Bugfix image
ESXi 8.0 U3sg - 24853260	29 JUL 2025	Security	Security only image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

See these KBs for details on the authenticated online and offline download: [390121](#), [390098](#), and [389276](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually download the patch offline bundle ZIP file or see the KB [How to download ESXi ISO image for all releases including patch updates](#).

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues**ESXi_8.0.3-0.79.24859861**

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.79.24859861 - VMware_bootbank_vsan_8.0.3-0.79.24859861 - VMware_bootbank_infravisor_8.0.3-0.79.24859861 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.79.24859861 - VMware_bootbank_crx_8.0.3-0.79.24859861 - VMware_bootbank_bmcad_8.0.3-0.79.24859861 - VMware_bootbank_vsanhealth_8.0.3-0.79.24859861 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.79.24859861 - VMware_bootbank_gc_8.0.3-0.79.24859861 - VMware_bootbank_vdfs_8.0.3-0.79.24859861 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.79.24859861 - VMware_bootbank_esx-base_8.0.3-0.79.24859861 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.79.24859861 - VMware_bootbank_esx-xserver_8.0.3-0.79.24859861 - VMware_bootbank_vds-vsip_8.0.3-0.79.24859861

	• VMware_bootbank_cpu-microcode_8.0.3-0.79.24859861 • VMware_bootbank_esxio_8.0.3-0.79.24859861 • VMware_bootbank_clusterstore_8.0.3-0.79.24859861 • VMware_bootbank_esxio-base_8.0.3-0.79.24859861 • VMware_bootbank_trx_8.0.3-0.79.24859861 • VMware_bootbank_gc-esxio_8.0.3-0.79.24859861 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.79.24859861 • VMware_bootbank_bmc-al-esxio_8.0.3-0.79.24859861 • VMware_bootbank_vcls-pod-crx_8.0.3-0.79.24859861 • VMware_bootbank_native-misc-drivers_8.0.3-0.79.24859861
PRs Fixed	3511297, 3523908, 3523307, 3453391, 3502007, 3523288, 3496295, 3505543, 3523292, 3521243, 3549866, 3523289, 3526619, 3505417, 3547311, 3508977, 3514821, 3511112, 3525846, 3522258, 3522712, 3485301, 3522487, 3525785, 3527704, 3535747, 3524945, 3526457, 3499300, 3519342, 3485371, 3493836, 3469074, 3505628, 3500545, 3514647, 3516689, 3421841, 3506956, 3500400, 3517989, 3480453, 3505426, 3026630, 2916365, 3274848, 3033184, 3288124, 3273064, 3281300, 3151016, 3271023, 3040952
CVE numbers	N/A

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

This patch updates the esxio-combiner, vsan, infravisor, esxio-dvfilter-generic-fastpath, crx, bmc-al, vsanhealth, drivervm-gpu-base, gc, vdfs, native-misc-drivers-esxio, esx-base, esx-dvfilter-generic-fastpath, esx-xserver, vds-vsip, cpu-microcode, esxio, clusterstore, esxio-base, trx, gc-esxio, esxio-combiner-esxio, bmc-al-esxio, vcls-pod-crx, and native-misc-drivers VIBs.

This patch resolves the following issues:

PR 3511297: You see a no lock error on a vSAN NFS v3 volume

Internal network timeout issues might cause a no lock error on a vSAN NFS v3 volume. In the NFS client debug kernel logs, you see messages such as:

```
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334941] NFS: flock(/file1, t=1, fl=82).
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334945] lockd: nsm_monitor(acloud-fs11.kh.asegroup.com)
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334947] lockd: call procedure 2 on acloud-fs11.kh.asegroup.com.
Jan 13 14:27:44 khocpnfs01 kernel: [170419.335958] lockd: server returns status 2
Jan 13 14:27:44 khocpnfs01 kernel: [170419.335964] lockd: clnt proc returns -37.
```

This issue is resolved in this release.

PR 3523908: Virtual machines cannot boot when configured with a PCI passthrough device

VMs configured with a PCI passthrough device might consistently fail to boot, removing the device from its configuration and from the PCI device list of the ESXi host.

This issue is resolved in this release.

PR 3523307: A timing issue might cause vSAN ESA clusters to show as non-compliant with the SPBM policy

When you deploy a vSAN ESA cluster, a timing issue might cause the vSAN in-memory state for one of its modules to fail to initialize with the appropriate vSAN ESA configuration. As a result, the cluster is incorrectly identified as vSAN OSA and its objects show as non-compliant with the SPBM policy.

This issue is resolved in this release.

PR 3453391: You cannot create new virtual machines in vSAN clusters due to a networking issue

Due to a network issue between the storage cluster and a vSAN witness host, the Cluster Level Object Manager (CLOM) might create new witness components but the DOM does not delete the existing, which leads to exceeding the limit of witness component objects. As a result, creating new virtual machines on such vSAN clusters fails.

This issue is resolved in this release.

PR 3502007: Activation of TLS 1.3 in ClusterStore does not persist across ESXi host reboots

The `agent_settings` configuration on ESXi that controls whether TLS 1.3 is active on ClusterStore, which keeps the cluster-specific configuration data, reverts to the default of not active after a reboot of the ESX host.

This issue is resolved in this release.

PR 3523288: A virtual machine might fail to power on if the .vmx file is changed in the backend

In certain scenarios, when a VM is restored from a backup, applications external to the ESXi host might delete and recreate the .vmx file while the ESXi nfs41client is still accessing the file handle of the old .vmx file. As a result, a stale handle error causes the power-on of the VM to fail.

This issue is resolved in this release.

PR 3496295: When you use the suspend and resume feature on a virtual machine with Kubernetes control plane container configured, the VM might take long to respond or reset

During suspend and resume of a VM, the device state of a Kubernetes Control Plane container might change after a checkpoint and cause an inconsistent state of the device. As a result, the guest driver might reset the device.

This issue is resolved in this release.

PR 3505543: If you use ESXi host keys to encrypt virtual machines, the hostd service might fail

When you use an ESXi host key to encrypt a VM, if you delete the VM or change the key, the removal of the original host key might trigger an exception that can cause the hostd service to fail. As a result, the ESXi host temporarily disconnects from vCenter and you might not be able to manage VMs on the host if hostd does not restart after the failure.

This issue is resolved in this release.

PR 3523292: An ESXi host might fail with a purple diagnostic screen during migration of virtual machines from a Tintri NFSv4.1 datastore

During the migration of VMs by using vSphere vMotion from a Tintri NFSv4.1 datastore, the ESXi host might fail with a purple diagnostic screen. This issue is specific to how Tintri NFSv4.1 servers handle NULL values in certain scenarios.

This issue is resolved in this release.

PR 3521243: You cannot use public/private key pair to connect with SSH by using HTTP PUT

You cannot use authorized public/private keys to connect to ESXi hosts with SSH by using ConfigManager APIs (HTTP PUT).

This issue is resolved in this release.

PR 3549866: vSphere Virtual Volumes daemon fails after converting a VASA provider from external to internal

Due to an issue with handling IDs during an unbind operation for some vSphere API for Storage Awareness (VASA) providers, in some tasks such as converting a VASA provider from external to internal, the vSphere Virtual Volumes daemon vvolld might fail as it cannot map unique IDs.

This issue is resolved in this release. With the fix, vvolld gracefully handles such conditions and stops the unbind operation with message in the logs instead of failing.

PR 3523289: ClusterStore services might crashloop due to excessive memory usage and lead to network overload

A high default setting of a parameter in the ClusterStore capability might cause the service to consume more memory than necessary and reach the memory limit over time. As a result, ClusterStore services might crashloop and cause network overload as it attempts to communicate from one ESXi host to another during each startup.

This issue is resolved in this release.

PR 3526619: The authentication service of the Distributed Key Value Store, or ClusterStore, might cause DNS traffic overload

An issue with the authentication service of the Distributed Key Value Store, also called ClusterStore, that might keep trying to connect to a group of ESXi hosts even when that ClusterStore is no longer part of that group might cause excessive network utilization and eventually overload of certain infrastructure such as DNS.

This issue is resolved in this release. The fix implements a rate-limiter to reduce the retry frequency and network overhead.

PR 3505417: IF-MIB SNMP data intermittently reports hidden ports DVFilter Coalesce Portset and pps

When using third-party tools, the IF-MIB SNMP data might intermittently report hidden portsets, such as pps and DVFilter Coalesce Portset, even when they are not part of the ESXi host configuration.

This issue is resolved in this release.

PR 3547311: After a syslog configuration change, virtual machines lose connectivity to vSAN iSCSI targets

Configuring a remote syslog loghost might remove all non-syslog dynamic firewall rules, leading to service disruption such as VMs losing connection to vSAN iSCSI targets.

This issue is resolved in this release.

PR 3508977: During a snapshot operation of a VM with IOFilter attached, the VMX might fail due to an unresponsive thread

In very rare situations, once in 2000 snapshot operations, a race condition might occur in the I/O filter infrastructure cause the I/O filter watchdog to stop the VMX service and generate a VMX core dump.

This issue is resolved in this release.

PR 3514821: After a reboot of ESXi hosts, you see alarms in the vSphere Client that some uplinks are down

Due to some limitations in the VMkernel Observation (VOB) and hostd services, uplink up and down events might not be sent to vCenter in the correct order. As a result, when a task such as rebooting ESXi hosts completes, in the vSphere Client you see false alerts that some uplinks are down.

This issue is resolved in this release.

PR 3511112: After running the Linux fstrim command on a virtual machine, the ESXi host fails with a purple diagnostics screen

A VMFS heap memory exhaustion might cause memory access violations when clearing the pointer block cache. As a result, when you run the Linux `fstrim` command on a virtual machine, the ESXi host might fail with a purple diagnostics screen with a `PB3Manager` error in the backtrace.

This issue is resolved in this release.

For vSphere Quick Boot support in ESXi 8.0 Update 3g, see the list for each vendor in the Additional Information section of the [Understanding ESXi Quick Boot Compatibility](#) KB.

PR 3525846: Virtual machines with fixed passthrough hot-plugging enabled might fail due to misallocation of interrupts

For VMs with fixed passthrough hot-plugging enabled, after subsequent hot-adds of passthrough NVMe drives, the functionality might fail because interrupts might not be properly allocated.

This issue is resolved in this release.

PR 3522258: Mounting shares might fail due to rsyslogd CPU usage at 100%

During a restart of the rsyslogd service, CPU usage might reach 100% and you temporarily might not be able to mount shares.

This issue is resolved in this release.

PR 3522712: ESXi host reboot stops due to indefinite delay in the shutdown of the vSAN cluster monitoring, membership, and directory service (CMMDS)

When rebooting an ESXi host, network conditions can cause the vSAN CMMDS shutdown process to delay indefinitely.

This issue is resolved in this release.

PR 3485301: You see a delay in the power-on of vGPU virtual machines

You might see a delay in the power-on of virtual machines with NVIDIA vGPU devices and if many such tasks accumulate in vCenter, the vpxd service might be overloaded and vCenter becomes intermittently unresponsive. A restart of the vpxd service restores vCenter responsiveness.

This issue is resolved in this release.

PR 3522487: A routing issue with the Cisco Application Centric Infrastructure (ACI) switch might cause loss of vSAN traffic

To override normal IP routing, vSAN binds all outgoing traffic on TCP sockets and traffic runs over a vmkernel adapter dedicated to vSAN. In some cases, the vSAN binding might not be set or a TCP socket might be missing, so packets are sent with the normal IP routing. In such cases, Cisco ACI devices might change the routing table based on the incoming packets, which in turn causes traffic loss on the vSAN cluster.

This issue is resolved in this release.

PR 3525785: Excessive logs in the nsxdavim.log might cause an overload of the syslog server

Frequent updates from hostd on an NSX distributed virtual port group, such as add, update, and delete, might cause excessive logs in the `nsxdavim.log`, which in some cases might lead to an overload of the syslog server.

This issue is resolved in this release.

PR 3527704: ESXi host server firmware might not automatically make the PCIe Atomic Operations feature available to passthrough devices

GPU workloads on a virtual machine with AMD GPU VMDirectPath I/O might fail or hang because the ESXi host server firmware might not automatically activate PCIe atomic operations for AMD GPU devices.

This issue is resolved in this release. The fix adds the VMX configuration option `pciPassthru.exposeAtomicOps`, which you can set to `TRUE` to make sure VMs can use AMD GPUs as a passthrough device.

PR 3535747: On debug builds, an ESXi host might fail with a purple diagnostic screen due to a race condition

On debug builds, a race condition between the power-off task on a virtual machine and ESXi memory scrubbing might cause the ESXi host to fail with a purple diagnostic screen and the following error:

```
ASSERT bin/bora/vmkernel/private/_virtual_includes/hdrs/mpage2M.h:264
```

The issue does not affect release builds.

This issue is resolved in this release.

PR 3524945: Adding new disks external to the home directory of a virtual machine might fail with an error for reaching the security policy limit

When a disk external to the home directory of a VM is added or removed, the security domain policy of the VM updates to give or revoke access to the external directory. It is possible that if an error occurs during disk load, to avoid issues if it is resolved, external directories are no longer removed from the security policy, even after a resolution. In time, the rules in the security policy might reach the max limit. As a result, adding new disks external to the home directory of a virtual machine fails and in the `sandboxd.log` you see errors such as:

```
2025-05-11T08:14:55.306Z Er(163) sandboxd[2099161]: [Originator@6876
sub=SecurityDomainsImpl] AddUpdateFileRules: Failed to add file policies to domain :X:
world ID :Y:Argument list too long
```

This issue is resolved in this release. The fix makes sure that once a disk load issue is resolved, external directories are removed from the security policy. If a VM is already facing the issue, either restart the VM or migrate it to another ESXi host to trigger a regeneration of the security policy.

PR 3526457: You see multiple resets of a virtual machine in case of a machine check exception (MCE) error

If an uncorrectable memory error occurs on the guest OS of a virtual machine, memory access on the error page either by the guest or the ESXi hypervisor triggers a machine check exception. In such cases, the hypervisor might try to inject a virtual machine check error to the guest, which leads to guest OS failure. Depending on the guest OS, you might see multiple resets of the VM or it might become unresponsive.

This issue is resolved in this release.

PR 3499300: Virtual machine migration fails during upgrade of NSX or Distributed Firewall (DFW) reconfiguration

If the destination host of a VM migration operation is in the process of an upgrade of NSX or DFW reconfiguration, VM migrations to that host might fail.

This issue is resolved in this release.

PR 3519342: ESXi host fails with purple diagnostic screen during unmap operations on NVMe 2.0 devices due to misinterpreted unmap limits

Incorrect handling of the `maxUnmapDescriptorCount` and `maxUnmapLbaCount` values, which might be set as 0 by an NVMe target, can cause ESXi hosts to fail with a purple diagnostic screen when performing UNMAP operations on NVMe 2.0 devices. The NVMe 2.0 specification defines 0 in the `maxUnmapDescriptorCount` and `maxUnmapLbaCount` fields as indicating no limits on UNMAP operations, but ESXi reads them as indicating lack of UNMAP support. As a result, device-level UNMAP operation exits prematurely, leading to the failure.

This issue is resolved in this release.

PR 3485371: The VMX service might unexpectedly end replication-based snapshot operation due to a storage system failure

A storage system failure during a replication-based snapshot operation can lead to a logic error in the data protection component, potentially resulting in invalid memory access. As a result, the VMX service might fail and end the task.

This issue is resolved in this release.

PR 3493836: Virtual machine fails due to Signal 11, segmentation violation, during power on

In rare cases, a virtual machine might fail with Signal 11 error due to an issue with initializing VM statistics during power on.

This issue is resolved in this release.

PR 3469074: NVME over TCP adapter occasionally fails to reconnect to some storage array ports during storage array upgrade and controller failover

Controllers on an ESXi host might consistently fail to come online and require a system reboot to restore their operational state due to intermittent failures of the NVME over TCP adapter to reconnect to some storage array ports during storage array upgrade and controller failover. As a result, you see VMs on such hosts becoming unresponsive.

This issue is resolved in this release.

PR 3505628: During upgrade of ToR switches, the teamPolicyUpDelay capability of the "Route based on physical NIC load" teaming policy might not work consistently

If you set up the "Route based on physical NIC load" teaming policy to provide a window for ToR switches upgrade by using the teamPolicyUpDelay capability, in some cases this capability might not work. As a result, you might see a temporary network outage as the physical network is not ready.

This issue is resolved in this release.

PR 3500545: You cannot apply a vSAN license on 4Kn Software Emulated (SWE) disks due to insufficient capacity

When a vSAN cluster claims a 4Kn SWE disk, the capacity license takes into account the logical block size instead of the physical block size. As a result, the capacity value is much higher than the total size of disks claimed and in some cases, you might not be able to apply the vSAN license to such devices.

This issue is resolved in this release. The fix uses the logical block size and logical block count in pair for capacity calculation.

PR 3514647: You do not see an option to select more than 8 GPUs for vGPU device groups

When setting up a vGPU VM, you do not see an option to select more than 8 GPUs for vGPU device groups. For systems such as NVIDIA HGX, there could be more than 8 physical GPUs with corresponding device groups.

This issue is resolved in this release. The fix adds the option to select 16 GPUs.

PR 3516689: After upgrading to ESXi 8.0 Update 3 or later, PCIe devices are reported as non-passthrough capable or SR-IOV cannot be activated

Due to overly strict sanity checks, after upgrading your system to ESXi 8.0 Update 3 or later, some PCIe devices might be reported as non-passthrough capable or SR-IOV cannot be activated.

This issue is resolved in this release.

PR 3421841: ESXi fails to boot on systems with more than 512 logical CPUs in one NUMA node

Prior to ESXi 8.0 Update 8g, ESXi fails to boot on systems with more than 512 logical processors, or hyperthreads, in one NUMA node. Such systems are very unusual, but configuring the NUMA Nodes per Socket BIOS setting to 0 (NPS0) and deactivating the CCX as NUMA Domain on AMD system can cause all logical processors to be in a single NUMA node.

This issue is resolved in this release.

PR 3506956: You do not see any vGPU device groups for NVIDIA HGX H100 systems

The vGPU device group filtering code does not ensure that each NIC which is a candidate for a NIC-to-GPU device group has the required SR-IOV virtual function active. For NVIDIA HGX H100 systems this issue can cause all vGPU device groups to be ignored.

This issue is resolved in this release.

PR 3500400: Booting virtual machines with vGPU on an NVIDIA Blackwell GPU fails

vGPU VMs that reference a physical NVIDIA Blackwell GPU cannot power on because the vmx process fails.

This issue is resolved in this release.

PR 3517989: ESXi performance degrades with Hitachi NVMe-oF storage target

When using the NVMe/TCP protocol, creation and deployment of virtual machines on datastores backed by a NVMe namespace on Hitachi NVMe-oF target might be slow. when using the default settings.

This issue is resolved in this release.

PR 3480453: Virtual machine vNIC fails to activate after a vSphere vMotion operation

When system memory is highly fragmented, a VM vNIC might fail to allocate continuous physical memory and the vNIC cannot activate after a vSphere vMotion operation. As a result, VMs are randomly going into unresponsive state after successfully migrating to another ESXi host.

This issue is resolved in this release. The fix avoids allocation of continuous physical memory during vNIC activation.

PR 3505426: You might see multiple hardware sensor alerts due to a Baseboard Management Controller (BMC) issue

An underlying BMC issue might cause ESXi hardware health monitoring to produce multiple alerts for hardware sensors on ESXi hosts, which you see in the **vSphere Client under Monitor > Tasks and Events**.

This issue is resolved in this release. The fix logs raw sensor data directly from IPMI to provide a data point to debug possible BMC issues that cause false alerts, but cannot prevent BMC issues to occur.

PR 3026630: You see error messages when trying to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

This issue is resolved in this release.

PR 2916365: VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

This issue is resolved in this release.

PR 3274848: Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

This issue is resolved in this release.

PR 3033184: If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade.`

This issue is resolved in this release.

PR 3288124: You cannot update to ESXi 8.0 Update 2b or later by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

This issue is resolved in this release.

PR 3273064: If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

This issue is resolved in this release.

PR 3281300: When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

This issue is resolved in this release.

PR 3151016: Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a vSphere Lifecycle Manager image in the vSphere Client and then navigate away.
2. You trigger a page refresh.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session.

If you use API, firmware compliance details are missing from the response.

This issue is resolved in this release.

PR 3271023: TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

This issue is resolved in this release.

PR 3040952: In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

This issue is resolved in this release.

esx-update_8.0.3-0.79.24859861

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesx_8.0.3-0.79.24859861 - VMware_bootbank_esx-update_8.0.3-0.79.24859861
PRs Fixed	3499150
CVE numbers	N/A

Updates the `loadesx` and `esx-update` VIBs to resolve the following issue:

PR 3499150: NSX deployment fails during nsx-datapath VIB installation due to an out-of-memory condition

Insufficient sizing of the memory resource pool of one of the services during the installation of VIBs in an NSX deployment might cause the deployment to fail. In the `vmkernel.log` you see an error such as:

```
2025-02-11T19:39:18.172Z In(182) vmkernel: cpu68:2107749)Admission failure in path: host/
vim/vmvisor/settingsd-task-forks/python.2107554:python.2107764:uw.2107764
```

This issue is resolved in this release.

esxio-update_8.0.3-0.79.24859861

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesxio_8.0.3-0.79.24859861 - VMware_bootbank_esxio-update_8.0.3-0.79.24859861
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

Intel-Volume-Mgmt-Device_3.0.0.1010-13vmw.803.0.79.24859861

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMW_bootbank_iavmd_3.0.0.1010-13vmw.803.0.79.24859861
PRs Fixed	3518255
CVE numbers	N/A

Updates the `iavmd` VIB to resolve the following issue:

PR 3518255: During booting or when the Intel Volume Management Device (VMD) driver first loads, NVMe devices might fail in the enumeration of namespaces

If an NVMe drive has 512 namespaces, the VMD driver runs out of memory when trying to enumerate them all.

This issue is resolved in this release. ESXi 8.0 Update 3g contains VMD driver version 3.0.0.1010-13 that provides the fix. If you do not update your system to ESXi 8.0 Update 3g, you can just update the VMD driver to version 3.0.0.1010-13.

VMware-NVMeoF-TCP_1.0.1.36-1vmw.803.0.79.24859861

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMW_bootbank_nvmetcp_1.0.1.36-1vmw.803.0.79.24859861 - VMW_bootbank_nvmetcp-esxio_1.0.1.36-1vmw.803.0.79.24859861
PRs Fixed	3517150
CVE numbers	N/A

Updates the `nvmetcp` VIB to resolve the following issue:

PR 3517150: When data digests on NVMe over TCP datastores are active, you might see frequent connection drops and performance degradation

Various applications can modify the write I/O data that is in transit between ESXi hosts before the `nvmetcp` driver completes a task. When data digests on NVMe over TCP datastores are active, if data is modified after the digest is calculated within the `nvmetcp` driver, the target receives the modified I/O data. This leads to a digest mismatch, which in turn causes the controller connection to drop, the controller to reset, and ultimately, I/O performance degradation.

This issue is resolved in this release.

VMware-SDHCI-Driver_1.0.3-7vmw.803.0.79.24859861

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMW_bootbank_vmksdhci_1.0.3-7vmw.803.0.79.24859861 - VMW_bootbank_vmksdhci-esxio_1.0.3-7vmw.803.0.79.24859861
PRs Fixed	3518979
CVE numbers	N/A

Updates the `vmksdhci` and `vmksdhci-esxio` VIBs to resolve the following issue:

PR 3518979: ESX hosts on dual-DPU systems might fail with a purple diagnostic screen due to an interrupt storm caused by the MultiMediaCard (MMC) controller

Interrupt storm caused by the MMC controller on a dual-DPU systems completely occupies one or more CPUs for long periods, which might lead to failure of ESX hosts with a purple diagnostic screen. In the kernel log, you see errors such as:

```
2025-04-30T01:26:26.364Z cpu2:528386) 0x4538e841b530: [0x4200374c98a0] VMKExcEntry@vmkernel#nover+0x6c stack: 0x1
2025-04-30T01:26:26.364Z cpu2:528386) 0x4538e841b530: [0x4200374c98a0] VMKExcEntry@vmkernel#nover+0x6c stack: 0x1
2025-04-30T01:26:26.371Z cpu2:528386) 0x4538e841b660: [0x420036f50d40] Heap_Free@vmkernel#nover+0x26c stack: 0x0
2025-04-30T01:26:26.377Z cpu2:528386) 0x4538e841b6b0: [0x420037dc1598] sleep_list_remove@com.vmware.vmkbs-
d#1+0x104 stack: 0x0
2025-04-30T01:26:26.385Z cpu2:528386) 0x4538e841b700: [0x420037f66970] dwmmc_intr@ (vmksdhci) #<None>+0x200 stack:
0x0
2025-04-30T01:26:26.392Z cpu2:528386) 0x4538e841b730: [0x420036f78454] IntrCookieBH@vmkernel#nover+0x208 stack:
0x0
2025-04-30T01:26:26.398Z cpu2:528386) 0x4538e841b7e0: [0x420036f556e0] BH_DrainAndDisableInterrupts@vmker-
nel#nover+0x21c stack: 0x0
```

This issue rarely occurs in single-DPU systems.

This issue is resolved in this release.

ESXi_8.0.3-0.76.24853260

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_drivervm-gpu-base_8.0.3-0.76.24853260 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.76.24853260 - VMware_bootbank_vsan_8.0.3-0.76.24853260

	- VMware_bootbank_esxio-combiner_8.0.3-0.76.24853260 - VMware_bootbank_bmcad-esxio_8.0.3-0.76.24853260 - VMware_bootbank_trx_8.0.3-0.76.24853260 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.76.24853260 - VMware_bootbank_crx_8.0.3-0.76.24853260 - VMware_bootbank_cpu-microcode_8.0.3-0.76.24853260 - VMware_bootbank_clusterstore_8.0.3-0.76.24853260 - VMware_bootbank_esxio_8.0.3-0.76.24853260 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.76.24853260 - VMware_bootbank_vcls-pod-crx_8.0.3-0.76.24853260 - VMware_bootbank_gc_8.0.3-0.76.24853260 - VMware_bootbank_vds-vsip_8.0.3-0.76.24853260 - VMware_bootbank_gc-esxio_8.0.3-0.76.24853260 - VMware_bootbank_infravisor_8.0.3-0.76.24853260 - VMware_bootbank_vsanhealth_8.0.3-0.76.24853260 - VMware_bootbank_esx-xserver_8.0.3-0.76.24853260 - VMware_bootbank_native-misc-drivers_8.0.3-0.76.24853260 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.76.24853260 - VMware_bootbank_bmcad_8.0.3-0.76.24853260 - VMware_bootbank_vdfs_8.0.3-0.76.24853260 - VMware_bootbank_esxio-base_8.0.3-0.76.24853260 - VMware_bootbank_esx-base_8.0.3-0.76.24853260
PRs Fixed	3492319, 3506965
CVE numbers	N/A

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching. This patch updates the `drivervm-gpu-base`, `esxio-dvfilter-generic-fastpath`, `vsan`, `esxio-combiner`, `bmcad-esxio`, `trx`, `esx-dvfilter-generic-fastpath`, `crx`, `cpu-microcode`, `clusterstore`, `esxio`, `native-misc-drivers-esxio`, `vcls-pod-crx`, `gc`, `vds-vsip`, `gc-esxio`, `infravisor`, `vsanhealth`, `esx-xserver`, `native-misc-drivers`, `esxio-combiner-esxio`, `bmcad`, `vdfs`, `esxio-base`, and `esx-base` VIBs.

This patch resolves the following issues:

ESXi 8.0 Update 3g includes the following AMD microcode:

ESXi builds provide microcode for most supported CPUs, and if the BIOS of your server has an older microcode version than the one provided with the ESXi build, ESXi attempts to load the newer version. Because of an AMD change in [microcode signing](#), ESXi might not be able to apply the latest microcode. In such a case, ESXi 8.0 Update 3g applies the newest microcode that it can and if the version is not the latest, ESXi issues a SysAlert such as `This machine's BIOS should be updated to pick up the latest microcode`. Updating to the latest BIOS is always best practice as usually the latest BIOS also carries the latest microcode.

Code Name	FMS	MCU Rev	MCU Date	Brand Names
Bulldozer	0x00600f12 (15/01/2)	0x0600063e	2/7/2018	Opteron 6200/4200/3200 Series
Piledriver	0x00600f20 (15/02/0)	0x06000852	2/6/2018	Opteron 6300/4300/3300 Series
Zen-Naples	0x00800f12 (17/01/2)	0x08001278	11/11/2024	EPYC 7001 Series

Code Name	FMS	MCU Rev	MCU Date	Brand Names
Zen2-Rome	0x00830f10 (17/31/0)	0x0830107c	12/18/2023	EPYC 7002/7Fx2/7Hx2 Series
Zen3-Milan	0x00a00f11 (19/01/1)	0x0a0011db	11/11/2024	EPYC 7003 Series
Zen3-Milan (old BIOS)	0x00a00f11 (19/01/1)	0x0a0011d5	2/23/2024	EPYC 7003 Series
Zen3-Milan-X	0x00a00f12 (19/01/2)	0x0a001244	11/11/2024	EPYC 7003X Series
Zen3-Milan-X (old BIOS)	0x00a00f12 (19/01/2)	0x0a001238	2/26/2024	EPYC 7003X Series
Zen4-Genoa	0x00a10f11 (19/11/1)	0x0a101154	11/12/2024	EPYC 9004 Series
Zen4-Genoa (old BIOS)	0x00a10f11 (19/11/1)	0x0a101148	2/23/2024	EPYC 9004 Series
Zen4-Genoa-X	0x00a10f12 (19/11/2)	0x0a10124f	11/12/2024	EPYC 9004X Series
Zen4-Genoa-X (old BIOS)	0x00a10f12 (19/11/2)	0x0a101248	2/26/2024	EPYC 9004X Series
Zen4-Bergamo/Siena	0x00aa0f02 (19/a0/2)	0x0aa00219	11/13/2024	EPYC 8004/9004 Series
Zen4-Bergamo/Siena (old BIOS)	0x00aa0f02 (19/a0/2)	0x0aa00215	2/28/2024	EPYC 8004/9004 Series

ESXi 8.0 Update 3g includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003901	12/12/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003901	12/12/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/ b)	0xbf	baseline	0x7002b01	12/12/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd000404	1/7/2025	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002d0	1/7/2025	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xbc	12/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x3c	12/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x56	12/11/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	baseline	0x2c0003e0	7/30/2024	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b000639	1/28/2025	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xfa	7/28/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x104	11/14/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x64	12/1/2024	Intel Xeon E-2300 Series
Granite Rapids SP R1S	0xa06d1 (06/ad/1)	0x20	baseline	0xa0000d1	2/7/2025	Intel Xeon 6500P/6700P/6900P Series
Granite Rapids AP/SP	0xa06d1 (06/ad/1)	0x95	baseline	0x10003a2	2/7/2025	Intel Xeon 6500P/6700P/6900P Series
Sierra Forest	0xa06f3 (06/af/3)	0x01	baseline	0x3000341	2/10/2025	Intel Xeon 6700E/6900E Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x12e	1/15/2025	Intel Xeon E-2400 Series
Raptor Lake E/HX/S	0xb0674 (06/b7/4)	0x32	baseline	0x12e	1/15/2025	Intel Xeon E-2400 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x210002a9	3/14/2025	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series

esx-update_8.0.3-0.76.24853260

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesx_8.0.3-0.76.24853260 - VMware_bootbank_esx-update_8.0.3-0.76.24853260
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesx` and `esx-update` VIBs.

esxio-update_8.0.3-0.76.24853260

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesxio_8.0.3-0.76.24853260 - VMware_bootbank_esxio-update_8.0.3-0.76.24853260
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

VMware-VM-Tools_12.5.3.24819442-24853260

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	No
Virtual Machine Migration or Shutdown Required	No

Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMware_locker_tools-light_12.5.3.24819442-24853260
PRs Fixed	3518336
CVE numbers	N/A

Updates the `locker_tools-light` VIB.

VMware Tools Bundling Changes in ESXi 8.0 Update 3g

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 3g:

- **windows.iso**: VMware Tools 12.5.3 supports 64-bit Windows versions - Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- **windows-x86.iso**: VMware Tools 12.4.7 supports 32-bit Windows versions.
- **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.5.3 Release Notes](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Upgrading VMware Tools](#)

ESXi-8.0U3g-24859861-standard

Profile Name	ESXi-8.0U3g-24859861-standard
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 29, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.79.24859861 - VMware_bootbank_vsan_8.0.3-0.79.24859861 - VMware_bootbank_infravisor_8.0.3-0.79.24859861 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.79.24859861 - VMware_bootbank_crx_8.0.3-0.79.24859861 - VMware_bootbank_bmcad_8.0.3-0.79.24859861

	• VMware_bootbank_vsanhealth_8.0.3-0.79.24859861 • VMware_bootbank_drivervm-gpu-base_8.0.3-0.79.24859861 • VMware_bootbank_gc_8.0.3-0.79.24859861 • VMware_bootbank_vdfs_8.0.3-0.79.24859861 • VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.79.24859861 • VMware_bootbank_esx-base_8.0.3-0.79.24859861 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.79.24859861 • VMware_bootbank_esx-xserver_8.0.3-0.79.24859861 • VMware_bootbank_vds-vsip_8.0.3-0.79.24859861 • VMware_bootbank_cpu-microcode_8.0.3-0.79.24859861 • VMware_bootbank_esxio_8.0.3-0.79.24859861 • VMware_bootbank_clusterstore_8.0.3-0.79.24859861 • VMware_bootbank_esxio-base_8.0.3-0.79.24859861 • VMware_bootbank_trx_8.0.3-0.79.24859861 • VMware_bootbank_gc-esxio_8.0.3-0.79.24859861 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.79.24859861 • VMware_bootbank_bmc-al-esxio_8.0.3-0.79.24859861 • VMware_bootbank_vcls-pod-crx_8.0.3-0.79.24859861 • VMware_bootbank_native-misc-drivers_8.0.3-0.79.24859861 • VMware_bootbank_loadesx_8.0.3-0.79.24859861 • VMware_bootbank_esx-update_8.0.3-0.79.24859861 • VMware_bootbank_loadesxio_8.0.3-0.79.24859861 • VMware_bootbank_esxio-update_8.0.3-0.79.24859861 • VMW_bootbank_iavmd_3.0.0.1010-13vmw.803.0.79.24859861 • VMW_bootbank_nvmetcp_1.0.1.36-1vmw.803.0.79.24859861 • VMW_bootbank_nvmetcp-esxio_1.0.1.36-1vmw.803.0.79.24859861 • VMW_bootbank_vmksdhci_1.0.3-7vmw.803.0.79.24859861 • VMW_bootbank_vmksdhci-esxio_1.0.3-7vmw.803.0.79.24859861 • VMware_locker_tools-light_12.5.3.24819442-24853260
PRs Fixed	3511297, 3523908, 3523307, 3453391, 3502007, 3523288, 3496295, 3505543, 3523292, 3521243, 3549866, 3523289, 3526619, 3505417, 3547311, 3508977, 3514821, 3511112, 3525846, 3522258, 3522712, 3485301, 3522487, 3525785, 3527704, 3535747, 3524945, 3526457, 3499300, 3519342, 3485371, 3493836, 3469074, 3505628, 3500545, 3514647, 3516689, 3421841, 3506956, 3500400, 3517989, 3480453, 3505426, 3026630, 2916365, 3274848, 3033184, 3288124, 3273064, 3281300, 3151016, 3271023, 3040952, 3499150, 3518255, 3517150, 3518979
Related CVE numbers	N/A

This patch updates the following issues:

PR 3511297: You see a no lock error on a vSAN NFS v3 volume

Internal network timeout issues might cause a no lock error on a vSAN NFS v3 volume. In the NFS client debug kernel logs, you see messages such as:

```
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334941] NFS: flock(/file1, t=1, fl=82).
```

```

Jan 13 14:27:41 khocpnfs01 kernel: [170416.334945] lockd: nsm_monitor(acloud-fs11.kh.asegroup.com)
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334947] lockd: call procedure 2 on acloud-fs11.kh.asegroup.com.
Jan 13 14:27:44 khocpnfs01 kernel: [170419.335958] lockd: server returns status 2
Jan 13 14:27:44 khocpnfs01 kernel: [170419.335964] lockd: clnt proc returns -37.

```

This issue is resolved in this release.

PR 3523908: Virtual machines cannot boot when configured with a PCI passthrough device

VMs configured with a PCI passthrough device might consistently fail to boot, removing the device from its configuration and from the PCI device list of the ESXi host.

This issue is resolved in this release.

PR 3523307: A timing issue might cause vSAN ESA clusters to show as non-compliant with the SPBM policy

When you deploy a vSAN ESA cluster, a timing issue might cause the vSAN in-memory state for one of its modules to fail to initialize with the appropriate vSAN ESA configuration. As a result, the cluster is incorrectly identified as vSAN OSA and its objects show as non-compliant with the SPBM policy.

This issue is resolved in this release.

PR 3453391: You cannot create new virtual machines in vSAN clusters due to a networking issue

Due to a network issue between the storage cluster and a vSAN witness host, the Cluster Level Object Manager (CLOM) might create new witness components but the DOM does not delete the existing, which leads to exceeding the limit of witness component objects. As a result, creating new virtual machines on such vSAN clusters fails.

This issue is resolved in this release.

PR 3453391: You cannot create new virtual machines in vSAN clusters due to a networking issue

Due to a network issue between the storage cluster and a vSAN witness host, the Cluster Level Object Manager (CLOM) might create new witness components but the DOM does not delete the existing, which leads to exceeding the limit of witness component objects. As a result, creating new virtual machines on such vSAN clusters fails.

This issue is resolved in this release.

PR 3521243: You cannot use public/private key pair to connect with SSH by using HTTP PUT

You cannot use authorized public/private keys to connect to ESXi hosts with SSH by using ConfigManager APIs (HTTP PUT).

This issue is resolved in this release.

PR 3518979: ESX hosts on dual-DPU systems might fail with a purple diagnostic screen due to an interrupt storm caused by the MultiMediaCard (MMC) controller

Interrupt storm caused by the MMC controller on a dual-DPU systems completely occupies one or more CPUs for long periods, which might lead to failure of ESX hosts with a purple diagnostic screen. In the kernel log, you see errors such as:

```

2025-04-30T01:26:26.364Z cpu2:528386) 0x4538e841b530:[0x4200374c98a0]VMKExcEntry@vmkernel#nover+0x6c stack: 0x1
2025-04-30T01:26:26.364Z cpu2:528386) 0x4538e841b530:[0x4200374c98a0]VMKExcEntry@vmkernel#nover+0x6c stack: 0x1
2025-04-30T01:26:26.371Z cpu2:528386) 0x4538e841b660:[0x420036f50d40]Heap_Free@vmkernel#nover+0x26c stack: 0x0
2025-04-30T01:26:26.377Z cpu2:528386) 0x4538e841b6b0:[0x420037dc1598]sleep_list_remove@com.vmware.vmkbs-
d#1+0x104 stack: 0x0
2025-04-30T01:26:26.385Z cpu2:528386) 0x4538e841b700:[0x420037f66970]dwmmc_intr@(vmksdhci)#<None>+0x200 stack:
0x0
2025-04-30T01:26:26.392Z cpu2:528386) 0x4538e841b730:[0x420036f78454]IntrCookieBH@vmkernel#nover+0x208 stack:
0x0
2025-04-30T01:26:26.398Z cpu2:528386) 0x4538e841b7e0:[0x420036f556e0]BH_DrainAndDisableInterrupts@vmker-
nel#nover+0x21c stack: 0x0

```

This issue rarely occurs in single-DPU systems.

This issue is resolved in this release.

PR 3517150: When data digests on NVMe over TCP datastores are active, you might see frequent connection drops and performance degradation

Various applications can modify the write I/O data that is in transit between ESXi hosts before the nvmetcp driver completes a task. When data digests on NVMe over TCP datastores are active, if data is modified after the digest is calculated within the nvmetcp driver, the target receives the modified I/O data. This leads to a digest mismatch, which in turn causes the controller connection to drop, the controller to reset, and ultimately, I/O performance degradation.

This issue is resolved in this release.

PR 3518255: During booting or when the Intel Volume Management Device (VMD) driver first loads, NVMe devices might fail in the enumeration of namespaces

If an NVMe drive has 512 namespaces, the VMD driver runs out of memory when trying to enumerate them all.

This issue is resolved in this release. ESXi 8.0 Update 3g contains VMD driver version 3.0.0.1010-13 that provides the fix. If you do not update your system to ESXi 8.0 Update 3g, you can just update the VMD driver to version 3.0.0.1010-13.

PR 3499150: NSX deployment fails during nsx-datapath VIB installation due to an out-of-memory condition

Insufficient sizing of the memory resource pool of one of the services during the installation of VIBs in an NSX deployment might cause the deployment to fail. In the `vmkernel.log` you see an error such as:

```
2025-02-11T19:39:18.172Z In(182) vmkernel: cpu68:2107749)Admission failure in path: host/
vim/vmvisor/settingsd-task-forks/python.2107554:python.2107764:uw.2107764
```

This issue is resolved in this release.

PR 3502007: Activation of TLS 1.3 in ClusterStore does not persist across ESXi host reboots

The `agent_settings` configuration on ESXi that controls whether TLS 1.3 is active on ClusterStore, which keeps the cluster-specific configuration data, reverts to the default of not active after a reboot of the ESX host.

This issue is resolved in this release.

PR 3523288: A virtual machine might fail to power on if the .vmx file is changed in the backend

In certain scenarios, when a VM is restored from a backup, applications external to the ESXi host might delete and recreate the .vmx file while the ESXi nfs41client is still accessing the file handle of the old .vmx file. As a result, a stale handle error causes the power-on of the VM to fail.

This issue is resolved in this release.

PR 3496295: When you use the suspend and resume feature on a virtual machine with Kubernetes control plane container configured, the VM might take long to respond or reset

During suspend and resume of a VM, the device state of a Kubernetes Control Plane container might change after a checkpoint and cause an inconsistent state of the device. As a result, the guest driver might reset the device.

This issue is resolved in this release.

PR 3505543: If you use ESXi host keys to encrypt virtual machines, the hostd service might fail

When you use an ESXi host key to encrypt a VM, if you delete the VM or change the key, the removal of the original host key might trigger an exception that can cause the hostd service to fail. As a result, the ESXi host temporarily disconnects from vCenter and you might not be able to manage VMs on the host if hostd does not restart after the failure.

This issue is resolved in this release.

PR 3523292: An ESXi host might fail with a purple diagnostic screen during migration of virtual machines from a Tintri NFSv4.1 datastore

During the migration of VMs by using vSphere vMotion from a Tintri NFSv4.1 datastore, the ESXi host might fail with a purple diagnostic screen. This issue is specific to how Tintri NFSv4.1 servers handle NULL values in certain scenarios.

This issue is resolved in this release.

PR 3549866: vSphere Virtual Volumes daemon fails after converting a VASA provider from external to internal

Due to an issue with handling IDs during an unbind operation for some vSphere API for Storage Awareness (VASA) providers, in some tasks such as converting a VASA provider from external to internal, the vSphere Virtual Volumes daemon vvolld might fail as it cannot map unique IDs.

This issue is resolved in this release. With the fix, vvolld gracefully handles such conditions and stops the unbind operation with message in the logs instead of failing.

PR 3523289: ClusterStore services might crashloop due to excessive memory usage and lead to network overload

A high default setting of a parameter in the ClusterStore capability might cause the service to consume more memory than necessary and reach the memory limit over time. As a result, ClusterStore services might crashloop and cause network overload as it attempts to communicate from one ESXi host to another during each startup.

This issue is resolved in this release.

PR 3526619: The authentication service of the Distributed Key Value Store, or ClusterStore, might cause DNS traffic overload

An issue with the authentication service of the Distributed Key Value Store, also called ClusterStore, that might keep trying to connect to a group of ESXi hosts even when that ClusterStore is no longer part of that group might cause excessive network utilization and eventually overload of certain infrastructure such as DNS.

This issue is resolved in this release. The fix implements a rate-limiter to reduce the retry frequency and network overhead.

PR 3505417: IF-MIB SNMP data intermittently reports hidden ports DVFilter Coalesce Portset and pps

When using third-party tools, the IF-MIB SNMP data might intermittently report hidden portsets, such as pps and DVFilter Coalesce Portset, even when they are not part of the ESXi host configuration.

This issue is resolved in this release.

PR 3547311: After a syslog configuration change, virtual machines lose connectivity to vSAN iSCSI targets

Configuring a remote syslog loghost might remove all non-syslog dynamic firewall rules, leading to service disruption such as VMs losing connection to vSAN iSCSI targets.

This issue is resolved in this release.

PR 3508977: During a snapshot operation of a VM with IOFilter attached, the VMX might fail due to an unresponsive thread

In very rare situations, once in 2000 snapshot operations, a race condition might occur in the I/O filter infrastructure cause the I/O filter watchdog to stop the VMX service and generate a VMX core dump.

This issue is resolved in this release.

PR 3514821: After a reboot of ESXi hosts, you see alarms in the vSphere Client that some uplinks are down

Due to some limitations in the VMkernel Observation (VOB) and hostd services, uplink up and down events might not be sent to vCenter in the correct order. As a result, when a task such as rebooting ESXi hosts completes, in the vSphere Client you see false alerts that some uplinks are down.

This issue is resolved in this release.

PR 3511112: After running the Linux `fstrim` command on a virtual machine, the ESXi host fails with a purple diagnostics screen

A VMFS heap memory exhaustion might cause memory access violations when clearing the pointer block cache. As a result, when you run the Linux `fstrim` command on a virtual machine, the ESXi host might fail with a purple diagnostics screen with a `PB3Manager` error in the backtrace.

This issue is resolved in this release.

For vSphere Quick Boot support in ESXi 8.0 Update 3g, see the list for each vendor in the Additional Information section of the [Understanding ESXi Quick Boot Compatibility](#) KB.

PR 3525846: Virtual machines with fixed passthrough hot-plugging enabled might fail due to misallocation of interrupts

For VMs with fixed passthrough hot-plugging enabled, after subsequent hot-adds of passthrough NVMe drives, the functionality might fail because interrupts might not be properly allocated.

This issue is resolved in this release.

PR 3522258: Mounting shares might fail due to `rsyslogd` CPU usage at 100%

During a restart of the `rsyslogd` service, CPU usage might reach 100% and you temporarily might not be able to mount shares.

This issue is resolved in this release.

PR 3522712: ESXi host reboot stops due to indefinite delay in the shutdown of the vSAN cluster monitoring, membership, and directory service (CMMDS)

When rebooting an ESXi host, network conditions can cause the vSAN CMMDS shutdown process to delay indefinitely.

This issue is resolved in this release.

PR 3485301: You see a delay in the power-on of vGPU virtual machines

You might see a delay in the power-on of virtual machines with NVIDIA vGPU devices and if many such tasks accumulate in vCenter, the `vpdx` service might be overloaded and vCenter becomes intermittently unresponsive. A restart of the `vpdx` service restores vCenter responsiveness.

This issue is resolved in this release.

PR 3522487: A routing issue with the Cisco Application Centric Infrastructure (ACI) switch might cause loss of vSAN traffic

To override normal IP routing, vSAN binds all outgoing traffic on TCP sockets and traffic runs over a vmkernel adapter dedicated to vSAN. In some cases, the vSAN binding might not be set or a TCP socket might be missing, so packets are sent with the normal IP routing. In such cases, Cisco ACI devices might change the routing table based on the incoming packets, which in turn causes traffic loss on the vSAN cluster.

This issue is resolved in this release.

PR 3525785: Excessive logs in the `nsxdavim.log` might cause an overload of the syslog server

Frequent updates from `hostd` on an NSX distributed virtual port group, such as add, update, and delete, might cause excessive logs in the `nsxdavim.log`, which in some cases might lead to an overload of the syslog server.

This issue is resolved in this release.

PR 3527704: ESXi host server firmware might not automatically make the PCIe Atomic Operations feature available to passthrough devices

GPU workloads on a virtual machine with AMD GPU VMDirectPath I/O might fail or hang because the ESXi host server firmware might not automatically activate PCIe atomic operations for AMD GPU devices.

This issue is resolved in this release. The fix adds the VMX configuration option `pciPassthru.exposeAtomicOps`, which you can set to `TRUE` to make sure VMs can use AMD GPUs as a passthrough device.

PR 3535747: On debug builds, an ESXi host might fail with a purple diagnostic screen due to a race condition

On debug builds, a race condition between the power-off task on a virtual machine and ESXi memory scrubbing might cause the ESXi host to fail with a purple diagnostic screen and the following error:

```
ASSERT bin/bora/vmkernel/private/_virtual_includes/hdrs/mpage2M.h:264
```

The issue does not affect release builds.

This issue is resolved in this release.

PR 3524945: Adding new disks external to the home directory of a virtual machine might fail with an error for reaching the security policy limit

When a disk external to the home directory of a VM is added or removed, the security domain policy of the VM updates to give or revoke access to the external directory. It is possible that if an error occurs during disk load, to avoid issues if it is resolved, external directories are no longer removed from the security policy, even after a resolution. In time, the rules in the security policy might reach the max limit. As a result, adding new disks external to the home directory of a virtual machine fails and in the `sandboxd.log` you see errors such as:

```
2025-05-11T08:14:55.306Z Er(163) sandboxd[2099161]: [Originator@6876
sub=SecurityDomainsImpl] AddUpdateFileRules: Failed to add file policies to domain :X:
world ID :Y:Argument list too long
```

This issue is resolved in this release. The fix makes sure that once a disk load issue is resolved, external directories are removed from the security policy. If a VM is already facing the issue, either restart the VM or migrate it to another ESXi host to trigger a regeneration of the security policy.

PR 3526457: You see multiple resets of a virtual machine in case of a machine check exception (MCE) error

If an uncorrectable memory error occurs on the guest OS of a virtual machine, memory access on the error page either by the guest or the ESXi hypervisor triggers a machine check exception. In such cases, the hypervisor might try to inject a virtual machine check error to the guest, which leads to guest OS failure. Depending on the guest OS, you might see multiple resets of the VM or it might become unresponsive.

This issue is resolved in this release.

PR 3499300: Virtual machine migration fails during upgrade of NSX or Distributed Firewall (DFW) reconfiguration

If the destination host of a VM migration operation is in the process of an upgrade of NSX or DFW reconfiguration, VM migrations to that host might fail.

This issue is resolved in this release.

PR 3519342: ESXi host fails with purple diagnostic screen during unmap operations on NVMe 2.0 devices due to misinterpreted unmap limits

Incorrect handling of the `maxUnmapDescriptorCount` and `maxUnmapLbaCount` values, which might be set as 0 by an NVMe target, can cause ESXi hosts to fail with a purple diagnostic screen when performing UNMAP operations on NVMe 2.0 devices. The NVMe 2.0 specification defines 0 in the `maxUnmapDescriptorCount` and `maxUnmapLbaCount` fields as indicating no limits on UNMAP operations, but ESXi reads them as indicating lack of UNMAP support. As a result, device-level UNMAP operation exits prematurely, leading to the failure.

This issue is resolved in this release.

PR 3485371: The VMX service might unexpectedly end replication-based snapshot operation due to a storage system failure

A storage system failure during a replication-based snapshot operation can lead to a logic error in the data protection component, potentially resulting in invalid memory access. As a result, the VMX service might fail and end the task.

This issue is resolved in this release.

PR 3493836: Virtual machine fails due to Signal 11, segmentation violation, during power on

In rare cases, a virtual machine might fail with Signal 11 error due to an issue with initializing VM statistics during power on.

This issue is resolved in this release.

PR 3469074: NVME over TCP adapter occasionally fails to reconnect to some storage array ports during storage array upgrade and controller failover

Controllers on an ESXi host might consistently fail to come online and require a system reboot to restore their operational state due to intermittent failures of the NVME over TCP adapter to reconnect to some storage array ports during storage array upgrade and controller failover. As a result, you see VMs on such hosts becoming unresponsive.

This issue is resolved in this release.

PR 3505628: During upgrade of ToR switches, the teamPolicyUpDelay capability of the "Route based on physical NIC load" teaming policy might not work consistently

If you set up the "Route based on physical NIC load" teaming policy to provide a window for ToR switches upgrade by using the teamPolicyUpDelay capability, in some cases this capability might not work. As a result, you might see a temporary network outage as the physical network is not ready.

This issue is resolved in this release.

PR 3500545: You cannot apply a vSAN license on 4Kn Software Emulated (SWE) disks due to insufficient capacity

When a vSAN cluster claims a 4Kn SWE disk, the capacity license takes into account the logical block size instead of the physical block size. As a result, the capacity value is much higher than the total size of disks claimed and in some cases, you might not be able to apply the vSAN license to such devices.

This issue is resolved in this release. The fix uses the logical block size and logical block count in pair for capacity calculation.

PR 3514647: You do not see an option to select more than 8 GPUs for vGPU device groups

When setting up a vGPU VM, you do not see an option to select more than 8 GPUs for vGPU device groups. For systems such as NVIDIA HGX, there could be more than 8 physical GPUs with corresponding device groups.

This issue is resolved in this release. The fix adds the option to select 16 GPUs.

PR 3516689: After upgrading to ESXi 8.0 Update 3 or later, PCIe devices are reported as non-passthrough capable or SR-IOV cannot be activated

Due to overly strict sanity checks, after upgrading your system to ESXi 8.0 Update 3 or later, some PCIe devices might be reported as non-passthrough capable or SR-IOV cannot be activated.

This issue is resolved in this release.

PR 3421841: ESXi fails to boot on systems with more than 512 logical CPUs in one NUMA node

Prior to ESXi 8.0 Update 8g, ESXi fails to boot on systems with more than 512 logical processors, or hyperthreads, in one NUMA node. Such systems are very unusual, but configuring the NUMA Nodes per Socket BIOS setting to 0 (NPS0) and deactivating the CCX as NUMA Domain on AMD system can cause all logical processors to be in a single NUMA node.

This issue is resolved in this release.

PR 3506956: You do not see any vGPU device groups for NVIDIA HGX H100 systems

The vGPU device group filtering code does not ensure that each NIC which is a candidate for a NIC-to-GPU device group has the required SR-IOV virtual function active. For NVIDIA HGX H100 systems this issue can cause all vGPU device groups to be ignored.

This issue is resolved in this release.

PR 3500400: Booting virtual machines with vGPU on an NVIDIA Blackwell GPU fails

vGPU VMs that reference a physical NVIDIA Blackwell GPU cannot power on because the vmx process fails.

This issue is resolved in this release.

PR 3517989: ESXi performance degrades with Hitachi NVMe-oF storage target

When using the NVMe/TCP protocol, creation and deployment of virtual machines on datastores backed by a NVMe namespace on Hitachi NVMe-oF target might be slow. when using the default settings.

This issue is resolved in this release.

PR 3480453: Virtual machine vNIC fails to activate after a vSphere vMotion operation

When system memory is highly fragmented, a VM vNIC might fail to allocate continuous physical memory and the vNIC cannot activate after a vSphere vMotion operation. As a result, VMs are randomly going into unresponsive state after successfully migrating to another ESXi host.

This issue is resolved in this release. The fix avoids allocation of continuous physical memory during vNIC activation.

PR 3505426: You might see multiple hardware sensor alerts due to a Baseboard Management Controller (BMC) issue

An underlying BMC issue might cause ESXi hardware health monitoring to produce multiple alerts for hardware sensors on ESXi hosts, which you see in the **vSphere Client under Monitor > Tasks and Events**.

This issue is resolved in this release. The fix logs raw sensor data directly from IPMI to provide a data point to debug possible BMC issues that cause false alerts, but cannot prevent BMC issues to occur.

PR 3026630: You see error messages when trying to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

This issue is resolved in this release.

PR 2916365: VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

This issue is resolved in this release.

PR 3274848: Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

This issue is resolved in this release.

PR 3033184: If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade.`

This issue is resolved in this release.

PR 3288124: You cannot update to ESXi 8.0 Update 2b or later by using `esxcli software vib` commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

This issue is resolved in this release.

PR 3273064: If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

This issue is resolved in this release.

PR 3281300: When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

This issue is resolved in this release.

PR 3151016: Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a vSphere Lifecycle Manager image in the vSphere Client and then navigate away.
2. You trigger a page refresh.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session.

If you use API, firmware compliance details are missing from the response.

This issue is resolved in this release.

PR 3271023: TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as 0xFFFF might cause the end system to drop or delay the TCP packet.

This issue is resolved in this release.

PR 3040952: In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

This issue is resolved in this release.

ESXi-8.0U3g-24859861-no-tools

Profile Name	ESXi-8.0U3g-24859861-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 29, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.79.24859861 - VMware_bootbank_vsan_8.0.3-0.79.24859861 - VMware_bootbank_infravisor_8.0.3-0.79.24859861 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.79.24859861 - VMware_bootbank_crx_8.0.3-0.79.24859861 - VMware_bootbank_bmcad_8.0.3-0.79.24859861 - VMware_bootbank_vsanhealth_8.0.3-0.79.24859861 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.79.24859861 - VMware_bootbank_gc_8.0.3-0.79.24859861 - VMware_bootbank_vdfs_8.0.3-0.79.24859861 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.79.24859861 - VMware_bootbank_esx-base_8.0.3-0.79.24859861 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.79.24859861 - VMware_bootbank_esx-xserver_8.0.3-0.79.24859861 - VMware_bootbank_vds-vsip_8.0.3-0.79.24859861

	• VMware_bootbank_cpu-microcode_8.0.3-0.79.24859861 • VMware_bootbank_esxio_8.0.3-0.79.24859861 • VMware_bootbank_clusterstore_8.0.3-0.79.24859861 • VMware_bootbank_esxio-base_8.0.3-0.79.24859861 • VMware_bootbank_trx_8.0.3-0.79.24859861 • VMware_bootbank_gc-esxio_8.0.3-0.79.24859861 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.79.24859861 • VMware_bootbank_bmcad-esxio_8.0.3-0.79.24859861 • VMware_bootbank_vcls-pod-crx_8.0.3-0.79.24859861 • VMware_bootbank_native-misc-drivers_8.0.3-0.79.24859861 • VMware_bootbank_loadesx_8.0.3-0.79.24859861 • VMware_bootbank_esx-update_8.0.3-0.79.24859861 • VMware_bootbank_loadesxio_8.0.3-0.79.24859861 • VMware_bootbank_esxio-update_8.0.3-0.79.24859861 • VMW_bootbank_iavmd_3.0.0.1010-13vmw.803.0.79.24859861 • VMW_bootbank_nvmetcp_1.0.1.36-1vmw.803.0.79.24859861 • VMW_bootbank_nvmetcp-esxio_1.0.1.36-1vmw.803.0.79.24859861 • VMW_bootbank_vmksdhci_1.0.3-7vmw.803.0.79.24859861 • VMW_bootbank_vmksdhci-esxio_1.0.3-7vmw.803.0.79.24859861
PRs Fixed	3511297, 3523908, 3523307, 3453391, 3502007, 3523288, 3496295, 3505543, 3523292, 3521243, 3549866, 3523289, 3526619, 3505417, 3547311, 3508977, 3514821, 3511112, 3525846, 3522258, 3522712, 3485301, 3522487, 3525785, 3527704, 3535747, 3524945, 3526457, 3499300, 3519342, 3485371, 3493836, 3469074, 3505628, 3500545, 3514647, 3516689, 3421841, 3506956, 3500400, 3517989, 3480453, 3505426, 3026630, 2916365, 3274848, 3033184, 3288124, 3273064, 3281300, 3151016, 3271023, 3040952, 3499150, 3518255, 3517150, 3518979
Related CVE numbers	N/A

This patch updates the following issues:

PR 3511297: You see a no lock error on a vSAN NFS v3 volume

Internal network timeout issues might cause a no lock error on a vSAN NFS v3 volume. In the NFS client debug kernel logs, you see messages such as:

```
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334941] NFS: flock(/file1, t=1, fl=82).
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334945] lockd: nsm_monitor(accloud-fs11.kh.asegroup.com)
Jan 13 14:27:41 khocpnfs01 kernel: [170416.334947] lockd: call procedure 2 on accloud-fs11.kh.asegroup.com.
Jan 13 14:27:44 khocpnfs01 kernel: [170419.335958] lockd: server returns status 2
Jan 13 14:27:44 khocpnfs01 kernel: [170419.335964] lockd: clnt proc returns -37.
```

This issue is resolved in this release.

PR 3523908: Virtual machines cannot boot when configured with a PCI passthrough device

VMs configured with a PCI passthrough device might consistently fail to boot, removing the device from its configuration and from the PCI device list of the ESXi host.

This issue is resolved in this release.

PR 3523307: A timing issue might cause vSAN ESA clusters to show as non-compliant with the SPBM policy

When you deploy a vSAN ESA cluster, a timing issue might cause the vSAN in-memory state for one of its modules to fail to initialize with the appropriate vSAN ESA configuration. As a result, the cluster is incorrectly identified as vSAN OSA and its objects show as non-compliant with the SPBM policy.

This issue is resolved in this release.

PR 3453391: You cannot create new virtual machines in vSAN clusters due to a networking issue

Due to a network issue between the storage cluster and a vSAN witness host, the Cluster Level Object Manager (CLOM) might create new witness components but the DOM does not delete the existing, which leads to exceeding the limit of witness component objects. As a result, creating new virtual machines on such vSAN clusters fails.

This issue is resolved in this release.

PR 3453391: You cannot create new virtual machines in vSAN clusters due to a networking issue

Due to a network issue between the storage cluster and a vSAN witness host, the Cluster Level Object Manager (CLOM) might create new witness components but the DOM does not delete the existing, which leads to exceeding the limit of witness component objects. As a result, creating new virtual machines on such vSAN clusters fails.

This issue is resolved in this release.

PR 3521243: You cannot use public/private key pair to connect with SSH by using HTTP PUT

You cannot use authorized public/private keys to connect to ESXi hosts with SSH by using ConfigManager APIs (HTTP PUT).

This issue is resolved in this release.

PR 3518979: ESX hosts on dual-DPU systems might fail with a purple diagnostic screen due to an interrupt storm caused by the MultiMediaCard (MMC) controller

Interrupt storm caused by the MMC controller on a dual-DPU systems completely occupies one or more CPUs for long periods, which might lead to failure of ESX hosts with a purple diagnostic screen. In the kernel log, you see errors such as:

```
2025-04-30T01:26:26.364Z cpu2:528386) 0x4538e841b530:[0x4200374c98a0]VMKExcEntry@vmkernel#nover+0x6c stack: 0x1
2025-04-30T01:26:26.364Z cpu2:528386) 0x4538e841b530:[0x4200374c98a0]VMKExcEntry@vmkernel#nover+0x6c stack: 0x1
2025-04-30T01:26:26.371Z cpu2:528386) 0x4538e841b660:[0x420036f50d40]Heap_Free@vmkernel#nover+0x26c stack: 0x0
2025-04-30T01:26:26.377Z cpu2:528386) 0x4538e841b6b0:[0x420037dc1598]sleep_list_remove@com.vmware.vmkbs-
d#1+0x104 stack: 0x0
2025-04-30T01:26:26.385Z cpu2:528386) 0x4538e841b700:[0x420037f66970]dwmmc_intr@ (vmksdhci) #<None>+0x200 stack:
0x0
2025-04-30T01:26:26.392Z cpu2:528386) 0x4538e841b730:[0x420036f78454]IntrCookieBH@vmkernel#nover+0x208 stack:
0x0
2025-04-30T01:26:26.398Z cpu2:528386) 0x4538e841b7e0:[0x420036f556e0]BH_DrainAndDisableInterrupts@vmker-
nel#nover+0x21c stack: 0x0
```

This issue rarely occurs in single-DPU systems.

This issue is resolved in this release.

PR 3517150: When data digests on NVMe over TCP datastores are active, you might see frequent connection drops and performance degradation

Various applications can modify the write I/O data that is in transit between ESXi hosts before the nvmetcp driver completes a task. When data digests on NVMe over TCP datastores are active, if data is modified after the digest is calculated within the nvmetcp driver, the target receives the modified I/O data. This leads to a digest mismatch, which in turn causes the controller connection to drop, the controller to reset, and ultimately, I/O performance degradation.

This issue is resolved in this release.

PR 3518255: During booting or when the Intel Volume Management Device (VMD) driver first loads, NVMe devices might fail in the enumeration of namespaces

If an NVMe drive has 512 namespaces, the VMD driver runs out of memory when trying to enumerate them all.

This issue is resolved in this release. ESXi 8.0 Update 3g contains VMD driver version 3.0.0.1010-13 that provides the fix. If you do not update your system to ESXi 8.0 Update 3g, you can just update the VMD driver to version 3.0.0.1010-13.

PR 3499150: NSX deployment fails during nsx-datapath VIB installation due to an out-of-memory condition

Insufficient sizing of the memory resource pool of one of the services during the installation of VIBs in an NSX deployment might cause the deployment to fail. In the `vmkernel.log` you see an error such as:

```
2025-02-11T19:39:18.172Z In(182) vmkernel: cpu68:2107749)Admission failure in path: host/
vim/vmvisor/settingsd-task-forks/python.2107554:python.2107764:uw.2107764
```

This issue is resolved in this release.

PR 3502007: Activation of TLS 1.3 in ClusterStore does not persist across ESXi host reboots

The `agent_settings` configuration on ESXi that controls whether TLS 1.3 is active on ClusterStore, which keeps the cluster-specific configuration data, reverts to the default of not active after a reboot of the ESX host.

This issue is resolved in this release.

PR 3523288: A virtual machine might fail to power on if the .vmx file is changed in the backend

In certain scenarios, when a VM is restored from a backup, applications external to the ESXi host might delete and recreate the .vmx file while the ESXi nfs41client is still accessing the file handle of the old .vmx file. As a result, a stale handle error causes the power-on of the VM to fail.

This issue is resolved in this release.

PR 3496295: When you use the suspend and resume feature on a virtual machine with Kubernetes control plane container configured, the VM might take long to respond or reset

During suspend and resume of a VM, the device state of a Kubernetes Control Plane container might change after a checkpoint and cause an inconsistent state of the device. As a result, the guest driver might reset the device.

This issue is resolved in this release.

PR 3505543: If you use ESXi host keys to encrypt virtual machines, the hostd service might fail

When you use an ESXi host key to encrypt a VM, if you delete the VM or change the key, the removal of the original host key might trigger an exception that can cause the hostd service to fail. As a result, the ESXi host temporarily disconnects from vCenter and you might not be able to manage VMs on the host if hostd does not restart after the failure.

This issue is resolved in this release.

PR 3523292: An ESXi host might fail with a purple diagnostic screen during migration of virtual machines from a Tintri NFSv4.1 datastore

During the migration of VMs by using vSphere vMotion from a Tintri NFSv4.1 datastore, the ESXi host might fail with a purple diagnostic screen. This issue is specific to how Tintri NFSv4.1 servers handle NULL values in certain scenarios.

This issue is resolved in this release.

PR 3549866: vSphere Virtual Volumes daemon fails after converting a VASA provider from external to internal

Due to an issue with handling IDs during an unbind operation for some vSphere API for Storage Awareness (VASA) providers, in some tasks such as converting a VASA provider from external to internal, the vSphere Virtual Volumes daemon `vvold` might fail as it cannot map unique IDs.

This issue is resolved in this release. With the fix, vvolld gracefully handles such conditions and stops the unbind operation with message in the logs instead of failing.

PR 3523289: ClusterStore services might crashloop due to excessive memory usage and lead to network overload

A high default setting of a parameter in the ClusterStore capability might cause the service to consume more memory than necessary and reach the memory limit over time. As a result, ClusterStore services might crashloop and cause network overload as it attempts to communicate from one ESXi host to another during each startup.

This issue is resolved in this release.

PR 3526619: The authentication service of the Distributed Key Value Store, or ClusterStore, might cause DNS traffic overload

An issue with the authentication service of the Distributed Key Value Store, also called ClusterStore, that might keep trying to connect to a group of ESXi hosts even when that ClusterStore is no longer part of that group might cause excessive network utilization and eventually overload of certain infrastructure such as DNS.

This issue is resolved in this release. The fix implements a rate-limiter to reduce the retry frequency and network overhead.

PR 3505417: IF-MIB SNMP data intermittently reports hidden ports DVFilter Coalesce Portset and pps

When using third-party tools, the IF-MIB SNMP data might intermittently report hidden portsets, such as pps and DVFilter Coalesce Portset, even when they are not part of the ESXi host configuration.

This issue is resolved in this release.

PR 3547311: After a syslog configuration change, virtual machines lose connectivity to vSAN iSCSI targets

Configuring a remote syslog loghost might remove all non-syslog dynamic firewall rules, leading to service disruption such as VMs losing connection to vSAN iSCSI targets.

This issue is resolved in this release.

PR 3508977: During a snapshot operation of a VM with IOFilter attached, the VMX might fail due to an unresponsive thread

In very rare situations, once in 2000 snapshot operations, a race condition might occur in the I/O filter infrastructure cause the I/O filter watchdog to stop the VMX service and generate a VMX core dump.

This issue is resolved in this release.

PR 3514821: After a reboot of ESXi hosts, you see alarms in the vSphere Client that some uplinks are down

Due to some limitations in the VMkernel Observation (VOB) and hostd services, uplink up and down events might not be sent to vCenter in the correct order. As a result, when a task such as rebooting ESXi hosts completes, in the vSphere Client you see false alerts that some uplinks are down.

This issue is resolved in this release.

PR 3511112: After running the Linux fstrim command on a virtual machine, the ESXi host fails with a purple diagnostics screen

A VMFS heap memory exhaustion might cause memory access violations when clearing the pointer block cache. As a result, when you run the Linux `fstrim` command on a virtual machine, the ESXi host might fail with a purple diagnostics screen with a `PB3Manager` error in the backtrace.

This issue is resolved in this release.

For vSphere Quick Boot support in ESXi 8.0 Update 3g, see the list for each vendor in the Additional Information section of the [Understanding ESXi Quick Boot Compatibility](#) KB.

PR 3525846: Virtual machines with fixed passthrough hot-plugging enabled might fail due to misallocation of interrupts

For VMs with fixed passthrough hot-plugging enabled, after subsequent hot-adds of passthrough NVMe drives, the functionality might fail because interrupts might not be properly allocated.

This issue is resolved in this release.

PR 3522258: Mounting shares might fail due to rsyslogd CPU usage at 100%

During a restart of the rsyslogd service, CPU usage might reach 100% and you temporarily might not be able to mount shares.

This issue is resolved in this release.

PR 3522712: ESXi host reboot stops due to indefinite delay in the shutdown of the vSAN cluster monitoring, membership, and directory service (CMMDS)

When rebooting an ESXi host, network conditions can cause the vSAN CMMDS shutdown process to delay indefinitely.

This issue is resolved in this release.

PR 3485301: You see a delay in the power-on of vGPU virtual machines

You might see a delay in the power-on of virtual machines with NVIDIA vGPU devices and if many such tasks accumulate in vCenter, the vpxd service might be overloaded and vCenter becomes intermittently unresponsive. A restart of the vpxd service restores vCenter responsiveness.

This issue is resolved in this release.

PR 3522487: A routing issue with the Cisco Application Centric Infrastructure (ACI) switch might cause loss of vSAN traffic

To override normal IP routing, vSAN binds all outgoing traffic on TCP sockets and traffic runs over a vmkernel adapter dedicated to vSAN. In some cases, the vSAN binding might not be set or a TCP socket might be missing, so packets are sent with the normal IP routing. In such cases, Cisco ACI devices might change the routing table based on the incoming packets, which in turn causes traffic loss on the vSAN cluster.

This issue is resolved in this release.

PR 3525785: Excessive logs in the nsxdavim.log might cause an overload of the syslog server

Frequent updates from hostd on an NSX distributed virtual port group, such as add, update, and delete, might cause excessive logs in the `nsxdavim.log`, which in some cases might lead to an overload of the syslog server.

This issue is resolved in this release.

PR 3527704: ESXi host server firmware might not automatically make the PCIe Atomic Operations feature available to passthrough devices

GPU workloads on a virtual machine with AMD GPU VMDirectPath I/O might fail or hang because the ESXi host server firmware might not automatically activate PCIe atomic operations for AMD GPU devices.

This issue is resolved in this release. The fix adds the VMX configuration option `pciPassthru.exposeAtomicOps`, which you can set to `TRUE` to make sure VMs can use AMD GPUs as a passthrough device.

PR 3535747: On debug builds, an ESXi host might fail with a purple diagnostic screen due to a race condition

On debug builds, a race condition between the power-off task on a virtual machine and ESXi memory scrubbing might cause the ESXi host to fail with a purple diagnostic screen and the following error:

```
ASSERT bin/bora/vmkernel/private/_virtual_includes/hdrs/mpage2M.h:264
```

The issue does not affect release builds.

This issue is resolved in this release.

PR 3524945: Adding new disks external to the home directory of a virtual machine might fail with an error for reaching the security policy limit

When a disk external to the home directory of a VM is added or removed, the security domain policy of the VM updates to give or revoke access to the external directory. It is possible that if an error occurs during disk load, to avoid issues if it is resolved, external directories are no longer removed from the security policy, even after a resolution. In time, the rules in the security policy might reach the max limit. As a result, adding new disks external to the home directory of a virtual machine fails and in the `sandboxd.log` you see errors such as:

```
2025-05-11T08:14:55.306Z Er(163) sandboxd[2099161]: [Originator@6876
sub=SecurityDomainsImpl] AddUpdateFileRules: Failed to add file policies to domain :X:
world ID :Y:Argument list too long
```

This issue is resolved in this release. The fix makes sure that once a disk load issue is resolved, external directories are removed from the security policy. If a VM is already facing the issue, either restart the VM or migrate it to another ESXi host to trigger a regeneration of the security policy.

PR 3526457: You see multiple resets of a virtual machine in case of a machine check exception (MCE) error

If an uncorrectable memory error occurs on the guest OS of a virtual machine, memory access on the error page either by the guest or the ESXi hypervisor triggers a machine check exception. In such cases, the hypervisor might try to inject a virtual machine check error to the guest, which leads to guest OS failure. Depending on the guest OS, you might see multiple resets of the VM or it might become unresponsive.

This issue is resolved in this release.

PR 3499300: Virtual machine migration fails during upgrade of NSX or Distributed Firewall (DFW) reconfiguration

If the destination host of a VM migration operation is in the process of an upgrade of NSX or DFW reconfiguration, VM migrations to that host might fail.

This issue is resolved in this release.

PR 3519342: ESXi host fails with purple diagnostic screen during unmap operations on NVMe 2.0 devices due to misinterpreted unmap limits

Incorrect handling of the `maxUnmapDescriptorCount` and `maxUnmapLbaCount` values, which might be set as 0 by an NVMe target, can cause ESXi hosts to fail with a purple diagnostic screen when performing UNMAP operations on NVMe 2.0 devices. The NVMe 2.0 specification defines 0 in the `maxUnmapDescriptorCount` and `maxUnmapLbaCount` fields as indicating no limits on UNMAP operations, but ESXi reads them as indicating lack of UNMAP support. As a result, device-level UNMAP operation exits prematurely, leading to the failure.

This issue is resolved in this release.

PR 3485371: The VMX service might unexpectedly end replication-based snapshot operation due to a storage system failure

A storage system failure during a replication-based snapshot operation can lead to a logic error in the data protection component, potentially resulting in invalid memory access. As a result, the VMX service might fail and end the task.

This issue is resolved in this release.

PR 3493836: Virtual machine fails due to Signal 11, segmentation violation, during power on

In rare cases, a virtual machine might fail with Signal 11 error due to an issue with initializing VM statistics during power on.

This issue is resolved in this release.

PR 3469074: NVME over TCP adapter occasionally fails to reconnect to some storage array ports during storage array upgrade and controller failover

Controllers on an ESXi host might consistently fail to come online and require a system reboot to restore their operational state due to intermittent failures of the NVME over TCP adapter to reconnect to some storage array ports during storage array upgrade and controller failover. As a result, you see VMs on such hosts becoming unresponsive.

This issue is resolved in this release.

PR 3505628: During upgrade of ToR switches, the teamPolicyUpDelay capability of the "Route based on physical NIC load" teaming policy might not work consistently

If you set up the "Route based on physical NIC load" teaming policy to provide a window for ToR switches upgrade by using the teamPolicyUpDelay capability, in some cases this capability might not work. As a result, you might see a temporary network outage as the physical network is not ready.

This issue is resolved in this release.

PR 3500545: You cannot apply a vSAN license on 4Kn Software Emulated (SWE) disks due to insufficient capacity

When a vSAN cluster claims a 4Kn SWE disk, the capacity license takes into account the logical block size instead of the physical block size. As a result, the capacity value is much higher than the total size of disks claimed and in some cases, you might not be able to apply the vSAN license to such devices.

This issue is resolved in this release. The fix uses the logical block size and logical block count in pair for capacity calculation.

PR 3514647: You do not see an option to select more than 8 GPUs for vGPU device groups

When setting up a vGPU VM, you do not see an option to select more than 8 GPUs for vGPU device groups. For systems such as NVIDIA HGX, there could be more than 8 physical GPUs with corresponding device groups.

This issue is resolved in this release. The fix adds the option to select 16 GPUs.

PR 3516689: After upgrading to ESXi 8.0 Update 3 or later, PCIe devices are reported as non-passthrough capable or SR-IOV cannot be activated

Due to overly strict sanity checks, after upgrading your system to ESXi 8.0 Update 3 or later, some PCIe devices might be reported as non-passthrough capable or SR-IOV cannot be activated.

This issue is resolved in this release.

PR 3421841: ESXi fails to boot on systems with more than 512 logical CPUs in one NUMA node

Prior to ESXi 8.0 Update 8g, ESXi fails to boot on systems with more than 512 logical processors, or hyperthreads, in one NUMA node. Such systems are very unusual, but configuring the NUMA Nodes per Socket BIOS setting to 0 (NPS0) and deactivating the CCX as NUMA Domain on AMD system can cause all logical processors to be in a single NUMA node.

This issue is resolved in this release.

PR 3506956: You do not see any vGPU device groups for NVIDIA HGX H100 systems

The vGPU device group filtering code does not ensure that each NIC which is a candidate for a NIC-to-GPU device group has the required SR-IOV virtual function active. For NVIDIA HGX H100 systems this issue can cause all vGPU device groups to be ignored.

This issue is resolved in this release.

PR 3500400: Booting virtual machines with vGPU on an NVIDIA Blackwell GPU fails

vGPU VMs that reference a physical NVIDIA Blackwell GPU cannot power on because the vmx process fails.

This issue is resolved in this release.

PR 3517989: ESXi performance degrades with Hitachi NVMe-oF storage target

When using the NVMe/TCP protocol, creation and deployment of virtual machines on datastores backed by a NVMe namespace on Hitachi NVMe-oF target might be slow. when using the default settings.

This issue is resolved in this release.

PR 3480453: Virtual machine vNIC fails to activate after a vSphere vMotion operation

When system memory is highly fragmented, a VM vNIC might fail to allocate continuous physical memory and the vNIC cannot activate after a vSphere vMotion operation. As a result, VMs are randomly going into unresponsive state after successfully migrating to another ESXi host.

This issue is resolved in this release. The fix avoids allocation of continuous physical memory during vNIC activation.

PR 3505426: You might see multiple hardware sensor alerts due to a Baseboard Management Controller (BMC) issue

An underlying BMC issue might cause ESXi hardware health monitoring to produce multiple alerts for hardware sensors on ESXi hosts, which you see in the **vSphere Client under Monitor > Tasks and Events**.

This issue is resolved in this release. The fix logs raw sensor data directly from IPMI to provide a data point to debug possible BMC issues that cause false alerts, but cannot prevent BMC issues to occur.

PR 3026630: You see error messages when trying to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

This issue is resolved in this release.

PR 2916365: VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

This issue is resolved in this release.

PR 3274848: Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

This issue is resolved in this release.

PR 3033184: If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade.`

This issue is resolved in this release.

PR 3288124: You cannot update to ESXi 8.0 Update 2b or later by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

This issue is resolved in this release.

PR 3273064: If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

This issue is resolved in this release.

PR 3281300: When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

This issue is resolved in this release.

PR 3151016: Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a vSphere Lifecycle Manager image in the vSphere Client and then navigate away.
2. You trigger a page refresh.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session.

If you use API, firmware compliance details are missing from the response.

This issue is resolved in this release.

PR 3271023: TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

This issue is resolved in this release.

PR 3040952: In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

This issue is resolved in this release.

ESXi-8.0U3sg-24853260-standard

Profile Name	ESXi-8.0U3sg-24853260-standard
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 29, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_drivervm-gpu-base_8.0.3-0.76.24853260 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.76.24853260 - VMware_bootbank_vsan_8.0.3-0.76.24853260 - VMware_bootbank_esxio-combiner_8.0.3-0.76.24853260 - VMware_bootbank_bmcad-esxio_8.0.3-0.76.24853260 - VMware_bootbank_trx_8.0.3-0.76.24853260 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.76.24853260 - VMware_bootbank_crx_8.0.3-0.76.24853260 - VMware_bootbank_cpu-microcode_8.0.3-0.76.24853260 - VMware_bootbank_clusterstore_8.0.3-0.76.24853260 - VMware_bootbank_esxio_8.0.3-0.76.24853260 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.76.24853260 - VMware_bootbank_vcls-pod-crx_8.0.3-0.76.24853260 - VMware_bootbank_gc_8.0.3-0.76.24853260 - VMware_bootbank_vds-vsip_8.0.3-0.76.24853260 - VMware_bootbank_gc-esxio_8.0.3-0.76.24853260 - VMware_bootbank_infravisor_8.0.3-0.76.24853260 - VMware_bootbank_vsanhealth_8.0.3-0.76.24853260 - VMware_bootbank_esx-xserver_8.0.3-0.76.24853260 - VMware_bootbank_native-misc-drivers_8.0.3-0.76.24853260 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.76.24853260 - VMware_bootbank_bmcad_8.0.3-0.76.24853260 - VMware_bootbank_vdfs_8.0.3-0.76.24853260 - VMware_bootbank_esxio-base_8.0.3-0.76.24853260 - VMware_bootbank_esx-base_8.0.3-0.76.24853260

	- VMware_bootbank_loadesx_8.0.3-0.76.24853260 - VMware_bootbank_esx-update_8.0.3-0.76.24853260 - VMware_bootbank_loadesxio_8.0.3-0.76.24853260 - VMware_bootbank_esxio-update_8.0.3-0.76.24853260 - VMware_locker_tools-light_12.5.3.24819442-24853260
PRs Fixed	3492319, 3506965, 3518336
Related CVE numbers	N/A

This patch updates the following issues:

VMware Tools Bundling Changes in ESXi 8.0 Update 3g

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 3g:

- windows.iso:** VMware Tools 12.5.3 supports 64-bit Windows versions - Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- windows-x86.iso:** VMware Tools 12.4.7 supports 32-bit Windows versions.
- linux.iso:** VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - windows.iso:** for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - winPreVista.iso:** for Windows 2000, Windows XP, and Windows 2003.
 - linuxPreGLibc25.iso:** supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- solaris.iso:** VMware Tools image 10.3.10 for Solaris.
- darwin.iso:** Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.5.3 Release Notes](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Upgrading VMware Tools](#)

ESXi 8.0 Update 3g includes the following AMD microcode:

ESXi builds provide microcode for most supported CPUs, and if the BIOS of your server has an older microcode version than the one provided with the ESXi build, ESXi attempts to load the newer version. Because of an AMD change in [microcode signing](#), ESXi might not be able to apply the latest microcode. In such a case, ESXi 8.0 Update 3g applies the newest microcode that it can and if the version is not the latest, ESXi issues a SysAlert such as `This machine's BIOS should be updated to pick up the latest microcode`. Updating to the latest BIOS is always best practice as usually the latest BIOS also carries the latest microcode.

Code Name	FMS	MCU Rev	MCU Date	Brand Names
Bulldozer	0x00600f12 (15/01/2)	0x0600063e	2/7/2018	Opteron 6200/4200/3200 Series
Piledriver	0x00600f20 (15/02/0)	0x06000852	2/6/2018	Opteron 6300/4300/3300 Series
Zen-Naples	0x00800f12 (17/01/2)	0x08001278	11/11/2024	EPYC 7001 Series
Zen2-Rome	0x00830f10 (17/31/0)	0x0830107c	12/18/2023	EPYC 7002/7Fx2/7Hx2 Series

Code Name	FMS	MCU Rev	MCU Date	Brand Names
Zen3-Milan	0x00a00f11 (19/01/1)	0x0a0011db	11/11/2024	EPYC 7003 Series
Zen3-Milan (old BIOS)	0x00a00f11 (19/01/1)	0x0a0011d5	2/23/2024	EPYC 7003 Series
Zen3-Milan-X	0x00a00f12 (19/01/2)	0x0a001244	11/11/2024	EPYC 7003X Series
Zen3-Milan-X (old BIOS)	0x00a00f12 (19/01/2)	0x0a001238	2/26/2024	EPYC 7003X Series
Zen4-Genoa	0x00a10f11 (19/11/1)	0x0a101154	11/12/2024	EPYC 9004 Series
Zen4-Genoa (old BIOS)	0x00a10f11 (19/11/1)	0x0a101148	2/23/2024	EPYC 9004 Series
Zen4-Genoa-X	0x00a10f12 (19/11/2)	0x0a10124f	11/12/2024	EPYC 9004X Series
Zen4-Genoa-X (old BIOS)	0x00a10f12 (19/11/2)	0x0a101248	2/26/2024	EPYC 9004X Series
Zen4-Bergamo/Siena	0x00aa0f02 (19/a0/2)	0x0aa00219	11/13/2024	EPYC 8004/9004 Series
Zen4-Bergamo/Siena (old BIOS)	0x00aa0f02 (19/a0/2)	0x0aa00215	2/28/2024	EPYC 8004/9004 Series

ESXi 8.0 Update 3g includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003901	12/12/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003901	12/12/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/ b)	0xbf	baseline	0x7002b01	12/12/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd000404	1/7/2025	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002d0	1/7/2025	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xbc	12/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x3c	12/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x56	12/11/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	baseline	0x2c0003e0	7/30/2024	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b000639	1/28/2025	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xfa	7/28/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x104	11/14/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x64	12/1/2024	Intel Xeon E-2300 Series
Granite Rapids SP R1S	0xa06d1 (06/ad/1)	0x20	baseline	0xa0000d1	2/7/2025	Intel Xeon 6500P/6700P/6900P Series
Granite Rapids AP/SP	0xa06d1 (06/ad/1)	0x95	baseline	0x10003a2	2/7/2025	Intel Xeon 6500P/6700P/6900P Series
Sierra Forest	0xa06f3 (06/af/3)	0x01	baseline	0x3000341	2/10/2025	Intel Xeon 6700E/6900E Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x12e	1/15/2025	Intel Xeon E-2400 Series
Raptor Lake E/HX/S	0xb0674 (06/b7/4)	0x32	baseline	0x12e	1/15/2025	Intel Xeon E-2400 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x210002a9	3/14/2025	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series

ESXi-8.0U3sg-24853260-no-tools

Profile Name	ESXi-8.0U3sg-24853260-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 29, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_drivervm-gpu-base_8.0.3-0.76.24853260 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.76.24853260 - VMware_bootbank_vsan_8.0.3-0.76.24853260 - VMware_bootbank_esxio-combiner_8.0.3-0.76.24853260 - VMware_bootbank_bmc-al-esxio_8.0.3-0.76.24853260 - VMware_bootbank_trx_8.0.3-0.76.24853260 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.76.24853260 - VMware_bootbank_crx_8.0.3-0.76.24853260 - VMware_bootbank_cpu-microcode_8.0.3-0.76.24853260 - VMware_bootbank_clusterstore_8.0.3-0.76.24853260 - VMware_bootbank_esxio_8.0.3-0.76.24853260 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.76.24853260 - VMware_bootbank_vcls-pod-crx_8.0.3-0.76.24853260 - VMware_bootbank_gc_8.0.3-0.76.24853260 - VMware_bootbank_vds-vsip_8.0.3-0.76.24853260 - VMware_bootbank_gc-esxio_8.0.3-0.76.24853260 - VMware_bootbank_infravisor_8.0.3-0.76.24853260 - VMware_bootbank_vsanhealth_8.0.3-0.76.24853260 - VMware_bootbank_esx-xserver_8.0.3-0.76.24853260 - VMware_bootbank_native-misc-drivers_8.0.3-0.76.24853260 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.76.24853260 - VMware_bootbank_bmc-al_8.0.3-0.76.24853260 - VMware_bootbank_vdfs_8.0.3-0.76.24853260 - VMware_bootbank_esxio-base_8.0.3-0.76.24853260 - VMware_bootbank_esx-base_8.0.3-0.76.24853260 - VMware_bootbank_loadesx_8.0.3-0.76.24853260 - VMware_bootbank_esx-update_8.0.3-0.76.24853260

	- VMware_bootbank_loadesxio_8.0.3-0.76.24853260 - VMware_bootbank_esxio-update_8.0.3-0.76.24853260
PRs Fixed	3492319, 3506965, 3518336
Related CVE numbers	N/A

This patch updates the following issues:

VMware Tools Bundling Changes in ESXi 8.0 Update 3g

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 3g:

- windows.iso**: VMware Tools 12.5.3 supports 64-bit Windows versions - Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- windows-x86.iso**: VMware Tools 12.4.7 supports 32-bit Windows versions.
- linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.5.3 Release Notes](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Upgrading VMware Tools](#)

ESXi 8.0 Update 3g includes the following AMD microcode:

ESXi builds provide microcode for most supported CPUs, and if the BIOS of your server has an older microcode version than the one provided with the ESXi build, ESXi attempts to load the newer version. Because of an AMD change in [microcode signing](#), ESXi might not be able to apply the latest microcode. In such a case, ESXi 8.0 Update 3g applies the newest microcode that it can and if the version is not the latest, ESXi issues a SysAlert such as `This machine's BIOS should be updated to pick up the latest microcode`. Updating to the latest BIOS is always best practice as usually the latest BIOS also carries the latest microcode.

Code Name	FMS	MCU Rev	MCU Date	Brand Names
Bulldozer	0x00600f12 (15/01/2)	0x0600063e	2/7/2018	Opteron 6200/4200/3200 Series
Piledriver	0x00600f20 (15/02/0)	0x06000852	2/6/2018	Opteron 6300/4300/3300 Series
Zen-Naples	0x00800f12 (17/01/2)	0x08001278	11/11/2024	EPYC 7001 Series
Zen2-Rome	0x00830f10 (17/31/0)	0x0830107c	12/18/2023	EPYC 7002/7Fx2/7Hx2 Series
Zen3-Milan	0x00a00f11 (19/01/1)	0x0a0011db	11/11/2024	EPYC 7003 Series
Zen3-Milan (old BIOS)	0x00a00f11 (19/01/1)	0x0a0011d5	2/23/2024	EPYC 7003 Series

Code Name	FMS	MCU Rev	MCU Date	Brand Names
Zen3-Milan-X	0x00a00f12 (19/01/2)	0x0a001244	11/11/2024	EPYC 7003X Series
Zen3-Milan-X (old BIOS)	0x00a00f12 (19/01/2)	0x0a001238	2/26/2024	EPYC 7003X Series
Zen4-Genoa	0x00a10f11 (19/11/1)	0x0a101154	11/12/2024	EPYC 9004 Series
Zen4-Genoa (old BIOS)	0x00a10f11 (19/11/1)	0x0a101148	2/23/2024	EPYC 9004 Series
Zen4-Genoa-X	0x00a10f12 (19/11/2)	0x0a10124f	11/12/2024	EPYC 9004X Series
Zen4-Genoa-X (old BIOS)	0x00a10f12 (19/11/2)	0x0a101248	2/26/2024	EPYC 9004X Series
Zen4-Bergamo/Siena	0x00aa0f02 (19/a0/2)	0x0aa00219	11/13/2024	EPYC 8004/9004 Series
Zen4-Bergamo/Siena (old BIOS)	0x00aa0f02 (19/a0/2)	0x0aa00215	2/28/2024	EPYC 8004/9004 Series

ESXi 8.0 Update 3g includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
						Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
						v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003901	12/12/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003901	12/12/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	baseline	0x7002b01	12/12/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd000404	1/7/2025	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002d0	1/7/2025	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xbc	12/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x3c	12/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x56	12/11/2024	Intel Xeon W-11000E Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	baseline	0x2c0003e0	7/30/2024	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b000639	1/28/2025	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xfa	7/28/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x104	11/14/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x64	12/1/2024	Intel Xeon E-2300 Series
Granite Rapids SP R1S	0xa06d1 (06/ad/1)	0x20	baseline	0xa0000d1	2/7/2025	Intel Xeon 6500P/6700P/6900P Series
Granite Rapids AP/SP	0xa06d1 (06/ad/1)	0x95	baseline	0x10003a2	2/7/2025	Intel Xeon 6500P/6700P/6900P Series
Sierra Forest	0xa06f3 (06/af/3)	0x01	baseline	0x3000341	2/10/2025	Intel Xeon 6700E/6900E Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x12e	1/15/2025	Intel Xeon E-2400 Series
Raptor Lake E/HX/S	0xb0674 (06/b7/4)	0x32	baseline	0x12e	1/15/2025	Intel Xeon E-2400 Series
Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x210002a9	3/14/2025	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series

ESXi 8.0 U3g - 24859861

Name	ESXi
Version	ESXi 8.0 U3g - 24859861
Release Date	July 29, 2025
Category	Bugfix
Affected Components	• ESXi Component - core ESXi VIBs • ESXi Install/Upgrade Component • ESXi Install/Upgrade Component • Intel NVME Driver with VMD Technology • VMware NVMe over TCP Driver • MMC/SD HCI Driver • ESXi Tools Component
PRs Fixed	3511297, 3523908, 3523307, 3453391, 3502007, 3523288, 3496295, 3505543, 3523292, 3521243, 3549866, 3523289, 3526619, 3505417, 3547311, 3508977, 3514821, 3511112, 3525846, 3522258, 3522712, 3485301, 3522487, 3525785, 3527704, 3535747, 3524945, 3526457, 3499300, 3519342, 3485371, 3493836, 3469074, 3505628, 3500545, 3514647, 3516689, 3421841, 3506956, 3500400, 3517989, 3480453, 3505426, 3026630, 2916365, 3274848, 3033184, 3288124, 3273064, 3281300, 3151016, 3271023, 3040952, 3499150, 3518255, 3517150, 3518979
Related CVE numbers	N/A

This patch resolves the issues listed in the ESXi_8.0.3-0.79.24859861 bulletin.

ESXi 8.0 U3sg - 24853260

Name	ESXi
Version	ESXi 8.0 U3sg - 24853260
Release Date	July 29, 2025
Category	Security
Affected Components	• ESXi Component - core ESXi VIBs • ESXi Install/Upgrade Component • ESXi Tools Component
PRs Fixed	3492319, 3506965, 3518336
Related CVE numbers	N/A

This patch resolves the issues listed in the ESXi_8.0.3-0.76.24853260 bulletin.

Known Issues from Previous Releases**Installation, Upgrade, and Migration Issues**

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

The Cancel option in an interactive ESXi installation might not work as expected

Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings > <specific broadcom NIC> > Device Configuration Menu > Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some nmlx5_core driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

The irdman driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic

If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the irdman driver might fail. This issue is unlikely to occur, as the irdman driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

Workaround: Use ULPs with RC transport mode.

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with

uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1

The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out

vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD

In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See [Reconciling Discrepancies in the Managed Virtual Disk Catalog](#).

vSphere Lifecycle Manager Issues

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 3f Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 3f | 15 JUL 2025 | Build 24784735
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

Patches Contained in This Release

VMware ESXi 8.0 Update 3f

Build Details

VMware vSphere Hypervisor (ESXi) Offline Bundle

Download Filename:	VMware-ESXi-8.0U3f-24784735-depot.zip
Build:	24784735
Download Size:	640.9 MB
SHA256 checksum:	2c35d498540de2fd1dc8217b52cf7c71e6a69b8117253b10abe349b7344686be
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.3-0.73.24784735	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.73.24784735	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.73.24784735	Security	Critical

Rollup Bulletins

These rollup bulletins contain the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U3f-24784735	Security	Critical	Security only image

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name

ESXi-8.0U3f-24784735-standard

ESXi-8.0U3f-24784735-no-tools

ESXi Images

Name and Version	Release Date	Category	Detail
ESXi 8.0 U3f - 24784735	15 JUL 2025	General	Security only image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

See these KBs for details on the authenticated online and offline download: [390121](#), [390098](#), and [389276](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually download the patch offline bundle ZIP file or see the KB [How to download ESXi ISO image for all releases including patch updates](#).

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.3-0.73.24784735

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.73.24784735 - VMware_bootbank_vsan_8.0.3-0.73.24784735 - VMware_bootbank_infravisor_8.0.3-0.73.24784735 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.73.24784735 - VMware_bootbank_crx_8.0.3-0.73.24784735 - VMware_bootbank_bmcal_8.0.3-0.73.24784735 - VMware_bootbank_vsanhealth_8.0.3-0.73.24784735 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.73.24784735 - VMware_bootbank_gc_8.0.3-0.73.24784735 - VMware_bootbank_vdfs_8.0.3-0.73.24784735 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.73.24784735 - VMware_bootbank_esx-base_8.0.3-0.73.24784735 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.73.24784735 - VMware_bootbank_esx-xserver_8.0.3-0.73.24784735 - VMware_bootbank_vds-vsip_8.0.3-0.73.24784735

	- VMware_bootbank_cpu-microcode_8.0.3-0.73.24784735 - VMware_bootbank_esxio_8.0.3-0.73.24784735 - VMware_bootbank_clusterstore_8.0.3-0.73.24784735 - VMware_bootbank_esxio-base_8.0.3-0.73.24784735 - VMware_bootbank_trx_8.0.3-0.73.24784735 - VMware_bootbank_gc-esxio_8.0.3-0.73.24784735 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.73.24784735 - VMware_bootbank_bmcad-esxio_8.0.3-0.73.24784735 - VMware_bootbank_vcls-pod-crx_8.0.3-0.73.24784735 - VMware_bootbank_native-misc-drivers_8.0.3-0.73.24784735
PRs Fixed	N/A
CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

This patch updates the esx-base VIB. Due to their dependency with the esx-base VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: esxio-combiner, vsan, infravisor, esxio-dvfilter-generic-fastpath, crx, bmcad, vsanhealth, drivervm-gpu-base, gc, vdfs, native-misc-drivers-esxio, esx-base, esx-dvfilter-generic-fastpath, esx-xserver, vds-vsip, cpu-microcode, esxio, clusterstore, esxio-base, trx, gc-esxio, esxio-combiner-esxio, bmcad-esxio, vcls-pod-crx, and native-misc-drivers .

This patch resolves the following issues:

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

esx-update_8.0.3-0.73.24784735

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesx_8.0.3-0.73.24784735 - VMware_bootbank_esx-update_8.0.3-0.73.24784735
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the esx-base VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: loadesx and esx-update .

esxio-update_8.0.3-0.73.24784735

Patch Category	Security
-----------------------	----------

Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesxio_8.0.3-0.73.24784735 - VMware_bootbank_esxio-update_8.0.3-0.73.24784735
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesxio` and `esxio-update`.

ESXi-8.0U3f-24784735-standard

Profile Name	ESXi-8.0U3f-24784735-standard
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 15, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.73.24784735 - VMware_bootbank_vsan_8.0.3-0.73.24784735 - VMware_bootbank_infravisor_8.0.3-0.73.24784735 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.73.24784735 - VMware_bootbank_crx_8.0.3-0.73.24784735 - VMware_bootbank_bmcad_8.0.3-0.73.24784735 - VMware_bootbank_vsanhealth_8.0.3-0.73.24784735 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.73.24784735 - VMware_bootbank_gc_8.0.3-0.73.24784735 - VMware_bootbank_vdfs_8.0.3-0.73.24784735 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.73.24784735 - VMware_bootbank_esx-base_8.0.3-0.73.24784735 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.73.24784735 - VMware_bootbank_esx-xserver_8.0.3-0.73.24784735 - VMware_bootbank_vds-vsip_8.0.3-0.73.24784735 - VMware_bootbank_cpu-microcode_8.0.3-0.73.24784735 - VMware_bootbank_esxio_8.0.3-0.73.24784735 - VMware_bootbank_clusterstore_8.0.3-0.73.24784735 - VMware_bootbank_esxio-base_8.0.3-0.73.24784735 - VMware_bootbank_trx_8.0.3-0.73.24784735 - VMware_bootbank_gc-esxio_8.0.3-0.73.24784735 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.73.24784735

	- VMware_bootbank_bmcad-esxio_8.0.3-0.73.24784735 - VMware_bootbank_vcls-pod-crx_8.0.3-0.73.24784735 - VMware_bootbank_native-misc-drivers_8.0.3-0.73.24784735 - VMware_bootbank_loadesx_8.0.3-0.73.24784735 - VMware_bootbank_esx-update_8.0.3-0.73.24784735 - VMware_bootbank_loadesxio_8.0.3-0.73.24784735 - VMware_bootbank_esxio-update_8.0.3-0.73.24784735
PRs Fixed	N/A
Related CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

This patch updates the following issues:

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

ESXi-8.0U3f-24784735-no-tools

Profile Name	ESXi-8.0U3f-24784735-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 15, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.73.24784735 - VMware_bootbank_vsan_8.0.3-0.73.24784735 - VMware_bootbank_infravisor_8.0.3-0.73.24784735 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.73.24784735 - VMware_bootbank_crx_8.0.3-0.73.24784735 - VMware_bootbank_bmcad_8.0.3-0.73.24784735 - VMware_bootbank_vsanhealth_8.0.3-0.73.24784735 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.73.24784735 - VMware_bootbank_gc_8.0.3-0.73.24784735 - VMware_bootbank_vdfs_8.0.3-0.73.24784735 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.73.24784735 - VMware_bootbank_esx-base_8.0.3-0.73.24784735 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.73.24784735 - VMware_bootbank_esx-xserver_8.0.3-0.73.24784735 - VMware_bootbank_vds-vsip_8.0.3-0.73.24784735 - VMware_bootbank_cpu-microcode_8.0.3-0.73.24784735 - VMware_bootbank_esxio_8.0.3-0.73.24784735 - VMware_bootbank_clusterstore_8.0.3-0.73.24784735 - VMware_bootbank_esxio-base_8.0.3-0.73.24784735 - VMware_bootbank_trx_8.0.3-0.73.24784735

	- VMware_bootbank_gc-esxio_8.0.3-0.73.24784735 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.73.24784735 - VMware_bootbank_bmcad-esxio_8.0.3-0.73.24784735 - VMware_bootbank_vcls-pod-crx_8.0.3-0.73.24784735 - VMware_bootbank_native-misc-drivers_8.0.3-0.73.24784735 - VMware_bootbank_loadesx_8.0.3-0.73.24784735 - VMware_bootbank_esx-update_8.0.3-0.73.24784735 - VMware_bootbank_loadesxio_8.0.3-0.73.24784735 - VMware_bootbank_esxio-update_8.0.3-0.73.24784735
PRs Fixed	N/A
Related CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

This patch updates the following issues:

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

ESXi 8.0 U3f - 24784735

Name	ESXi
Version	ESXi 8.0 U3f-24784735
Release Date	July 15, 2025
Category	Security
Affected Components	- ESXi Component - core ESXi VIBs - ESXi Install/Upgrade Component - ESXi Install/Upgrade Component
PRs Fixed	N/A
Related CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

This patch resolves the issues listed in ESXi_8.0.3-0.73.24784735.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

You cannot update to ESXi 8.0 Update 2b by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see [Upgrade or Update a Host with Image Profiles](#).

The Cancel option in an interactive ESXi installation might not work as expected

Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some nmlx5_core driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.

- `drss` : Use the DRSS parameter.
- `rss` : Use the RSS parameter.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

The irdman driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic

If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the irdman driver might fail. This issue is unlikely to occur, as the irdman driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

Workaround: Use ULPs with RC transport mode.

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1

The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out

vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD

In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See [Reconciling Discrepancies in the Managed Virtual Disk Catalog](#).

vCenter Server and vSphere Client Issues

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

VMware ESXi 8.0 Update 3e Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 3e | 10 APR 2025 | Build 24674464
Check for additions and updates to these release notes.

What's New

This release resolves CVE-2025-41226, CVE-2025-41227, and CVE-2025-41228. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0010](#).

Support for Communication Device Class Network Control Model (CDC-NCM) in the ESXi USB driver:

Starting with ESXi 8.0 Update 3e, the ESXi USB driver supports the CDC-NCM protocol for compatibility with HPE Gen12 iLO Virtual NIC and interoperability with HPE Agentless Management (AMS), Integrated Smart Update Tools (iSUT), the iLOrest config tool, Intelligent Provisioning, and DPUs.

ESXi 8.0 Update 3e adds support for vSphere Quick Boot to:

- Intel vRAN Baseband Driver
- Intel Platform Monitoring Technology Driver
- Intel Data Center Graphics Driver
- AMD Instinct MI Series Driver

Broadcom makes available the VMware vSphere Hypervisor version 8, an entry-level hypervisor. No Broadcom support is available for this offering and it is for non-production use. vSphere Hypervisor cannot connect to vCenter and therefore cannot be centrally managed. You can remotely manage individual vSphere Hypervisor hosts by using the VMware Host Client. vSphere Hypervisor supports a maximum of 8 virtual CPUs per virtual machine. You can download it free of charge from the [Broadcom Support portal](#).

Patches Contained in This Release

VMware ESXi 8.0 Update 3e

Build Details

VMware vSphere Hypervisor (ESXi) Offline Bundle

Download Filename:	VMware-ESXi-8.0U3e-24674464-depot.zip
Build:	24674464
Download Size:	1012.8 MB
SHA256 checksum:	cd3301c8cba325787942cf1619a7ef8ab529818091645f66ab63997776a5dde4
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.3-0.70.24674464	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.70.24674464	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.70.24674464	Bugfix	Critical

Component	Bulletin	Category	Severity
Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers	Mellanox-nmlx5_4.23.6.5-1vmw.803.0.70.24674464	Bugfix	Critical
Broadcom NetXtreme I ESX VMKAPI ethernet driver	Broadcom-ntg3_4.1.15.0-4vmw.803.0.70.24674464	Bugfix	Critical
Broadcom Emulex Connectivity Division FC and FCoE Driver	Broadcom-ELX-lpfc_14.4.0.40-35vmw.803.0.70.24674464	Bugfix	Critical
Intel NVME Driver with VMD Technology	Intel-Volume-Mgmt-Device_3.0.0.1010-12vmw.803.0.70.24674464	Bugfix	Critical
VMware NVMe over TCP Driver	VMware-NVMeoF-TCP_1.0.1.32-1vmw.803.0.70.24674464	Bugfix	Critical
USB Native Driver for VMware	VMware-vmkusb_0.1-23vmw.803.0.70.24674464	Bugfix	Critical
ESXi component with core ESXi VIBs	ESXi_8.0.3-0.65.24659227	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.65.24659227	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.65.24659227	Security	Critical
ESXi Tools Component	VMware-VM-Tools_12.5.1.24649672-24659227	Security	Critical

Rollup Bulletins

These rollup bulletins contain the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U3e-24674464	Bugfix	Critical	Security and Bugfix image
ESXi80U3se-24659227	Security	Critical	Security only image

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U3e-24674464-standard
ESXi-8.0U3e-24674464-no-tools
ESXi-8.0U3se-24659227-standard
ESXi-8.0U3se-24659227-no-tools

ESXi Images

Name and Version	Release Date	Category	Detail
ESXi 8.0 U3e - 24674464	10 APR 2025	General	Security and Bugfix image
ESXi 8.0 U3se - 24659227	10 APR 2025	Security	Security only image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

See these KBs for details on the authenticated online and offline download: [390121](#), [390098](#), and [389276](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually download the patch offline bundle ZIP file or see the KB [How to download ESXi ISO image for all releases including patch updates](#).

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.3-0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.70.24674464 - VMware_bootbank_vsan_8.0.3-0.70.24674464 - VMware_bootbank_infravisor_8.0.3-0.70.24674464 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.70.24674464 - VMware_bootbank_crx_8.0.3-0.70.24674464 - VMware_bootbank_bmcad_8.0.3-0.70.24674464 - VMware_bootbank_vsanhealth_8.0.3-0.70.24674464 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.70.24674464 - VMware_bootbank_gc_8.0.3-0.70.24674464 - VMware_bootbank_vdfs_8.0.3-0.70.24674464 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.70.24674464 - VMware_bootbank_esx-base_8.0.3-0.70.24674464 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.70.24674464 - VMware_bootbank_esx-xserver_8.0.3-0.70.24674464 - VMware_bootbank_vds-vsip_8.0.3-0.70.24674464 - VMware_bootbank_cpu-microcode_8.0.3-0.70.24674464 - VMware_bootbank_esxio_8.0.3-0.70.24674464 - VMware_bootbank_clusterstore_8.0.3-0.70.24674464 - VMware_bootbank_esxio-base_8.0.3-0.70.24674464 - VMware_bootbank_trx_8.0.3-0.70.24674464 - VMware_bootbank_gc-esxio_8.0.3-0.70.24674464 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.70.24674464 - VMware_bootbank_bmcad-esxio_8.0.3-0.70.24674464

	- VMware_bootbank_vcls-pod-crx_8.0.3-0.70.24674464 - VMware_bootbank_native-misc-drivers_8.0.3-0.70.24674464
PRs Fixed	3498633, 3341273, 3410296, 3414552, 3415171, 3425316, 3427281, 3480453, 3431433, 3432906, 3433295, 3433436, 3433610, 3434171, 3434270, 3437835, 3437903, 3438291, 3439279, 3444039, 3444137, 3444408, 3445065, 3445276, 3446362, 3447505, 3447686, 3447687, 3447888, 3448872, 3449523, 3450290, 3450374, 3450501, 3451651, 3451669, 3451849, 3452081, 3452460, 3452487, 3452906, 3453391, 3453981, 3454108, 3454123, 3454926, 3454942, 3454983, 3455021, 3455315, 3456018, 3457039, 3457042, 3457050, 3457350, 3457744, 3457900, 3457959, 3458521, 3459100, 3459378, 3459526, 3459825, 3460312, 3460819, 3463099, 3463147, 3463364, 3465047, 3465070, 3465259, 3465275, 3465373, 3465624, 3465965, 3466011, 3466176, 3466283, 3466538, 3467137, 3467521, 3467708, 3468136, 3468139, 3468215, 3469014, 3469036, 3469145, 3469467, 3469652, 3470385, 3470427, 3470500, 3470714, 3470732, 3472623, 3473154, 3473626, 3473692, 3473767, 3473768, 3473860, 3474427, 3474443, 3474967, 3475015, 3475152, 3475393, 3476179, 3476630, 3476631, 3476632, 3476634, 3476963, 3477021, 3477062, 3477385, 3477409, 3477725, 3478001, 3478002, 3478236, 3478276, 3478486, 3479338, 3479340, 3479464, 3480154, 3480456, 3480968, 3481216, 3481241, 3481310, 3481386, 3481475, 3481507, 3482259, 3482802, 3482963, 3483589, 3483610, 3483714, 3484669, 3484787, 3484836, 3485671, 3485971, 3486600, 3486647, 3486747, 3486955, 3487351, 3487380, 3488415, 3488480, 3494691, 3497983, 3462191, 3442016, 3478908, 3444303, 3454941, 3469074, 3473257, 3458162, 3474996
CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

This patch updates the esxio-combiner, vsan, infravisor, esxio-dvfilter-generic-fastpath, crx, bmc, vsanhealth, drivervm-gpu-base, gc, vdfs, native-misc-drivers-esxio, esx-base, esx-dvfilter-generic-fastpath, esx-xserver, vds-vsip, cpu-microcode, esxio, clusterstore, esxio-base, trx, gc-esxio, esxio-combiner-esxio, bmc-esxio, vcls-pod-crx, and native-misc-drivers VIBs.

This patch resolves the following issues:

PR 3476963: An ESXi module might fail to load due to exceeding symbol size limit

If the symbol size of an ESXi module, such as an NVIDIA driver, exceeds an internal limit, loading such module fails with a vmkernel error log such as *module symbol space too large*.

This issue is resolved in this release.

PR 3447687: Virtual machines cannot power on or migrate in a VxRail environment after a VIB change

If you install or remove a VIB on an ESX host in a VxRail environment followed by the `backup.sh 1` command without rebooting the host, an issue in the ConfigStore framework might cause the host to enter into an inconsistent state. As a result, virtual machines on such hosts cannot power on or migrate.

This issue is resolved this release.

PR 3442191: Update of a vSAN environment with Original Storage Architecture (OSA) to version 8.0 Update 3 might be slower compared to previous updates

The first reboot of the vSphere system in a vSAN OSA environment following upgrade or update to 8.0 Update 3 might cause a delay in the update process, which depends on the workload running at the time of update. Nodes with heavy workloads running before the update are likely to cause higher delays. This is expected and for updates to 8.0 Update 3, environments with large install base must consider the possible delay in the update process.

This issue is resolved in this release.

PR 3479338: Delayed cleanup of a deleted component can lead to high space usage reports for vSAN datastores

The vSAN I/O scheduler manages in-flight I/Os for performance and fairness. Under certain conditions, the scheduler might get an incorrect counter for the in-flight resync I/Os, which blocks new resync I/Os. As a result, the cleanup of blocked resync I/Os is delayed and might lead to high space usage.

This issue is resolved in this release.

PR 3476179: A memory leak might prevent the Object Storage File System Daemon (osfsd) from processing requests for vSphere Virtual Volumes

A rare memory leak condition in the osfsd service after prolonged uptime might prevent vSphere Virtual Volumes from provisioning virtual machines, including creation and cloning operations.

This issue is resolved in this release.

PR 3459100: Virtual machine migration in a vSphere Virtual Volumes datastore might cause an ESXi host to fail with a purple diagnostic screen

During virtual machine migration from a vSphere Virtual Volumes datastore, a rare race condition in Device Block allocation tracking might lead to an ESXi host failure, resulting in a purple diagnostic screen.

This issue is resolved in this release. The fix prevents the race condition.

PR 3484787: If you reconfigure the scratch location and reboot an ESXi host, all logging stops and the host becomes unresponsive

If you have audit recording to local storage enabled for an ESXi host on a directory configured at the default location `/scratch/auditLog` or at a subdirectory of `/scratch`, and then reconfigure `/scratch`, the syslog daemon fails after a host reboot, because it cannot find the record directory and files in the new location. As a result, all logging on such hosts stops and eventually the host becomes unresponsive.

This issue is resolved in this release. The fix handles such scenarios and deactivates audit records to local storage for security reasons, generating a VMkernel Observation (VOB) message such as `vSphere error: (esx.problem.vmsyslogd.auditrecord.local.disabled)` but preventing a failure of the vmsyslog daemon. You can re-enable audit recording in the new `/scratch` location after the host reboots. Old audit record files are available in the previously configured scratch path, if it is still accessible.

PR 3455315: The PowerSvc plug-in of the sensord service stops polling data by using Intelligent Platform Management Interface (IPMI) without a retry

Due to a change related to the Running Average Power Limit (RAPL) host power reporting in ESXi 8.0 Update 3, the sensord service might deactivate power monitoring by using IPMI if only one IPMI power data poll fails. The PowerSvc plug-in does not retry the poll and IPMI polling stops. PowerSvc goes offline and logs `failing to read from the power provider in the /var/log/sensord.log`.

This issue is resolved in this release. The fix makes sure the PowerSvc plug-in retries IPMI power data polls in case of failure and IPMI data collection continues.

PR 3460819: Virtual machines with UnixWare OS and a vlane virtual device might repeatedly reset when running on an ESXi 8.x host

An issue with vance virtual devices might cause virtual machines with UnixWare OS on ESXi 8.x hosts to repeatedly fail due to a reboot of the guest OS with the error `Your guest operating system has accessed an I/O range where multiple devices overlap`.

This issue is resolved in this release.

PR 3479340: You see checksum errors for some objects in a stretched cluster after a network outage

After a network outage or some other critical event on a stretched cluster, the vSAN Distributed Object Manager (DOM) might fail to get the correct checksums for objects that do not have sub-fault tolerance. As a result, you might see in tools such as VMware Skyline Health Diagnostics checksum errors for many objects in the stretched cluster.

This issue is resolved in this release.

PR 3461166: In the vSphere Client, you cannot change a VMNIC on the virtual switch of an ESXi host

The cache of the hostd service might not be up-to-date after you add a VMNIC to the virtual switch on an ESXi host due to a bug in the teaming policy of port groups. As a result, when you remove another VMNIC from the switch, both the newly added and the removed VMNICs are removed from the teaming policy of the port groups. When you work in the vSphere Client, the task seems to revert.

This issue is resolved in this release.

PR 3463099: Guest OS might mark file system read-only after repeated IO failure which can cause Kubernetes control plane nodes to go into Read-only mode

An issue with the check for max allowed memory address in memory mapping code during I/O processing at virtual SCSI might cause a guest OS on an ESXi host to mark guest file system as read-only. As a result, Kubernetes control plane nodes might go into Read-only mode. In the vmkernel logs, you see multiple error messages such as `failed to pin: I/O error`.

This issue is resolved in this release.

PR 3461321: Custom TLS settings of the Fault Domain Manager (FDM) might not work on ESXi hosts of earlier versions that you manage with a vCenter of version 8.0 Update 3

Custom TLS settings of the vmware-fdm service at port 8182 might not work on hosts of earlier versions in a cluster with active High Availability that you manage with a vCenter of version 8.0 Update 3.

This issue is resolved in this release. For more information on the workflows and the use of such custom settings, see KB [320798](#).

PR 3488415: ESXi host fails with a purple diagnostic screen with messages such as #PF Exception 14 in world XXXXXXXX:vmknvmeGener

In the discovery phase of NVMe over Fabrics (NVMe-oF), some targets might report to an ESXi host some subsystem NVMe Qualified Names (NQN) to which that ESXi host cannot connect due to missing permissions or unavailable route. The ESXi host continues attempting to connect to such NQNs and if a network issue occurs, in rare cases, the host might fail with a purple diagnostic screen with messages such as `#PF Exception 14 in world XXXXXXXX:vmknvmeGener`.

This issue is resolved in this release.

PR 3486647: Copying files with the cp command creates files with unexpected access timestamp (atime) and modification timestamp (mtime) values

If vSAN File Services NFS 4.1 is mounted to a RHEL 8 or later NFS client, when copying a file by using the `cp` command, you might see `atime` or `mtime` such as 1970 instead of the current date.

This issue is resolved in this release.

PR 3486025: Virtual machines and interactive shells might become unresponsive due to an issue with the initialization of user worlds in ESXi

By default, the initialization of user worlds in ESXi adds an active group that requires the allocation of some CPU resources. If the initialization fails for some reason before the CPU resource is available, the creation of new user worlds starts failing. As a result, virtual machines and interactive shells might become unresponsive.

This issue is resolved in this release.

PR 3469652: Logging stops on all log files and configured remote host targets on ESXi hosts and hosts might become unresponsive

A rare deadlock situation in the ESXi log daemon vmsyslogd might cause logging to stop on all log files and configured remote host targets on ESXi hosts. As a result, the hosts might become unresponsive. The issue might occur as the result of multiple syslog reloads.

This issue is resolved in this release.

PR 3466011: You see the datastore name in storage VMkernel Observation (VOB) reports as Unknown

Due to a missing function to grant access required to fetch datastore names, you might see the datastore name in `esx.problem.storage.vob` messages as unknown.

For example: 2024-08-27T04:31:29.150Z In(14) vobd[2097668]: [scsiCorrelator] 5815563033719us: [esx.problem.storage.redundancy.lost] Lost path redundancy to storage device naa.60060e80072b600000302b6000000026. Path vmhba1:C0:T1:L0 is down. Affected datastores: Unknown.

This issue is resolved in this release.

PR 3482259: NVMe over TCP controllers might stay offline and not recover even when the network and the target port have no issues

Due to issues with the network or the target, NVMe over TCP controllers in an ESXi host might enter an offline state. Usually, controllers recover and turn online as soon as the issue is resolved, but in some very rare scenarios, the controllers might stay offline and do not recover, which makes the NVMe/TCP namespaces inaccessible. The issue occurs due to a dead lock when many controllers are being reset and at the same time many I/O requests time out, which triggers multiple taskmgmt requests that compete for critical resource. The only way to exit this state is to reboot the ESXi host.

This issue is resolved in this release. The fix avoids such dead locks.

PR 3481216: "Route based on physical NIC load" policy does not work as expected

An issue with the permissions for the net-blt daemon might cause load balancing based on the **Routing based on PNIC load** teaming policy to not work as expected.

This issue is resolved in this release.

PR 3486955: Storage path unclaiming requests might be blocked due to a Fabric Performance Impact Notification (FPIN) issue

FPIN messages serve as an early warning system to relieve negative effects on the performance of Fibre Channel devices, such as optimizing I/O behavior and avoiding impaired paths by notifying devices on the current fabric conditions. In some cases, when memory allocation fails, FPIN might keep a reference on a storage path. As a result, path scanning and claiming routine returns `BUSY`, as the reference no longer exists, which blocks path claiming and unclaiming, and FPIN-related information cannot go through to ESXi hosts. In the ESXCLI command output you do not see FPIN messages. When the issue occurs, you see a message in the VMkernel logs such as: `WARNING: StorageFPIN: 521: Failed to allocate memory.`

Another symptom is when after a SAN maintenance or in case of a storage path outage FPIN holds a reference count on the storage paths and the FC driver cannot allocate new paths or reestablish existing ones. In `/var/log/vmkernel.log`, you see the following sequence repeating:

WARNING: nfnic: <2>: fnic_handle_report_lun: 1467: lun add failure! in_remove: 0
ioAllowed: 1

WARNING: nfnic: <2>: fnic_tport_event_handler: 2130: lunmap update failed, retry ...

This issue is resolved in this release.

PR 3341273: Migration of virtual machines fails with an error "Failed to allocate migration heap" on the source ESXi host

In very rare cases, if a virtual machine powers off before a running migration completes, ESXi might fail to clean up the migration heap. If the memory leak continues over time, the migration heap might be exhausted, and virtual machines cannot migrate from such ESXi hosts.

This issue is resolved in this release.

PR 3467137: A race condition during heap cleanup might cause an ESXi host to fail with a purple diagnostic screen

A rare race condition when reading the VMX file descriptor table occurs at the time VMX shuts down might cause an ESXi host to fail with a purple diagnostic screen. The error message is similar to `PANIC bora/vmkernel/main/dlmalloc.c:4944 - Usage error in dlmalloc.`

This issue is resolved in this release.

PR 3483610: If you insert an invalid or blank DVD on an ESXi host and configure it as a host device on virtual machines, the VM performance degrades

If a blank or incompatible disc is in the DVD drive of an ESXi host, the ESXi HPP module checks for the SCSI sense key and additional sense keys to determine next steps. Due to a missing check to handle Active State Change (ASC) codes 0x30 and 0x3a, the check fails, causing a failover that keeps repeating. As a result, you see a degradation of performance on the virtual machines on such hosts.

This issue is resolved in this release. The fix handles the ASC.

PR 3475015: Duplicate traffic on multiple uplinks might cause traffic outage

During an upgrade of the physical switch infrastructure, if you use the `TeamPolicyUpDelay` setting and all uplinks on the distributed switch flap, a fallback mode during the time of delay might cause duplicate traffic on all uplinks of the switch, which leads to traffic outage.

This issue is resolved in this release. The fix adds the advanced config parameter `/Net/TeamingIgnoreShotgun` to prevent the fallback mode in such cases and avoid duplicate traffic.

PR 3451306: The usage scale in the CPU chart under Monitor > Performance Overview shows in thousands, and you cannot notice CPU usage spikes to 100%

The usage scale in the CPU chart under **Monitor > Performance Overview** for virtual machines might show in thousands, reaching up to 2400%, instead of hundreds to allow you to notice CPU usage spikes to 100%. The expanded scale also prevents realistic estimation of other key performance metrics on the graph, such as CPU Ready, as they are displayed at the bottom of the chart.

This issue is resolved in this release.

PR 3470732: Virtual machines using Raw Device Mapping (RDM) might become unresponsive due storage heap exhaustion

An issue in tracking SCSI Inquiry or Mode Sense requests that virtual machines issue to RDM LUNs might lead to heap exhaustion at the ESXi storage layer. As a result, such VMs become unresponsive. In the vSphere Client, you see series of **Lost access to volume** and **Successfully restored access** messages until you power off the VMs.

This issue is resolved in this release.

PR 3474443: A memory leak in the swapObj daemon (swapobjd) during object replication might cause a failure of the service due to stale metadata

An issue with stale metadata for swap objects might cause a memory leak in the swapobjd during object replication, leading to gradual memory exhaustion and ultimately to a failure of the service.

This issue is resolved in this release. The fix ensures stale swap object metadata is removed.

PR 3479464: ESXi hosts might fail with purple diagnostic screen while processing SCSI commands from the guest OS

An issue with the validation of SCSI commands from the guest OS at the virtual SCSI layer might cause an ESXi host to fail with a purple diagnostic screen.

This issue is resolved in this release. The fix adds enhanced command validation for SCSI commands at the virtual SCSI layer.

PR 3481507: Discovery of iSCSI dynamic targets, LUNs, and datastores fails on ESXi host reboot

In some cases, dynamic discovery of targets in systems that use dynamic discovery addresses for iSCSI configuration might fail. As a result, you do not see iSCSI targets, LUNs, and datastores after a reboot of ESXi hosts. The dynamic discovery of iSCSI targets might also fail when the Challenge-Handshake Authentication Protocol (CHAP) is active on the ESXi host.

This issue is resolved in this release.

PR 3454942: In the vSphere Client, you cannot reconfigure a vGPU virtual machine if more than 128 vCPUs are active

If you add more than 128 vCPUs on a VM with a vGPU, you cannot manage the configuration from the vSphere Client. The issue occurs because input-output memory management unit (IOMMU) support for a vGPU with more than 128 vCPUs is not available.

This issue is resolved in this release. With the fix, you can manage the vGPU VM configuration from the vSphere Client. However, the VM might not be able to power on, since IOMMU support is not available. To make sure the reconfiguration succeeds, add to the .vmx file the following 2 parameters: `pciPassthru.vmiop.allowViommu="TRUE"` and `pciPassthru.vmiop.enableViommu="TRUE"`.

PR 3477772: Encrypted virtual machines with active Change Block Tracking (CBT) might intermittently power off after a rekey operation

Due to a race condition between the VMX and hostd services for a specific VMcrypt-related reconfiguration, encrypted VMs with active CBT might unexpectedly power off during a rekey operation.

This issue is resolved in this release.

PR 3469014: ESXi host might fail with a purple diagnostic screen if the Multipath Plug-in (MPP) does not implement the event handler for Fabric Performance Impact Notification (FPIN)

Description: ESXi 8.0 Update 3 added APIs to notify MPPs in case of Fibre channel FPIN events. However, if the MPP is compiled against an API SDK of version earlier than 8.0 Update 3, or if the MPP does not implement the new event handler, then the ESXi host fails with a purple diagnostic screen on an FPIN event.

This issue is resolved in this release.

PR 3465259: NVMe Dell PowerStore clustered VMDK datastore fails to recover after an All-Paths-Down (APD) situation

A reservation conflict might prevent a NVMe Dell PowerStore clustered VMDK datastore to take a reservation and cause failure to recover after losing the reservation.

This issue is resolved in this release.

PR 3480154: High VMFS heap memory consumption might cause virtual machines to become unresponsive

In environments with many large datastores and many VMs on thin-provisioned disks, VMFS resource allocation might consume a high amount of heap memory, proportional to the size of the datastores. As a result, VMFS heap memory might be exhausted and lead to VMs becoming unresponsive or vSphere vMotion tasks failing with errors for VM invalid state.

This issue is resolved in this release.

PR 3460312: Virtual machines on NFSv3 datastores might fail during snapshot consolidation if a third-party backup is also active

Snapshot consolidation during backup with third-party software, such as NetBackup, might fail due to extra read-only locks on the VM snapshot file while the snapshot consolidation is in progress and trying to upgrade that same lock. As a result, a lock conflict might cause a VM to fail.

This issue is resolved in this release. The fix prevents snapshot consolidation tasks to upgrade the read-only lock when more than 1 holders for that lock exist.

PR 3457744: Migration of virtual machines to 8.0 Update 3 and later ESXi host might result in TPM null hierarchy issues

A rare issue related to how applications use TPM null hierarchy might lead to failure to authenticate virtual machines after migrating them from an ESXi host of version earlier than 8.0 Update 3 to 8.0 Update 3 and later. You can identify the issue if the command `Get-TpmSupportedFeature` in Windows returns NULL or the command `sudo tpm2_createprimary -C n -c context.out` in Linux fails.

This issue is resolved in this release. If you already face the issue, power cycle the virtual machines.

PR 3470385: Restoring the iSCSI configurations on an ESXi host might fail during boot due to stale entries in the ConfigStore database

If the ConfigStore database on an ESXi host has stale iSCSI configurations related to hardware iSCSI adapters, the iSCSI jumpstart might fail and restoring the iSCSI configuration on that host also fails during boot.

This issue is resolved in this release.

PR 3477062: vSAN hosts might become unresponsive due to high latencies and high CPU usage during peaks of guest unmap traffic

Internal threads in vSAN might cause a CPU busy loop during a peak of unmap traffic. As a result, some hosts might become unresponsive.

This issue is resolved in this release. The fix is relinquishing CPUs when unmap requests start to slow down.

PR 3466538: A race condition between the TCP control path and the keep alive timer might cause an ESXi host to fail with a purple diagnostic screen

A rare race condition between the TCP control path during a disconnect operation and the keep alive timer might cause ESXi hosts of version 8.0 Update 3 to fail with a purple diagnostic screen. The issue occurs under specific network conditions when the timer and disconnect logic overlap.

This issue is resolved in this release. The fix prevents the purple diagnostic screen failure.

PR 3473626: If a command issued to a vSAN disk group complete after 120 seconds, the vSAN host might fail with a purple diagnostic screen

In a vSAN configuration, if a command issued to a vSAN disk group takes more than 120 sec to complete, in a corner case the vSAN host might fail with a purple diagnostic screen. The issue occurs because after 120 seconds a red alert that indicates stuck I/O condition is triggered to vSAN and the resources for that command are released. As a result, if a given command completes after this alert, the host fails.

This issue is resolved in this release.

PR 3459526: ESXi boot might take longer on hosts with iSCSI configuration

On ESXi hosts with iSCSI configuration, the dynamic discovery of iSCSI targets might time out or fail because the firewall related to the iSCSI ruleset is not active during early boot. As a result, if many discovery addresses are configured, ESXi boot takes more time as discovery is done in a serial order.

This issue is resolved in this release.

PR 3447505: Nonuniform Sub-NUMA Clustering (SNC) sizes might cause ESXi hosts to fail with a purple diagnostic screen post upgrade to ESXi 8.0 Update 3

On ESXi hosts that use the SNC technology, if the CPU is not manufactured with equal number of cores per sub-NUMA cluster, or some cores are deactivated from the BIOS, the host fails to boot with a purple diagnostic screen post upgrade to ESXi 8.0 Update 3.

This issue is resolved in this release. Instead of replacing the hardware, ESXi 8.0 Update 3e adds a software fix of the issue when CPUs have nonuniform sub-NUMA cluster size. If the root cause is that CPUs have cores deactivated from the BIOS, just activate such cores.

PR 3486556: You cannot retrieve Active Directory domain accounts from ESXi hosts by using PowerCLI

Due to changes in the Open LDAP library, the `Get-VIAccount` PowerCLI script might not run correctly and produce the following error:

```
A general system error occurred: Error accessing directory: Can't set LDAP options.
```

As a result, you cannot retrieve information about Active Directory accounts from ESXi hosts joined to the domain.

This issue is resolved in this release.

PR 3484669: You see intermittent gaps in vSAN performance statistics

In certain cases, you can see intermittent gaps in vSAN performance statistics dashboard charts. The issue occurs mostly when the vSAN cluster is integrated with VMware Aria Operations.

This issue is resolved in this release.

PR 3469036: ESXi installation on an iSCSI LUN fails as the ESXi iSCSI Boot Firmware Table (iBFT) module fails to discover the LUN

In some cases, the iBFT control structure might have an optional expansion for which the ESXi `vmkibft` module doesn't have a proper check. As a result, iSCSI LUN discovery by using iBFT might fail and cause failure of ESXi installation on the iSCSI LUN.

This issue is resolved in this release.

PR 3450501: ESXi hosts working with both VMFS5 and VMFS6 datastores might fail with a purple diagnostic screen due to rare issue

On ESXi hosts that work with both VMFS5 and VMFS6 datastores, a rare issue with the I/O completion path might cause incorrect access to the data structure of a VMFS5 datastore instead of a VMFS6 datastore. As a result, the host fails with a purple diagnostic screen. The issue affects only hosts working with both VMFS5 and VMFS6 volumes and can occur without any user intervention.

This issue is resolved in this release.

PR 3433295: Virtual machines residing on a vSphere Virtual Volumes datastore might become inaccessible due to a size reporting issue

An issue in the VASA provider might report an inflated size for the home directory that contains metadata files for a virtual machine and VMFS automatically expands to accommodate the larger size. However, if the VASA provider later reports the correct size, VMFS cannot shrink, and the home directory cannot open. As a result, all VM files stored in the home folder become inaccessible and you cannot access the VMs.

This issue is resolved in this release.

PR 3431433: Storage tests might fail when using NVMe as Guest Controllers with SCSI vSphere Virtual Volumes on Windows Server Failover Clusters (WSFC)

ESXi 8.0 Update 3 added support for NVMe reservation on NVMe-backed disks on vSphere Virtual Volumes in WSFC. However, NVMe guest controllers with SCSI-backed vSphere Virtual Volumes disks are not supported and this might result in failure while running WSFC cluster validation for storage tests. The specific test that might fail is `Validate SCSI-3 Persistent Reservation`. The issue occurs when the storage array does not support implementation of the `SCSI Read_full_status Persistent Reservation` command.

This issue is resolved in this release.

PR 3483589: In very rare cases, ESXi hosts might fail with a purple diagnostic screen due to a SCSI-3 reservation failure

In a very corner case scenario, ESXi hosts with active support of SCSI-3 reservation type WEAR (write exclusive all registrant) for VMFS might fail with a purple diagnostic screen due to an issue with handling the SCSI-3 reservation failures.

This issue is resolved in this release.

PR 3441150: In the vSphere Client, you see unusually high storage utilization of vSphere Replication persistent state files (.psf) in a vSAN datastore

For vSAN datastores, the provision type of .psf depends on the configured namespace storage policy and if it is set to Thick Provisioning, the space reservation of such .psf files might seem unusually high, regardless of the actual size of vSAN objects. For example, you might see a .psf file reservation for tens of TB in the vSphere Client, while the actual vSAN storage is in tens of GB.

This issue is resolved in this release. The fix makes sure .psf files always use Thin Provisioning policy.

PR 3450290: You cannot manually deactivate the DVSSync firewall ruleset

The DVSSync firewall ruleset is part of the system-owned rulesets that are active by default to avoid possible break of service but in some cases, you might need to deactivate the rule for security reasons.

This issue is resolved in this release. You can manually deactivate the ruleset as required.

PR 3451651: ESXi host fails with a purple diagnostic screen after unsuccessful rebind operation of vSphere Virtual Volumes

In the rebind orchestration of vSphere Virtual Volumes, all ESXi hosts that have bindings to the volume must respond in a time window defined by the number of Read Block Zero (RBZ) calls that the hosts send. If a migration target fails to complete the RBZ in time, ESXi hosts might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3465047: You see no UNMAP commands sent to NVMe 2.0 targets or targets supporting TP4040

The `identify controller` command to fetch the `Unmap` attribute from a target might fail with a status `0xf` due to a buffer mismatch on the targets which support TP4040 or NVMe 2.0. As a result, you see no UNMAP commands sent to NVMe 2.0 targets or targets supporting TP4040. In the vmkernel log, you see errors such as:

```
command <cmd> failed: ctrlr <ctrlnum>, queue <queue num>, psaCmd <psa cmd>, status
0xf, opc 0x6, cid <cid>, nsid <nsid>Complete vmkNvmeCmd: <cmd>, vmkPsaCmd: <psa cmd>,
cmdId.initiator=<init>, CmdSN: 0x0, status: 0xfGet DSM constraints failed for device
<device name> Status Transient storage condition, suggest retryUnable to fetch UNMAP
Attributes for device <device name>
```

This issue is resolved in this release.

PR 3427281: You see a Skyline Health alert for NICs outside the vSAN configuration

When you have active Remote Direct Memory Access (RDMA) for vSAN clusters and any cluster host has NICs outside the vSAN configuration, the health check **Network (RDMA NIC: RoCE v2) is vSAN certified** also lists the result for such NICs. This is not expected, but you can ignore the issue as it doesn't indicate any data path issue.

This issue is resolved in this release.

PR 3433436: In the vSphere Client, you see multiple alarms "Cannot login user root@xxxxx: no permission" for ESXi hosts in lockdown mode

In the **Tasks and Events** tab, you see multiple alarms `Cannot login user root@xxxxx: no permission for ESXi hosts in lockdown mode`.

This issue is resolved in this release.

PR 3433610: ESXi hosts fail with purple diagnostic screen when connecting to the network I/O scheduler

In very rare cases, the network I/O scheduler might hold a lock for a certain grace period in case of memory allocation failure during the connection. This lock might cause ESXi hosts to fail with a purple diagnostic screen and an error such as `VERIFY bora/vmkernel/sched/cpushed.c`.

In the backtrace, you see errors such as:

```
2024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2a0:
[0x420006423c31]RCU_WaitForGracePeriod@vmkernel#nover+0x42 stack: 0x1000000, 0x300000100,
0xbad0014, 0x420006518bc2, 0x0

2024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2d0:
[0x420006518bc1]vmk_HashRelease@vmkernel#nover+0x66 stack: 0x420006ac6eb5, 0x420007175487,
0x516574616572436b, 0x4307194064d0, 0x3030306

22024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2f0:
[0x420007175486]NetSchedHClkConnect@(netsched_hclk)#<None>+0x6ab stack: 0x30303062,
0x4300d3f27f48, 0x43126525a638, 0x0, 0x26
```

This issue is resolved in this release.

PR 3437835: AMD Turin processors with PCI Multi-segment Support preset in the BIOS might cause ESXi host failures

In AMD 5th generation EPYC 9005 series with code name Turin, **PCI Multi-segment Support** is active in the BIOS settings by default and as a result, you might see the following symptoms after installation:

- The ESXi host fails with a purple diagnostic screen.
- The ESXi host becomes unresponsive.
- Some I/O devices do not work.

The issue occurs because AMD Turin has preset PCIe multi-segment support, but ESXi does not provide PCIe multi-segment support for CPUs that have an input-output memory management unit (IOMMU). This issue is specific for AMD Turin CPUs.

This issue is resolved in this release.

PR 3469107: ESXi hosts intermittently disconnect from the Active Directory domain or from vCenter

Memory leaks might occur in Likewise during Active Directory operations or if smart card authentication is active on an ESXi host. As a result, Likewise processes might run out of memory and ESXi hosts intermittently disconnect from the Active Directory domain or from vCenter. In the `/var/log/vmkernel.log` file, you see messages similar to:

```
yyyy-mm-dd In(182) vmkernel: cpu72:2110465)uw.2110464 (44739) requires 1024 KB, asked 1024
KB from likewise (792) which has 93112 KB occupied and 72 KB available.

yyyy-mm-dd In(182) vmkernel: cpu72:2110465)Admission failure in path: host/vim/vmvisor/
likewise:lwsmd.2110464:uw.2110464
```


In the 'hostd-probe.log' file you see a message such as: `hostd detected to be non-responsive`.

This issue is resolved in this release. The fix adds a system service that monitors Likewise memory usage and restarts relevant services in case memory consumption approaches Likewise limits.

PR 3463364: Storage policy change in a vSAN ESA cluster with less than six hosts might cause performance issues during vSAN resync

If you try to change the storage policy of a large object from RAID-1 to RAID-5 in a vSAN ESA cluster with less than six hosts, the change might never complete. As a result, you might see degraded performance of the cluster and host, or disk maintenance operations might be blocked. The issue occurs because the policy change of a large object in a small vSAN ESA cluster executes in steps in smaller regions of the object and might cause repeating attempts to change the policy to RAID-5 of a region even if the change is already complete. This issue might ultimately cause VMs to become unresponsive during vSAN resync.

This issue is resolved in this release. The fix makes sure vSAN tracks storage policy changes correctly to prevent repeating attempts. If you do not update to ESXi 8.0 Update 3e, avoid changing the object policy from RAID-1 to RAID-5 in small vSAN ESA clusters.

PR 3477409: vSphere vMotion operations might fail due to a rare issue with no namespaces present in the namespacesmgr.db file

When a virtual machine is configured with namespaces and for some reason the namespaces table in the `namespacemgr.db` file is empty in the ESXi host database, migration of such a VM by using vSphere vMotion might fail. In the `vmware.log` file, you see errors such as:

```
vmx - NamespaceMgrCheckpoint: No valid queue found while restoring the namespace events.
The migrate type is 1.
```

This issue is resolved in this release.

PR 3456018: ESXi hosts fail with purple diagnostic screen and message such as "PCI bus error, undiagnosed"

On ESXi hosts with active PCIe multi-segment support and hot-plug slots with active PCIe Downstream Port Containment (DPC), any PCIe Uncorrectable Error reported by NVME drives plugged into the hot-plug slots can result in failure of the host with a purple diagnostics screen. The error message is `PCI bus error, undiagnosed`. This may be a hardware problem; please contact your hardware vendor.

This issue is resolved in this release.

PR 3444408: Allocation failure in Physical Region Pages (PRP) during I/O split might cause stuck I/O

Starting with vSphere 8.0 Update 3, the handling of I/O for PRP, which most local NVMe devices use to describe user data location, was transferred from the driver to the Pluggable Storage Architecture (PSA). A bug in PSA PRP handling, when a command is split into multiple child commands to handle the PRP requirements, might cause indefinitely stuck I/O. The issue occurs because if a child PRP memory allocation command fails, the parent command cannot complete. As a result, you might see an error such as `BlueScreen: ASSERT bora/modules/vmkernel/vmfs/fs3Misc.c:4245`. The issue affects only local PCIe NVMe disks.

This issue is resolved in this release. The fix adds the frame back in the path queue to make sure all child commands are processed and fixes the PRP child count when I/O fails due to issues such as path loss or no connectivity.

PR 3450374: A rare race between rebind and lazy unbind tasks in vSphere Virtual Volume might cause an ESXi host to fail with a purple diagnostic screen

In rare cases, a lazy unbind of a vSphere Virtual Volume might start while a rebind operation is ongoing, or the other way around. As a result, the ESXi host fails with a purple diagnostic screen.

This issue is resolved in this release.

PR 3445276: Storage path unclaiming requests might be blocked due to a Fabric Performance Impact Notification (FPIN) issue

FPIN messages serve as an early warning system to relieve negative effects on the performance of Fibre Channel devices, such as optimizing I/O behavior and avoiding impaired paths by notifying devices on the current fabric conditions. In some cases, when memory allocation fails, FPIN might keep a reference on a storage path. As a result, path scanning and claiming routine returns BUSY, as the reference is no longer existing, which blocks path claiming and unclaiming. When the issue occurs, you see the similar logs in vmkernel.log: `WARNING: StorageFPIN: 521: Failed to allocate memory.`

This issue is resolved in this release.

PR 3437903: The intrusion detection and prevention system (IDS/IPS) might report false positives for dropped packets in vSphere reporting tools

In environments with active IDS/IPS, ESXi might report false positives for dropped packets to network monitoring solutions such as Aria Operations. If you deactivate IDS/IPS, the count of dropped packets falls. The issue affects all virtual machines connected to NSX segments.

This issue is resolved in this release.

PR 3438291: ESXi hosts on a vSAN cluster might fail with purple diagnostic screen due to a rare timing issue

A rare timing issue in vSAN might cause ESXi host on vSAN cluster to fail with purple diagnostic screen. In the vmkernel logs, you see an error such as `@BlueScreen: #PF Exception 14 in world 2099267:VSAN_0x432a4`. In the backtrace, you see errors such as:

```
0x453af7b9bcc0:[0x42000d8e1fdf]PREF_InsertEx@com.vmware.vsanutil#0.0.0.1+0xab stack: 0x9,
0x45bc0767b7c0, 0x1a0c1841bdb50, 0x432a49db7140, 0x45bc2bd84280
```

```
0x453af7b9bd10:
[0x42000d8dcb3c]VSANServer_SetOperationTimerTC@com.vmware.vsanutil#0.0.0.1+0x69 stack:
0x45bc0767b7c0, 0x67143756669a719a, 0x45bc0767b7c0, 0xee7135183b47557, 0x45bc2bd84280
```

This issue is resolved in this release.

PR 3470500: After update to ESXi 8.0 Update 3b, you see an alarm requiring a hardware reset due to AMD erratum 1474 on all updated hosts

ESXi 8.0 Update 3b added a fix for vSphere systems becoming unresponsive in about 1044 days after the last system hardware reset due to AMD erratum 1474, related to the core-C6 (CC6) sleep state on EPYC processors of the series 7002, 7Fx2, 7Hx2, and 7001. The fix was to automatically deactivate CC6 in around 1000 days until the next hardware reset, and send alarms to notify about the required reset. However, due to an integer overflow, the timer callback triggers in about 6 days and generates the alert, deactivating CC6 and suggesting that you do a hardware reset.

This issue is resolved in this release. The fix makes sure that the deactivation of CC6 and triggering the alarm occurs in the scheduled 1000 days.

PR 3410296: You see gaps in the collection of metrics for some virtual machines

Due to a time sync issue, collection of data for some VMs might fail and you do not see any performance data in platforms such as Aria Operations. In the `/var/run/log/vobd.log` on the ESXi host, you see errors such as: `1149:2024-06-11T05:01:11.250Z: [ClockCorrelator] 7826272852729us: [vob.clock.correction.adjtime.sync] system clock synchronized to upstream time servers`

This issue is resolved in this release. The fix caches results from Stats Registry queries with selected dynamic counters to make sure metrics are consistent.

PR 3434172: You cannot apply a reverse path filter for virtual machines in promiscuous mode even /Net/ReversePathFwdCheckPromisc is active

In very specific conditions, after an upgrade of a distributed virtual switch or uninstalling an NSX transport node on an ESXi host with converged distributed virtual switch that contains a port group with active promiscuous mode, you might see some networking issues, including:

1. You cannot apply a reverse path filter for virtual machines in promiscuous mode even `/Net/ReversePathFwdCheckPromisc` is active.
2. Intermittent packet connectivity issues when using promiscuous mode on a distributed virtual switch with multiple uplinks.

This issue is resolved in this release.

PR 3442088: vSphere vMotion tasks fail with an error NamespaceMgr could not lock the db file

When virtual machines are configured with namespaces, an issue with the namespace database might cause migration of such VMs with vSphere vMotion to fail. In the vSphere Client, you see errors such as:

Failed to receive migration. The source detected that the destination failed to resume.

An error occurred restoring the virtual machine state during migration. NamespaceMgr could not lock the db file.

The destination ESXi host `vmware.log` shows the following messages:

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.checkpoint.migration.failedReceive] Failed to receive migration.
```

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.namespaceMgr.noLock] NamespaceMgr could not lock the db file.
```

The issue is more likely to occur on VMs with hardware versions earlier than 19 and impacts only VMs provisioned on shared datastores other than NFS, such as VMFS, vSAN, and vSphere Virtual Volumes.

This issue was reported as known in KB [369767](#) and is resolved in this release.

PR 3390022: Capture of network packets by using the PacketCapture tool on ESXi does not work

Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in [Collecting network packets using the lightweight PacketCapture on ESXi](#).

This issue is resolved in this release.

esx-update_8.0.3-0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesx_8.0.3-0.70.24674464 - VMware_bootbank_esx-update_8.0.3-0.70.24674464
PRs Fixed	3474996
CVE numbers	N/A

Updates the `loadesx` and `esx-update` VIBs.

ESXi 8.0 Update 3e adds support for vSphere Quick Boot to:

- Intel vRAN Baseband Driver
- Intel Platform Monitoring Technology Driver
- Intel Data Center Graphics Driver
- AMD Instinct MI Series Driver

esxio-update_8.0.3-0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	• VMware_bootbank_loadesxio_8.0.3-0.70.24674464 • VMware_bootbank_esxio-update_8.0.3-0.70.24674464
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

Broadcom-ntg3_4.1.15.0-4vmw.803.0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	• VMW_bootbank_ntg3_4.1.15.0-4vmw.803.0.70.24674464
PRs Fixed	3498633
CVE numbers	N/A

Updates the `ntg3` VIB.

Broadcom-ELX-lpfc_14.4.0.40-35vmw.803.0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	• VMW_bootbank_lpfc_14.4.0.40-35vmw.803.0.70.24674464
PRs Fixed	3451575

CVE numbers	N/A
-------------	-----

Updates the `lpfc` VIB.

PR 3451575: You don't see all paths to an IBM SAN Volume Controller (SVC) storage target

An issue in the `lpfc` driver might prevent the Emulex Fibre Channel Host Bus Adapter (FC HBA) from connecting to some IBM SVC storage targets. In a multipathing setup, you might see that not all storage paths are established.

This issue is resolved in this release.

Intel-Volume-Mgmt-Device_3.0.0.1010-12vmw.803.0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMW_bootbank_iavmd_3.0.0.1010-12vmw.803.0.70.24674464
PRs Fixed	3462191
CVE numbers	N/A

Updates the `iavmd` VIB.

PR 3462191: You do not see all drivers plugged in behind Intel Volume Management Device (VMD) controllers

When you have multiple drives plugged in behind VMD controllers, not all drives might be enumerating due to insufficient NRM heap memory. In the `vmkernel` log, you see errors such as:

```
2024-04-10T21:05:26Z In(182) vmkernel: NRM: memory pool with TAG 16 creation failed
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library is going to be stopped!
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library stopped
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library stopped
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library deinitialized started
```

This issue is resolved in this release.

Mellanox-nmlx5_4.23.6.5-1vmw.803.0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMW_bootbank_nmlx5-cc_4.23.6.5-1vmw.803.0.70.24674464 - VMW_bootbank_nmlx5-cc-esxio_4.23.6.5-1vmw.803.0.70.24674464 - VMW_bootbank_nmlx5-core_4.23.6.5-1vmw.803.0.70.24674464

	• VMW_bootbank_nmlx5-rdma_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-core-esxio_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-rdma-esxio_4.23.6.5-1vmw.803.0.70.24674464
PRs Fixed	3442016, 3478908, 3506135, 3483165, 3453042, 3453046, 3474329
CVE numbers	N/A

Updates the nmlx5 VIB.

PR 3478908: On ConnectX-4 and later physical network adapters, you see slower TX transmission when packets in the same queue have different values for the Priority Code Point (PCP)

Packets with different VLAN (CoS PCP/801.2p) priorities might be transmitted slower on ConnectX-4 and later physical network adapters.

This issue is resolved in this release. The fix adds the `trust_dscp_same_priority` nmlx5 driver module parameter to map all DSCP priorities to the same traffic class to avoid slowness in the NIC hardware.

PR 3506135: You see messages from the nmlx5 driver such as "Health: Miss counters detected" with "synd 0x0" and "IO was aborted"

The nmlx5 driver health logic might incorrectly identify the state of a NIC as faulty and fail firmware commands from the driver. In the vmkernel logs, you see similar messages at random times:

```
<NMLX_ERR> nmlx5_core: 0000:45:00.0: Health: Miss counters detected.
<NMLX_INF> synd 0x0: unrecognized error
<NMLX_INF> extSynd 0x0000
<NMLX_ERR> nmlx5_QueryNicVportContext:188 command failed: IO was aborted
<NMLX_ERR> nmlx5_QueryVportCounter:1851 command failed: IO was aborted
```

This issue is resolved in this release.

PR 3442016: You might see low throughput with North-South encapsulated traffic with NSX Edge VMs

In releases earlier than ESXi 8.0 Update 3e, the nmlx5 driver does not support decapsulation of GENEVE encapsulated packets with BF3 NICs. As a result, you might see low throughput with North-South encapsulated traffic with NSX Edge VMs.

This issue is resolved in this release. The fix adds support for decapsulation of GENEVE encapsulated packets with BF3 NICs for the nmlx5 driver.

PR 3483165: The Data Center Bridging group (DCB) state of nmlx5 driver NICs in ESXi hosts is always active in Hardware DCB mode

The DCB state of ESXi hosts is always active in VMNICs for the nmlx5 driver, irrespective of the DCB state of the physical switch.

This issue is resolved in this release. With the fix, the switch DCB state determines the DCB state of NICs in ESXi hosts when NICs are in Hardware DCB mode.

PR 3453042: An ESXi host fails with a purple diagnostic screen due to uplink watchdog resets on nmlx5 driver NICs

In rare cases, when ESXi on data processing units (DPU) fails with a kernel panic, the ESXi host also fails with a purple diagnostic screen. The issue occurs on ESXi hosts with virtual machines that have virtual NICs on a distributed switch backed by DPUs, when an uplink watchdog resets nmlx5 driver NICs.

This issue is resolved in this release.

PR 3453046: You might see low throughput with IPv6 UDP traffic in Enhanced Datapath mode for the nmlx5 driver

RSS logic in the nmlx5 driver incorrectly calculates hash for the IPv6 traffic which results in utilizing just a single RX queue.

This issue is resolved in this release.

PR 3474329: After deploying virtual machines, you see numerous log messages from the nmlx5 driver in the vmkernel.log

After the nmlx5 driver adds or removes the MAC address filter of a virtual machine or a VMkernel adapter, in the vmkernel.log, you might see multiple messages such as `nmlx5_en_L2TableIndexAdd` and `nmlx5_en_L2TableIndexRemove` after deploying virtual machines. These are informational logs from the nmlx5 driver indicating that filters are being applied or removed.

This issue is resolved in this release. With the fix, such messages are logged only in debug log level.

VMware-NVMeoF-TCP_1.0.1.32-1vmw.803.0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMW_bootbank_nvmetcp_1.0.1.32-1vmw.803.0.70.24674464 - VMW_bootbank_nvmetcp-esxio_1.0.1.32-1vmw.803.0.70.24674464
PRs Fixed	3444303, 3454941, 3469074, 3469074, 3473257
CVE numbers	N/A

Updates the `nvmetcp` VIB.

PR 3473257: NVMe/TCP controllers on an ESXi host might not recover connection when a target upgrades or restarts

When performing NVMe/TCP target upgrade or restart, controllers in the ESXi host might not recover their connection even after the target is online. As a result, you do not see the respective controllers from the vSphere Client. A rare race condition between the NVMe core layer and NVMe/TCP drivers causes the issue.

This issue is resolved in this release.

PR 3469074: ESXi NVMe/TCP controllers do not recover connectivity when a target upgrades or restarts

When performing NVMe/TCP target upgrade or restart, controllers in an ESXi host might not recover connectivity after the target is online. As a result, you do not see the respective controllers from the vSphere Client. A rare race condition between the NVMe core layer and NVMe/TCP drivers causes the issue.

This issue is resolved in this release.

PR 3454941: You cannot create datastores on namespaces backed by some kind of NVMe/TCP storage arrays in vSphere 8.0 Update 3b environments

An issue that affects only vSphere 8.0 Update 3b environments might cause tasks to create a datastore on namespaces backed by some kind of NVMe/TCP storage arrays to fail. Frequent messages in the vmkernel logs such as `nvmectp:nt_ReceiveR2TPdu:2667 [ctrlr 259, queue 3] txPdu 0x433685d6eb40 (cid 169) expects PDUs in types 0x28 but received PDU in type 9: Failure` are indicative for the issue.

This issue is resolved in this release.

PR 3444303: A rare race condition when connecting to a target NVMe over TCP controller might cause ESXi failure with a purple diagnostic screen

When connecting to a target NVMe over TCP controller, either for the first time or during a controller reset, the target might fail or not respond to the connect command for some reason. As a result, a race condition occurs in the ESXi host between the NVMe core layer, and the NVMe/TCP driver and the host fails with a purple diagnostic screen.

This issue is resolved in this release. This fix prevents the race condition.

VMware-vmkusb_0.1-23vmw.803.0.70.24674464

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMW_bootbank_vmkusb-esxio_0.1-23vmw.803.0.70.24674464 - VMW_bootbank_vmkusb_0.1-23vmw.803.0.70.24674464
PRs Fixed	3458162
CVE numbers	N/A

Updates the `vmkusb` VIB.

Support for Communication Device Class Network Control Model (CDC-NCM) in the ESXi USB driver:

Starting with ESX 8.0 Update 3e, the ESXi USB driver supports the CDC-NCM protocol for compatibility with HPE Gen12 iLO Virtual NIC and interoperability with HPE Agentless Management (AMS), Integrated Smart Update Tools (iSUT), the iLORest config tool, Intelligent Provisioning, and DPUs.

ESXi_8.0.3-0.65.24659227

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_drivervm-gpu-base_8.0.3-0.65.24659227 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.65.24659227

	- VMware_bootbank_vsan_8.0.3-0.65.24659227 - VMware_bootbank_esxio-combiner_8.0.3-0.65.24659227 - VMware_bootbank_bmc-al-esxio_8.0.3-0.65.24659227 - VMware_bootbank_trx_8.0.3-0.65.24659227 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.65.24659227 - VMware_bootbank_crx_8.0.3-0.65.24659227 - VMware_bootbank_cpu-microcode_8.0.3-0.65.24659227 - VMware_bootbank_clusterstore_8.0.3-0.65.24659227 - VMware_bootbank_esxio_8.0.3-0.65.24659227 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.65.24659227 - VMware_bootbank_vcls-pod-crx_8.0.3-0.65.24659227 - VMware_bootbank_gc_8.0.3-0.65.24659227 - VMware_bootbank_vds-vsip_8.0.3-0.65.24659227 - VMware_bootbank_gc-esxio_8.0.3-0.65.24659227 - VMware_bootbank_infravisor_8.0.3-0.65.24659227 - VMware_bootbank_vsanhealth_8.0.3-0.65.24659227 - VMware_bootbank_esx-xserver_8.0.3-0.65.24659227 - VMware_bootbank_native-misc-drivers_8.0.3-0.65.24659227 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.65.24659227 - VMware_bootbank_bmc-al_8.0.3-0.65.24659227 - VMware_bootbank_vdfs_8.0.3-0.65.24659227 - VMware_bootbank_esxio-base_8.0.3-0.65.24659227 - VMware_bootbank_esx-base_8.0.3-0.65.24659227
PRs Fixed	3404662, 3432571, 3432690, 3435458, 3458134, 3444037, 3444037, 3444626, 3447888, 3449523, 3452700, 3454671, 3454671, 3454671, 3457664, 3470207
CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching. This patch updates the `drivervm-gpu-base`, `esxio-dvfilter-generic-fastpath`, `vsan`, `esxio-combiner`, `bmc-al-esxio`, `trx`, `esx-dvfilter-generic-fastpath`, `crx`, `cpu-microcode`, `clusterstore`, `esxio`, `native-misc-drivers-esxio`, `vcls-pod-crx`, `gc`, `vds-vsip`, `gc-esxio`, `infravisor`, `vsanhealth`, `esx-xserver`, `native-misc-drivers`, `esxio-combiner-esxio`, `bmc-al`, `vdfs`, `esxio-base`, and `esx-base` VIBs.

This patch resolves the following issues:

ESXi 8.0 Update 3e includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor

Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C

Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series

Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003801	8/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003801	8/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/ b)	0xbf	baseline	0x7002a01	8/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series

Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd0003f5	8/2/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002c0	8/8/2024	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xba	8/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x3a	8/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x54	8/1/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	baseline	0x2c0003e0	7/30/2024	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b000620	7/30/2024	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xfa	7/28/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x102	7/31/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x63	8/1/2024	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x12c	9/25/2024	Intel Xeon E-2400 Series

Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x21000291	7/30/2024	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series
-------------------	-------------------	------	----------	------------	-----------	--

esx-update_8.0.3-0.65.24659227

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesx_8.0.3-0.65.24659227 - VMware_bootbank_esx-update_8.0.3-0.65.24659227
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesx` and `esx-update` VIBs.

esxio-update_8.0.3-0.65.24659227

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesxio_8.0.3-0.65.24659227 - VMware_bootbank_esxio-update_8.0.3-0.65.24659227
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

VMware-VM-Tools_12.5.1.24649672-24659227

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	No
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A

Affected Software	N/A
Affected VIBs Included	VMware_locker_tools-light_12.5.1.24649672-24659227
PRs Fixed	3442312, 3494209
CVE numbers	N/A

Updates the `locker_tools-light` VIB.

VMware Tools Bundling Changes in ESXi 8.0 Update 3e

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 3e:

- windows.iso:** VMware Tools 12.5.1 supports 64-bit Windows versions - Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- windows-x86.iso:** VMware Tools 12.4.6 supports 32-bit Windows versions.
- linux.iso:** VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - windows.iso:** for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - winPreVista.iso:** for Windows 2000, Windows XP, and Windows 2003.
 - linuxPreGLibc25.iso:** supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- solaris.iso:** VMware Tools image 10.3.10 for Solaris.
- darwin.iso:** Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.5.1 Release Notes](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Upgrading VMware Tools](#)

ESXi-8.0U3e-24674464-standard

Profile Name	ESXi-8.0U3e-24674464-standard
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	April 10, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.70.24674464 - VMware_bootbank_vsan_8.0.3-0.70.24674464 - VMware_bootbank_infravisor_8.0.3-0.70.24674464 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.70.24674464 - VMware_bootbank_crx_8.0.3-0.70.24674464 - VMware_bootbank_bmcal_8.0.3-0.70.24674464 - VMware_bootbank_vsanhealth_8.0.3-0.70.24674464

	• VMware_bootbank_drivervm-gpu-base_8.0.3-0.70.24674464 • VMware_bootbank_gc_8.0.3-0.70.24674464 • VMware_bootbank_vdfs_8.0.3-0.70.24674464 • VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.70.24674464 • VMware_bootbank_esx-base_8.0.3-0.70.24674464 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.70.24674464 • VMware_bootbank_esx-xserver_8.0.3-0.70.24674464 • VMware_bootbank_vds-vsip_8.0.3-0.70.24674464 • VMware_bootbank_cpu-microcode_8.0.3-0.70.24674464 • VMware_bootbank_esxio_8.0.3-0.70.24674464 • VMware_bootbank_clusterstore_8.0.3-0.70.24674464 • VMware_bootbank_esxio-base_8.0.3-0.70.24674464 • VMware_bootbank_trx_8.0.3-0.70.24674464 • VMware_bootbank_gc-esxio_8.0.3-0.70.24674464 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.70.24674464 • VMware_bootbank_bmc-esxio_8.0.3-0.70.24674464 • VMware_bootbank_vcls-pod-crx_8.0.3-0.70.24674464 • VMware_bootbank_native-misc-drivers_8.0.3-0.70.24674464 • VMware_bootbank_loadesx_8.0.3-0.70.24674464 • VMware_bootbank_esx-update_8.0.3-0.70.24674464 • VMware_bootbank_loadesxio_8.0.3-0.70.24674464 • VMware_bootbank_esxio-update_8.0.3-0.70.24674464 • VMW_bootbank_ntg3_4.1.15.0-4vmw.803.0.70.24674464 • VMW_bootbank_lpf_14.4.0.40-35vmw.803.0.70.24674464 • VMW_bootbank_iavmd_3.0.0.1010-12vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-cc_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-cc-esxio_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-core_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-rdma_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-core-esxio_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-rdma-esxio_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nvmetcp_1.0.1.32-1vmw.803.0.70.24674464 • VMW_bootbank_nvmetcp-esxio_1.0.1.32-1vmw.803.0.70.24674464 • VMW_bootbank_vmkusb-esxio_0.1-23vmw.803.0.70.24674464 • VMW_bootbank_vmkusb_0.1-23vmw.803.0.70.24674464 • VMware_locker_tools-light_12.5.1.24649672-24659227
PRs Fixed	3498633, 3341273, 3410296, 3414552, 3415171, 3425316, 3427281, 3480453, 3431433, 3432906, 3433295, 3433436, 3433610, 3434171, 3434270, 3437835, 3437903, 3438291, 3439279, 3444039, 3444137, 3444408, 3445065, 3445276, 3446362, 3447505, 3447686, 3447687, 3447888, 3448872, 3449523, 3450290, 3450374, 3450501, 3451651, 3451669,

	3451849, 3452081, 3452460, 3452487, 3452906, 3453391, 3453981, 3454108, 3454123, 3454926, 3454942, 3454983, 3455021, 3455315, 3456018, 3457039, 3457042, 3457050, 3457350, 3457744, 3457900, 3457959, 3458521, 3459100, 3459378, 3459526, 3459825, 3460312, 3460819, 3463099, 3463147, 3463364, 3465047, 3465070, 3465259, 3465275, 3465373, 3465624, 3465965, 3466011, 3466176, 3466283, 3466538, 3467137, 3467521, 3467708, 3468136, 3468139, 3468215, 3469014, 3469036, 3469145, 3469467, 3469652, 3470385, 3470427, 3470500, 3470714, 3470732, 3472623, 3473154, 3473626, 3473692, 3473767, 3473768, 3473860, 3474427, 3474443, 3474967, 3475015, 3475152, 3475393, 3476179, 3476630, 3476631, 3476632, 3476634, 3476963, 3477021, 3477062, 3477385, 3477409, 3477725, 3478001, 3478002, 3478236, 3478276, 3478486, 3479338, 3479340, 3479464, 3480154, 3480456, 3480968, 3481216, 3481241, 3481310, 3481386, 3481475, 3481507, 3482259, 3482802, 3482963, 3483589, 3483610, 3483714, 3484669, 3484787, 3484836, 3485671, 3485971, 3486600, 3486647, 3486747, 3486955, 3487351, 3487380, 3488415, 3488480, 3494691, 3497983, 3462191, 3442016, 3478908, 3444303, 3454941, 3469074, 3473257, 3458162, 3474996
Related CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

This patch updates the following issues:

PR 3478908: On ConnectX-4 and later physical network adapters, you see slower TX transmission when packets in the same queue have different values for the Priority Code Point (PCP)

Packets with different VLAN (CoS PCP/801.2p) priorities might be transmitted slower on ConnectX-4 and later physical network adapters.

This issue is resolved in this release. The fix adds the `trust_dscp_same_priority` nmlx5 driver module parameter to map all DSCP priorities to the same traffic class to avoid slowness in the NIC hardware.

PR 3442191: Update of a vSAN environment with Original Storage Architecture (OSA) to version 8.0 Update 3 might be slower compared to previous updates

The first reboot of the vSphere system in a vSAN OSA environment following upgrade or update to 8.0 Update 3 might cause a delay in the update process, which depends on the workload running at the time of update. Nodes with heavy workloads running before the update are likely to cause higher delays. This is expected and for updates to 8.0 Update 3, environments with large install base must consider the possible delay in the update process.

This issue is resolved in this release.

PR 3506135: You see messages from the nmlx5 driver such as "Health: Miss counters detected" with "synd 0x0" and "IO was aborted"

The nmlx5 driver health logic might incorrectly identify the state of a NIC as faulty and fail firmware commands from the driver. In the vmkernel logs, you see similar messages at random times:

```
<NMLX_ERR> nmlx5_core: 0000:45:00.0: Health: Miss counters detected.
<NMLX_INF> synd 0x0: unrecognized error
<NMLX_INF> extSynd 0x0000
<NMLX_ERR> nmlx5_QueryNicVportContext:188 command failed: IO was aborted
<NMLX_ERR> nmlx5_QueryVportCounter:1851 command failed: IO was aborted
```

This issue is resolved in this release.

ESXi 8.0 Update 3e adds support for vSphere Quick Boot to:

- Intel vRAN Baseband Driver
- Intel Platform Monitoring Technology Driver
- Intel Data Center Graphics Driver
- AMD Instinct MI Series Driver

Support for Communication Device Class Network Control Model (CDC-NCM) in the ESXi USB driver:

Starting with ESX 8.0 Update 3e, the ESXi USB driver supports the CDC-NCM protocol for compatibility with HPE Gen12 iLO Virtual NIC and interoperability with HPE Agentless Management (AMS), Integrated Smart Update Tools (iSUT), the iLORest config tool, Intelligent Provisioning, and DPUs.

PR 3479338: Delayed cleanup of a deleted component can lead to high space usage reports for vSAN datastores

The vSAN I/O scheduler manages in-flight I/Os for performance and fairness. Under certain conditions, the scheduler might get an incorrect counter for the in-flight resync I/Os, which blocks new resync I/Os. As a result, the cleanup of blocked resync I/Os is delayed and might lead to high space usage.

This issue is resolved in this release.

PR 3442016: You might see low throughput with North-South encapsulated traffic with NSX Edge VMs

In releases earlier than ESXi 8.0 Update 3e, the nmlx5 driver does not support decapsulation of GENEVE encapsulated packets with BF3 NICs. As a result, you might see low throughput with North-South encapsulated traffic with NSX Edge VMs.

This issue is resolved in this release. The fix adds support for decapsulation of GENEVE encapsulated packets with BF3 NICs for the nmlx5 driver.

PR 3476179: A memory leak might prevent the Object Storage File System Daemon (osfsd) from processing requests for vSphere Virtual Volumes

A rare memory leak condition in the osfsd service after prolonged uptime might prevent vSphere Virtual Volumes from provisioning virtual machines, including creation and cloning operations.

This issue is resolved in this release.

PR 3459100: Virtual machine migration in a vSphere Virtual Volumes datastore might cause an ESXi host to fail with a purple diagnostic screen

During virtual machine migration from a vSphere Virtual Volumes datastore, a rare race condition in Device Block allocation tracking might lead to an ESXi host failure, resulting in a purple diagnostic screen.

This issue is resolved in this release. The fix prevents the race condition.

PR 3484787: If you reconfigure the scratch location and reboot an ESXi host, all logging stops and the host becomes unresponsive

If you have audit recording to local storage enabled for an ESXi host on a directory configured at the default location `/scratch/auditLog` or at a subdirectory of `/scratch`, and then reconfigure `/scratch`, the syslog daemon fails after a host reboot, because it cannot find the record directory and files in the new location. As a result, all logging on such hosts stops and eventually the host becomes unresponsive.

This issue is resolved in this release. The fix handles such scenarios and deactivates audit records to local storage for security reasons, generating a VMkernel Observation (VOB) message such as `vSphere error: (esx.problem.vmsyslogd.auditrecord.local.disabled)` but preventing a failure of the vmsyslog daemon. You can re-enable audit recording in the new `/scratch` location after the host reboots. Old audit record files are available in the previously configured scratch path, if it is still accessible.

PR 3455315: The PowerSvc plug-in of the sensord service stops polling data by using Intelligent Platform Management Interface (IPMI) without a retry

Due to a change related to the Running Average Power Limit (RAPL) host power reporting in ESXi 8.0 Update 3, the sensor service might deactivate power monitoring by using IPMI if only one IPMI power data poll fails. The PowerSvc plug-in does not retry the poll and IPMI polling stops. PowerSvc goes offline and logs failing to read from the power provider in the `/var/log/sensord.log`.

This issue is resolved in this release. The fix makes sure the PowerSvc plug-in retries IPMI power data polls in case of failure and IPMI data collection continues.

PR 3460819: Virtual machines with UnixWare OS and a vlane virtual device might repeatedly reset when running on an ESXi 8.x host

An issue with vlane virtual devices might cause virtual machines with UnixWare OS on ESXi 8.x hosts to repeatedly fail due to a reboot of the guest OS with the error `Your guest operating system has accessed an I/O range where multiple devices overlap`.

This issue is resolved in this release.

PR 3483165: The Data Center Bridging group (DCB) state of nmlx5 driver NICs in ESXi hosts is always active in Hardware DCB mode

The DCB state of ESXi hosts is always active in VMNICs for the nmlx5 driver, irrespective of the DCB state of the physical switch.

This issue is resolved in this release. With the fix, the switch DCB state determines the DCB state of NICs in ESXi hosts when NICs are in Hardware DCB mode.

PR 3479340: You see checksum errors for some objects in a stretched cluster after a network outage

After a network outage or some other critical event on a stretched cluster, the vSAN Distributed Object Manager (DOM) might fail to get the correct checksums for objects that do not have sub-fault tolerance. As a result, you might see in tools such as VMware Skyline Health Diagnostics checksum errors for many objects in the stretched cluster.

This issue is resolved in this release.

PR 3461166: In the vSphere Client, you cannot change a VMNIC on the virtual switch of an ESXi host

The cache of the hostd service might not be up-to-date after you add a VMNIC to the virtual switch on an ESXi host due to a bug in the teaming policy of port groups. As a result, when you remove another VMNIC from the switch, both the newly added and the removed VMNICs are removed from the teaming policy of the port groups. When you work in the vSphere Client, the task seems to revert.

This issue is resolved in this release.

PR 3463099: Guest OS might mark file system read-only after repeated IO failure which can cause Kubernetes control plane nodes to go into Read-only mode

An issue with the check for max allowed memory address in memory mapping code during I/O processing at virtual SCSI might cause a guest OS on an ESXi host to mark guest file system as read-only. As a result, Kubernetes control plane nodes might go into Read-only mode. In the vmkernel logs, you see multiple error messages such as `failed to pin: I/O error`.

This issue is resolved in this release.

PR 3461321: Custom TLS settings of the Fault Domain Manager (FDM) might not work on ESXi hosts of earlier versions that you manage with a vCenter of version 8.0 Update 3

Custom TLS settings of the vmware-fdm service at port 8182 might not work on hosts of earlier versions in a cluster with active High Availability that you manage with a vCenter of version 8.0 Update 3.

This issue is resolved in this release. For more information on the workflows and the use of such custom settings, see [KB 320798](#).

PR 3488415: ESXi host fails with a purple diagnostic screen with messages such as #PF Exception 14 in world XXXXXXXX:vmknvmeGener

In the discovery phase of NVMe over Fabrics (NVMe-oF), some targets might report to an ESXi host some subsystem NVMe Qualified Names (NQN) to which that ESXi host cannot connect due to missing permissions or unavailable route. The ESXi host continues attempting to connect to such NQNs and if a network issue occurs, in rare cases, the host might fail with a purple diagnostic screen with messages such as #PF Exception 14 in world XXXXXXXX:vmknvmeGener.

This issue is resolved in this release.

PR 3486647: Copying files with the cp command creates files with unexpected access timestamp (atime) and modification timestamp (mtime) values

If vSAN File Services NFS 4.1 is mounted to a RHEL 8 or later NFS client, when copying a file by using the `cp` command, you might see `atime` or `mtime` such as 1970 instead of the current date.

This issue is resolved in this release.

PR 3453042: An ESXi host fails with a purple diagnostic screen due to uplink watchdog resets on nmlx5 driver NICs

In rare cases, when ESXi on data processing units (DPU) fails with a kernel panic, the ESXi host also fails with a purple diagnostic screen. The issue occurs on ESXi hosts with virtual machines that have virtual NICs on a distributed switch backed by DPUs, when an uplink watchdog resets nmlx5 driver NICs.

This issue is resolved in this release.

PR 3486025: Virtual machines and interactive shells might become unresponsive due to an issue with the initialization of user worlds in ESXi

By default, the initialization of user worlds in ESXi adds an active group that requires the allocation of some CPU resources. If the initialization fails for some reason before the CPU resource is available, the creation of new user worlds starts failing. As a result, virtual machines and interactive shells might become unresponsive.

This issue is resolved in this release.

PR 3469652: Logging stops on all log files and configured remote host targets on ESXi hosts

A rare deadlock situation in the ESXi log daemon `vm syslogd` might cause logging to stop on all log files and configured remote host targets on ESXi hosts. The issue might occur as the result of multiple syslog reloads.

This issue is resolved in this release.

PR 3466011: You see the datastore name in storage VMkernel Observation (VOB) reports as Unknown

Due to a missing function to grant access required to fetch datastore names, you might see the datastore name in `esx.problem.storage.vob` messages as unknown.

For example: 2024-08-27T04:31:29.150Z In(14) vobd[2097668]: [scsiCorrelator] 5815563033719us: [esx.problem.storage.redundancy.lost] Lost path redundancy to storage device naa.60060e80072b600000302b6000000026. Path vmhba1:C0:T1:L0 is down. Affected datastores: Unknown.

This issue is resolved in this release.

PR 3482259: NVMe over TCP controllers might stay offline and not recover even when the network and the target port have no issues

Due to issues with the network or the target, NVMe over TCP controllers in an ESXi host might enter an offline state. Usually, controllers recover and turn online as soon as the issue is resolved, but in some very rare scenarios, the controllers might stay offline and do not recover, which makes the NVMe/TCP namespaces inaccessible. The issue occurs due to a dead lock when many controllers are being reset and at the same time many I/O requests time out, which

triggers multiple taskmgmt requests that compete for critical resource. The only way to exit this state is to reboot the ESXi host.

This issue is resolved in this release. The fix avoids such dead locks.

PR 3481216: "Route based on physical NIC load" policy does not work as expected

An issue with the permissions for the net-blt daemon might cause load balancing based on the **Routing based on PNIC load** teaming policy to not work as expected.

This issue is resolved in this release.

PR 3486955: Storage path unclaiming requests might be blocked due to a Fabric Performance Impact Notification (FPIN) issue

FPIN messages serve as an early warning system to relieve negative effects on the performance of Fibre Channel devices, such as optimizing I/O behavior and avoiding impaired paths by notifying devices on the current fabric conditions. In some cases, when memory allocation fails, FPIN might keep a reference on a storage path. As a result, path scanning and claiming routine returns `BUSY`, as the reference no longer exists, which blocks path claiming and unclaiming, and FPIN-related information cannot go through to ESXi hosts. In the ESXCLI command output you do not see FPIN messages. When the issue occurs, you see a message in the VMkernel logs such as: `WARNING: StorageFPIN: 521: Failed to allocate memory.`

Another symptom is when after a SAN maintenance or in case of a storage path outage FPIN holds a reference count on the storage paths and the FC driver cannot allocate new paths or reestablish existing ones. In `/var/log/vmkernel.log`, you see the following sequence repeating:

```
WARNING: nfnic: <2>: fnic_handle_report_lun: 1467: lun add failure! in_remove: 0
ioAllowed: 1
```

```
WARNING: nfnic: <2>: fnic_tport_event_handler: 2130: lunmap update failed,retry ...
```

This issue is resolved in this release.

PR 3341273: Migration of virtual machines fails with an error "Failed to allocate migration heap" on the source ESXi host

In very rare cases, if a virtual machine powers off before a running migration completes, ESXi might fail to clean up the migration heap. If the memory leak continues over time, the migration heap might be exhausted, and virtual machines cannot migrate from such ESXi hosts.

This issue is resolved in this release.

PR 3467137: A race condition during heap cleanup might cause an ESXi host to fail with a purple diagnostic screen

A rare race condition when reading the VMX file descriptor table occurs at the time VMX shuts down might cause an ESXi host to fail with a purple diagnostic screen. The error message is similar to `PANIC bora/vmkernel/main/dlmalloc.c:4944 - Usage error in dlmalloc.`

This issue is resolved in this release.

PR 3483610: If you insert an invalid or blank DVD on an ESXi host and configure it as a host device on virtual machines, the VM performance degrades

If a blank or incompatible disc is in the DVD drive of an ESXi host, the ESXi HPP module checks for the SCSI sense key and additional sense keys to determine next steps. Due to a missing check to handle Active State Change (ASC) codes 0x30 and 0x3a, the check fails, causing a failover that keeps repeating. As a result, you see a degradation of performance on the virtual machines on such hosts.

This issue is resolved in this release. The fix handles the ASC.

PR 3475015: Duplicate traffic on multiple uplinks might cause traffic outage

During an upgrade of the physical switch infrastructure, if you use the `TeamPolicyUpDelay` setting and all uplinks on the distributed switch flap, a fallback mode during the time of delay might cause duplicate traffic on all uplinks of the switch, which leads to traffic outage.

This issue is resolved in this release. The fix adds the advanced config parameter `/Net/TeamingIgnoreShotgun` to prevent the fallback mode in such cases and avoid duplicate traffic.

PR 3451306: The usage scale in the CPU chart under Monitor > Performance Overview shows in thousands, and you cannot notice CPU usage spikes to 100%

The usage scale in the CPU chart under **Monitor > Performance Overview** for virtual machines might show in thousands, reaching up to 2400%, instead of hundreds to allow you to notice CPU usage spikes to 100%. The expanded scale also prevents realistic estimation of other key performance metrics on the graph, such as CPU Ready, as they are displayed at the bottom of the chart.

This issue is resolved in this release.

PR 3470732: Virtual machines using Raw Device Mapping (RDM) might become unresponsive due storage heap exhaustion

An issue in tracking SCSI Inquiry or Mode Sense requests that virtual machines issue to RDM LUNs might lead to heap exhaustion at the ESXi storage layer. As a result, such VMs become unresponsive. In the vSphere Client, you see series of **Lost access to volume** and **Successfully restored access** messages until you power off the VMs.

This issue is resolved in this release.

PR 3474443: A memory leak in the swapObj daemon (swapobjd) during object replication might cause a failure of the service due to stale metadata

An issue with stale metadata for swap objects might cause a memory leak in the swapobjd during object replication, leading to gradual memory exhaustion and ultimately to a failure of the service.

This issue is resolved in this release. The fix ensures stale swap object metadata is removed.

PR 3479464: ESXi hosts might fail with purple diagnostic screen while processing SCSI commands from the guest OS

An issue with the validation of SCSI commands from the guest OS at the virtual SCSI layer might cause an ESXi host to fail with a purple diagnostic screen.

This issue is resolved in this release. The fix adds enhanced command validation for SCSI commands at the virtual SCSI layer.

PR 3481507: Discovery of iSCSI dynamic targets, LUNs, and datastores fails on ESXi host reboot

In some cases, dynamic discovery of targets in systems that use dynamic discovery addresses for iSCSI configuration might fail. As a result, you do not see iSCSI targets, LUNs, and datastores after a reboot of ESXi hosts. The dynamic discovery of iSCSI targets might also fail when the Challenge-Handshake Authentication Protocol (CHAP) is active on the ESXi host.

This issue is resolved in this release.

PR 3454942: In the vSphere Client, you cannot reconfigure a vGPU virtual machine if more than 128 vCPUs are active

If you add more than 128 vCPUs on a VM with a vGPU, you cannot manage the configuration from the vSphere Client. The issue occurs because input-output memory management unit (IOMMU) support for a vGPU with more than 128 vCPUs is not available.

This issue is resolved in this release. With the fix, you can manage the vGPU VM configuration from the vSphere Client. However, the VM might not be able to power on, since IOMMU support is not available. To make sure the reconfiguration

succeeds, add to the .vmx file the following 2 parameters: `pciPassthru.vmiop.allowViommu="TRUE"` and `pciPassthru.vmiop.enableViommu="TRUE"`.

PR 3477772: Encrypted virtual machines with active Change Block Tracking (CBT) might intermittently power off after a rekey operation

Due to a race condition between the VMX and hostd services for a specific VMcrypt-related reconfiguration, encrypted VMs with active CBT might unexpectedly power off during a rekey operation.

This issue is resolved in this release.

PR 3469014: ESXi host might fail with a purple diagnostic screen if the Multipath Plug-in (MPP) does not implement the event handler for Fabric Performance Impact Notification (FPIN)

Description: ESXi 8.0 Update 3 added APIs to notify MPPs in case of Fibre channel FPIN events. However, if the MPP is compiled against an API SDK of version earlier than 8.0 Update 3, or if the MPP does not implement the new event handler, then the ESXi host fails with a purple diagnostic screen on an FPIN event.

This issue is resolved in this release.

PR 3465259: NVMe Dell PowerStore clustered VMDK datastore fails to recover after an All-Paths-Down (APD) situation

A reservation conflict might prevent a NVMe Dell PowerStore clustered VMDK datastore to take a reservation and cause failure to recover after losing the reservation.

This issue is resolved in this release.

PR 3480154: High VMFS heap memory consumption might cause virtual machines to become unresponsive

In environments with many large datastores and many VMs on thin-provisioned disks, VMFS resource allocation might consume a high amount of heap memory, proportional to the size of the datastores. As a result, VMFS heap memory might be exhausted and lead to VMs becoming unresponsive or vSphere vMotion tasks failing with errors for VM invalid state.

This issue is resolved in this release.

PR 3460312: Virtual machines on NFSv3 datastores might fail during snapshot consolidation if a third-party backup is also active

Snapshot consolidation during backup with third-party software, such as NetBackup, might fail due to extra read-only locks on the VM snapshot file while the snapshot consolidation is in progress and trying to upgrade that same lock. As a result, a lock conflict might cause a VM to fail.

This issue is resolved in this release. The fix prevents snapshot consolidation tasks to upgrade the read-only lock when more than 1 holders for that lock exist.

PR 3457744: Migration of virtual machines to 8.0 Update 3 and later ESXi host might result in TPM null hierarchy issues

A rare issue related to how applications use TPM null hierarchy might lead to failure to authenticate virtual machines after migrating them from an ESXi host of version earlier than 8.0 Update 3 to 8.0 Update 3 and later. You can identify the issue if the command `Get-TpmSupportedFeature` in Windows returns NULL or the command `sudo tpm2_createprimary -C n -c context.out` in Linux fails.

This issue is resolved in this release. If you already face the issue, power cycle the virtual machines.

PR 3470385: Restoring the iSCSI configurations on an ESXi host might fail during boot due to stale entries in the ConfigStore database

If the ConfigStore database on an ESXi host has stale iSCSI configurations related to hardware iSCSI adapters, the iSCSI jumpstart might fail and restoring the iSCSI configuration on that host also fails during boot.

This issue is resolved in this release.

PR 3477062: vSAN hosts might become unresponsive due to high latencies and high CPU usage during peaks of guest unmap traffic

Internal threads in vSAN might cause a CPU busy loop during a peak of unmap traffic. As a result, some hosts might become unresponsive.

This issue is resolved in this release. The fix is relinquishing CPUs when unmap requests start to slow down.

PR 3466538: A race condition between the TCP control path and the keep alive timer might cause an ESXi host to fail with a purple diagnostic screen

A rare race condition between the TCP control path during a disconnect operation and the keep alive timer might cause ESXi hosts of version 8.0 Update 3 to fail with a purple diagnostic screen. The issue occurs under specific network conditions when the timer and disconnect logic overlap.

This issue is resolved in this release. The fix prevents the purple diagnostic screen failure.

PR 3473626: If a command issued to a vSAN disk group complete after 120 seconds, the vSAN host might fail with a purple diagnostic screen

In a vSAN configuration, if a command issued to a vSAN disk group takes more than 120 sec to complete, in a corner case the vSAN host might fail with a purple diagnostic screen. The issue occurs because after 120 seconds a red alert that indicates stuck I/O condition is triggered to vSAN and the resources for that command are released. As a result, if a given command completes after this alert, the host fails.

This issue is resolved in this release.

PR 3459526: ESXi boot might take longer on hosts with iSCSI configuration

On ESXi hosts with iSCSI configuration, the dynamic discovery of iSCSI targets might time out or fail because the firewall related to the iSCSI ruleset is not active during early boot. As a result, if many discovery addresses are configured, ESXi boot takes more time as discovery is done in a serial order.

This issue is resolved in this release.

PR 3447505: Nonuniform Sub-NUMA Clustering (SNC) sizes might cause ESXi hosts to fail with a purple diagnostic screen post upgrade to ESXi 8.0 Update 3

On ESXi hosts that use the SNC technology, if the CPU is not manufactured with equal number of cores per sub-NUMA cluster, or some cores are deactivated from the BIOS, the host fails to boot with a purple diagnostic screen post upgrade to ESXi 8.0 Update 3.

This issue is resolved in this release. Instead of replacing the hardware, ESXi 8.0 Update 3e adds a software fix of the issue when CPUs have nonuniform sub-NUMA cluster size. If the root cause is that CPUs have cores deactivated from the BIOS, just activate such cores.

PR 3486556: You cannot retrieve Active Directory domain accounts from ESXi hosts by using PowerCLI

Due to changes in the Open LDAP library, the `Get-VIAccount` PowerCLI script might not run correctly and produce the following error:

```
A general system error occurred: Error accessing directory: Can't set LDAP options.
```

As a result, you cannot retrieve information about Active Directory accounts from ESXi hosts joined to the domain.

This issue is resolved in this release.

PR 3484669: You see intermittent gaps in vSAN performance statistics

In certain cases, you can see intermittent gaps in vSAN performance statistics dashboard charts. The issue occurs mostly when the vSAN cluster is integrated with VMware Aria Operations.

This issue is resolved in this release.

PR 3469036: ESXi installation on an iSCSI LUN fails as the ESXi iSCSI Boot Firmware Table (iBFT) module fails to discover the LUN

In some cases, the iBFT control structure might have an optional expansion for which the ESXi vmkibft module doesn't have a proper check. As a result, iSCSI LUN discovery by using iBFT might fail and cause failure of ESXi installation on the iSCSI LUN.

This issue is resolved in this release.

PR 3450501: ESXi hosts working with both VMFS5 and VMFS6 datastores might fail with a purple diagnostic screen due to rare issue

On ESXi hosts that work with both VMFS5 and VMFS6 datastores, a rare issue with the I/O completion path might cause incorrect access to the data structure of a VMFS5 datastore instead of a VMFS6 datastore. As a result, the host fails with a purple diagnostic screen. The issue affects only hosts working with both VMFS5 and VMFS6 volumes and can occur without any user intervention.

This issue is resolved in this release.

PR 3433295: Virtual machines residing on a vSphere Virtual Volumes datastore might become inaccessible due to a size reporting issue

An issue in the VASA provider might report an inflated size for the home directory that contains metadata files for a virtual machine and VMFS automatically expands to accommodate the larger size. However, if the VASA provider later reports the correct size, VMFS cannot shrink, and the home directory cannot open. As a result, all VM files stored in the home folder become inaccessible and you cannot access the VMs.

This issue is resolved in this release.

PR 3431433: Storage tests might fail when using NVMe as Guest Controllers with SCSI vSphere Virtual Volumes on Windows Server Failover Clusters (WSFC)

ESXi 8.0 Update 3 added support for NVMe reservation on NVMe-backed disks on vSphere Virtual Volumes in WSFC. However, NVMe guest controllers with SCSI-backed vSphere Virtual Volumes disks are not supported and this might result in failure while running WSFC cluster validation for storage tests. The specific test that might fail is `Validate SCSI-3 Persistent Reservation`. The issue occurs when the storage array does not support implementation of the `SCSI Read_full_status Persistent Reservation` command.

This issue is resolved in this release.

PR 3483589: In very rare cases, ESXi hosts might fail with a purple diagnostic screen due to a SCSI-3 reservation failure

In a very corner case scenario, ESXi hosts with active support of SCSI-3 reservation type WEAR (write exclusive all registrant) for VMFS might fail with a purple diagnostic screen due to an issue with handling the SCSI-3 reservation failures.

This issue is resolved in this release.

PR 3441150: In the vSphere Client, you see unusually high storage utilization of vSphere Replication persistent state files (.psf) in a vSAN datastore

For vSAN datastores, the provision type of .psf depends on the configured namespace storage policy and if it is set to Thick Provisioning, the space reservation of such .psf files might seem unusually high, regardless of the actual size of vSAN objects. For example, you might see a .psf file reservation for tens of TB in the vSphere Client, while the actual vSAN storage is in tens of GB.

This issue is resolved in this release. The fix makes sure .psf files always use Thin Provisioning policy.

PR 3450290: You cannot manually deactivate the DVSSync firewall ruleset

The DVSSync firewall ruleset is part of the system-owned rulesets that are active by default to avoid possible break of service but in some cases, you might need to deactivate the rule for security reasons.

This issue is resolved in this release. You can manually deactivate the ruleset as required.

PR 3451651: ESXi host fails with a purple diagnostic screen after unsuccessful rebind operation of vSphere Virtual Volumes

In the rebind orchestration of vSphere Virtual Volumes, all ESXi hosts that have bindings to the volume must respond in a time window defined by the number of Read Block Zero (RBZ) calls that the hosts send. If a migration target fails to complete the RBZ in time, ESXi hosts might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3465047: You see no UNMAP commands sent to NVMe 2.0 targets or targets supporting TP4040

The `identify controller` command to fetch the `Unmap` attribute from a target might fail with a status `0xf` due to a buffer mismatch on the targets which support TP4040 or NVMe 2.0. As a result, you see no UNMAP commands sent to NVMe 2.0 targets or targets supporting TP4040. In the `vmkernel` log, you see errors such as:

```
command <cmd> failed: ctrlr <ctrlnum>, queue <queue num>, psaCmd <psa cmd>, status
0xf, opc 0x6, cid <cid>, nsid <nsid>Complete vmkNvmeCmd: <cmd>, vmkPsaCmd: <psa cmd>,
cmdId.initiator=<init>, CmdSN: 0x0, status: 0xfGet DSM constraints failed for device
<device name> Status Transient storage condition, suggest retryUnable to fetch UNMAP
Attributes for device <device name>
```

This issue is resolved in this release.

PR 3453046: You might see low throughput with IPv6 UDP traffic in Enhanced Datapath mode for the nmlx5 driver

RSS logic in the `nmlx5` driver incorrectly calculates hash for the IPv6 traffic which results in utilizing just a single RX queue.

This issue is resolved in this release.

PR 3427281: You see a Skyline Health alert for NICs outside the vSAN configuration

When you have active Remote Direct Memory Access (RDMA) for vSAN clusters and any cluster host has NICs outside the vSAN configuration, the health check **Network (RDMA NIC: RoCE v2) is vSAN certified** also lists the result for such NICs. This is not expected, but you can ignore the issue as it doesn't indicate any data path issue.

This issue is resolved in this release.

PR 3433436: In the vSphere Client, you see multiple alarms "Cannot login user root@xxxxx: no permission" for ESXi hosts in lockdown mode

In the **Tasks and Events** tab, you see multiple alarms `Cannot login user root@xxxxx: no permission` for ESXi hosts in lockdown mode.

This issue is resolved in this release.

PR 3433610: ESXi hosts fail with purple diagnostic screen when connecting to the network I/O scheduler

In very rare cases, the network I/O scheduler might hold a lock for a certain grace period in case of memory allocation failure during the connection. This lock might cause ESXi hosts to fail with a purple diagnostic screen and an error such as `VERIFY bora/vmkernel/sched/cpushed.c`.

In the backtrace, you see errors such as:

```
2024-08-05T02:57:57.302Z cpu20:2097484) 0x45390a61b2a0:
[0x420006423c31]RCU_WaitForGracePeriod@vmkernel#nover+0x42 stack: 0x1000000, 0x300000100,
0xbad0014, 0x420006518bc2, 0x0
```

```
2024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2d0:
[0x420006518bc1]vmk_HashRelease@vmkernel#nover+0x66 stack: 0x420006ac6eb5, 0x420007175487,
0x516574616572436b, 0x4307194064d0, 0x3030306
```

```
22024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2f0:
[0x420007175486]NetSchedHClkConnect@(netsched_hclk)#<None>+0x6ab stack: 0x30303062,
0x4300d3f27f48, 0x43126525a638, 0x0, 0x26
```

This issue is resolved in this release.

PR 3437835: AMD Turin processors with PCI Multi-segment Support preset in the BIOS might cause ESXi host failures

In AMD 5th generation EPYC 9005 series with code name Turin, **PCI Multi-segment Support** is active in the BIOS settings by default and as a result, you might see the following symptoms after installation:

- The ESXi host fails with a purple diagnostic screen.
- The ESXi host becomes unresponsive.
- Some I/O devices do not work.

The issue occurs because AMD Turin has preset PCIe multi-segment support, but ESXi does not provide PCIe multi-segment support for CPUs that have an input-output memory management unit (IOMMU). This issue is specific for AMD Turin CPUs.

This issue is resolved in this release.

PR 3451575: You don't see all paths to an IBM SAN Volume Controller (SVC) storage target

An issue in the lpfc driver might prevent the Emulex Fibre Channel Host Bus Adapter (FC HBA) from connecting to some IBM SVC storage targets. In a multipathing setup, you might see that not all storage paths are established.

This issue is resolved in this release.

PR 3469107: ESXi hosts intermittently disconnect from the Active Directory domain or from vCenter

Memory leaks might occur in Likewise during Active Directory operations or if smart card authentication is active on an ESXi host. As a result, Likewise processes might run out of memory and ESXi hosts intermittently disconnect from the Active Directory domain or from vCenter. In the `/var/log/vmkernel.log` file, you see messages similar to:

```
yyyy-mm-dd In(182) vmkernel: cpu72:2110465)uw.2110464 (44739) requires 1024 KB, asked 1024
KB from likewise (792) which has 93112 KB occupied and 72 KB available.
```

```
yyyy-mm-dd In(182) vmkernel: cpu72:2110465)Admission failure in path: host/vim/vmvisor/
likewise:lwsmd.2110464:uw.2110464
```

In the 'hostd-probe.log' file you see a message such as: `hostd detected to be non-responsive`.

This issue is resolved in this release. The fix adds a system service that monitors Likewise memory usage and restarts relevant services in case memory consumption approaches Likewise limits.

PR 3463364: Storage policy change in a vSAN ESA cluster with less than six hosts might cause performance issues during vSAN resync

If you try to change the storage policy of a large object from RAID-1 to RAID-5 in a vSAN ESA cluster with less than six hosts, the change might never complete. As a result, you might see degraded performance of the cluster and host, or disk maintenance operations might be blocked. The issue occurs because the policy change of a large object in a small vSAN ESA cluster executes in steps in smaller regions of the object and might cause repeating attempts to change the policy to RAID-5 of a region even if the change is already complete. This issue might ultimately cause VMs to become unresponsive during vSAN resync.

This issue is resolved in this release. The fix makes sure vSAN tracks storage policy changes correctly to prevent repeating attempts. If you do not update to ESXi 8.0 Update 3e, avoid changing the object policy from RAID-1 to RAID-5 in small vSAN ESA clusters.

PR 3462191: You do not see all drivers plugged in behind Intel Volume Management Device (VMD) controllers

When you have multiple drives plugged in behind VMD controllers, not all drives might be enumerating due to insufficient NRM heap memory. In the vmkernel log, you see errors such as:

```
2024-04-10T21:05:26Z In(182) vmkernel: NRM: memory pool with TAG 16 creation failed
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library is going to be stopped!
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library stopped
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library stopped
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library deinitialized started
```

This issue is resolved in this release.

PR 3474329: After deploying virtual machines, you see numerous log messages from the nmlx5 driver in the vmkernel.log

After the nmlx5 driver adds or removes the MAC address filter of a virtual machine or a VMkernel adapter, in the vmkernel.log, you might see multiple messages such as `nmlx5_en_L2TableIndexAdd` and `nmlx5_en_L2TableIndexRemove` after deploying virtual machines. These are informational logs from the nmlx5 driver indicating that filters are being applied or removed.

This issue is resolved in this release. With the fix, such messages are logged only in debug log level.

PR 3473257: NVMe/TCP controllers on an ESXi host might not recover connection when a target upgrades or restarts

When performing NVMe/TCP target upgrade or restart, controllers in the ESXi host might not recover their connection even after the target is online. As a result, you do not see the respective controllers from the vSphere Client. A rare race condition between the NVMe core layer and NVMe/TCP drivers causes the issue.

This issue is resolved in this release.

PR 3469074: ESXi NVMe/TCP controllers do not recover connectivity when a target upgrades or restarts

When performing NVMe/TCP target upgrade or restart, controllers in an ESXi host might not recover connectivity after the target is online. As a result, you do not see the respective controllers from the vSphere Client. A rare race condition between the NVMe core layer and NVMe/TCP drivers causes the issue.

This issue is resolved in this release.

PR 3477409: vSphere vMotion operations might fail due to a rare issue with no namespaces present in the namespacesmgr.db file

When a virtual machine is configured with namespaces and for some reason the namespaces table in the `namespacesmgr.db` file is empty in the ESXi host database, migration of such a VM by using vSphere vMotion might fail. In the `vmware.log` file, you see errors such as:

```
vmx - NamespaceMgrCheckpoint: No valid queue found while restoring the namespace events.
The migrate type is 1.
```

This issue is resolved in this release.

PR 3456018: ESXi hosts fail with purple diagnostic screen and message such as "PCI bus error, undiagnosed"

On ESXi hosts with active PCIe multi-segment support and hot-plug slots with active PCIe Downstream Port Containment (DPC), any `PCIe Uncorrectable Error` reported by NVME drives plugged into the hot-plug slots can result in failure of the host with a purple diagnostics screen. The error message is `PCI bus error, undiagnosed`. This may be a hardware problem; please contact your hardware vendor.

This issue is resolved in this release.

PR 3454941: You cannot create datastores on namespaces backed by some kind of NVMe/TCP storage arrays in vSphere 8.0 Update 3b environments

An issue that affects only vSphere 8.0 Update 3b environments might cause tasks to create a datastore on namespaces backed by some kind of NVMe/TCP storage arrays to fail. Frequent messages in the vmkernel logs such as `nvmectcp:nt_ReceiveR2TPdu:2667 [ctrlr 259, queue 3] txPdu 0x433685d6eb40 (cid 169) expects PDUs in types 0x28 but received PDU in type 9: Failure` are indicative for the issue.

This issue is resolved in this release.

PR 3444408: Allocation failure in Physical Region Pages (PRP) during I/O split might cause stuck I/O

Starting with vSphere 8.0 Update 3, the handling of I/O for PRP, which most local NVMe devices use to describe user data location, was transferred from the driver to the Pluggable Storage Architecture (PSA). A bug in PSA PRP handling, when a command is split into multiple child commands to handle the PRP requirements, might cause indefinitely stuck I/O. The issue occurs because if a child PRP memory allocation command fails, the parent command cannot complete. As a result, you might see an error such as `BlueScreen: ASSERT bora/modules/vmkernel/vmfs/fs3Misc.c:4245`. The issue affects only local PCIe NVMe disks.

This issue is resolved in this release. The fix adds the frame back in the path queue to make sure all child commands are processed and fixes the PRP child count when I/O fails due to issues such as path loss or no connectivity.

PR 3444303: A rare race condition when connecting to a target NVMe over TCP controller might cause ESXi failure with a purple diagnostic screen

When connecting to a target NVMe over TCP controller, either for the first time or during a controller reset, the target might fail or not respond to the connect command for some reason. As a result, a race condition occurs in the ESXi host between the NVMe core layer, and the NVMe/TCP driver and the host fails with a purple diagnostic screen.

This issue is resolved in this release. This fix prevents the race condition.

PR 3450374: A rare race between rebind and lazy unbind tasks in vSphere Virtual Volume might cause an ESXi host to fail with a purple diagnostic screen

In rare cases, a lazy unbind of a vSphere Virtual Volume might start while a rebind operation is ongoing, or the other way around. As a result, the ESXi host fails with a purple diagnostic screen.

This issue is resolved in this release.

PR 3445276: Storage path unclaiming requests might be blocked due to a Fabric Performance Impact Notification (FPIN) issue

FPIN messages serve as an early warning system to relieve negative effects on the performance of Fibre Channel devices, such as optimizing I/O behavior and avoiding impaired paths by notifying devices on the current fabric conditions. In some cases, when memory allocation fails, FPIN might keep a reference on a storage path. As a result, path scanning and claiming routine returns BUSY, as the reference is no longer existing, which blocks path claiming and unclaiming. When the issue occurs, you see the similar logs in vmkernel.log: `WARNING: StorageFPIN: 521: Failed to allocate memory.`

This issue is resolved in this release.

PR 3437903: The intrusion detection and prevention system (IDS/IPS) might report false positives for dropped packets in vSphere reporting tools

In environments with active IDS/IPS, ESXi might report false positives for dropped packets to network monitoring solutions such as Aria Operations. If you deactivate IDS/IPS, the count of dropped packets falls. The issue affects all virtual machines connected to NSX segments.

This issue is resolved in this release.

PR 3438291: ESXi hosts on a vSAN cluster might fail with purple diagnostic screen due to a rare timing issue

A rare timing issue in vSAN might cause ESXi host on vSAN cluster to fail with purple diagnostic screen. In the vmkernel logs, you see an error such as @BlueScreen: #PF Exception 14 in world 2099267:VSAN_0x432a4. In the backtrace, you see errors such as:

```
0x453af7b9bcc0:[0x42000d8e1fdf]PREF_InsertEx@com.vmware.vsanutil#0.0.0.1+0xab stack: 0x9,
0x45bc0767b7c0, 0x1a0c1841bdb50, 0x432a49db7140, 0x45bc2bd84280
```

```
0x453af7b9bd10:
[0x42000d8dcb3c]VSANServer_SetOperationTimerTC@com.vmware.vsanutil#0.0.0.1+0x69 stack:
0x45bc0767b7c0, 0x67143756669a719a, 0x45bc0767b7c0, 0xee7135183b47557, 0x45bc2bd84280
```

This issue is resolved in this release.

PR 3470500: After update to ESXi 8.0 Update 3b, you see an alarm requiring a hardware reset due to AMD erratum 1474 on all updated hosts

ESXi 8.0 Update 3b added a fix for vSphere systems becoming unresponsive in about 1044 days after the last system hardware reset due to AMD erratum 1474, related to the core-C6 (CC6) sleep state on EPYC processors of the series 7002, 7Fx2, 7Hx2, and 7001. The fix was to automatically deactivate CC6 in around 1000 days until the next hardware reset, and send alarms to notify about the required reset. However, due to an integer overflow, the timer callback triggers in about 6 days and generates the alert, deactivating CC6 and suggesting that you do a hardware reset.

This issue is resolved in this release. The fix makes sure that the deactivation of CC6 and triggering the alarm occurs in the scheduled 1000 days.

PR 3410296: You see gaps in the collection of metrics for some virtual machines

Due to a time sync issue, collection of data for some VMs might fail and you do not see any performance data in platforms such as Aria Operations. In the `/var/run/log/vobd.log` on the ESXi host, you see errors such as: 1149:2024-06-11T05:01:11.250Z: [ClockCorrelator] 7826272852729us:

```
[vob.clock.correction.adjusttime.sync] system clock synchronized to upstream time servers
```

This issue is resolved in this release. The fix caches results from Stats Registry queries with selected dynamic counters to make sure metrics are consistent.

PR 3434172: You cannot apply a reverse path filter for virtual machines in promiscuous mode even /Net/ReversePathFwdCheckPromisc is active

In very specific conditions, after an upgrade of a distributed virtual switch or uninstalling an NSX transport node on an ESXi host with converged distributed virtual switch that contains a port group with active promiscuous mode, you might see some networking issues, including:

1. You cannot apply a reverse path filter for virtual machines in promiscuous mode even `/Net/ReversePathFwdCheckPromisc` is active.
2. Intermittent packet connectivity issues when using promiscuous mode on a distributed virtual switch with multiple uplinks.

This issue is resolved in this release.

PR 3442088: vSphere vMotion tasks fail with an error NamespaceMgr could not lock the db file

When virtual machines are configured with namespaces, an issue with the namespace database might cause migration of such VMs with vSphere vMotion to fail. In the vSphere Client, you see errors such as:

```
Failed to receive migration. The source detected that the destination failed to resume.
```

```
An error occurred restoring the virtual machine state during migration. NamespaceMgr could not lock the db file.
```

The destination ESXi host `vmware.log` shows the following messages:

[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.checkpoint.migration.failedReceive] Failed to receive migration.

[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.namespaceMgr.noLock] NamespaceMgr could not lock the db file.

The issue is more likely to occur on VMs with hardware versions earlier than 19 and impacts only VMs provisioned on shared datastores other than NFS, such as VMFS, vSAN, and vSphere Virtual Volumes.

This issue was reported as known in KB [369767](#) and is resolved in this release.

PR 339022: Capture of network packets by using the PacketCapture tool on ESXi does not work

Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in [Collecting network packets using the lightweight PacketCapture on ESXi](#).

This issue is resolved in this release.

ESXi-8.0U3e-24674464-no-tools

Profile Name	ESXi-8.0U3e-24674464-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	April 10, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner_8.0.3-0.70.24674464 - VMware_bootbank_vsan_8.0.3-0.70.24674464 - VMware_bootbank_infravisor_8.0.3-0.70.24674464 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.70.24674464 - VMware_bootbank_crx_8.0.3-0.70.24674464 - VMware_bootbank_bmcad_8.0.3-0.70.24674464 - VMware_bootbank_vsanhealth_8.0.3-0.70.24674464 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.70.24674464 - VMware_bootbank_gc_8.0.3-0.70.24674464 - VMware_bootbank_vdfs_8.0.3-0.70.24674464 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.70.24674464 - VMware_bootbank_esx-base_8.0.3-0.70.24674464 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.70.24674464 - VMware_bootbank_esx-xserver_8.0.3-0.70.24674464 - VMware_bootbank_vds-vsip_8.0.3-0.70.24674464 - VMware_bootbank_cpu-microcode_8.0.3-0.70.24674464 - VMware_bootbank_esxio_8.0.3-0.70.24674464 - VMware_bootbank_clusterstore_8.0.3-0.70.24674464 - VMware_bootbank_esxio-base_8.0.3-0.70.24674464 - VMware_bootbank_trx_8.0.3-0.70.24674464 - VMware_bootbank_gc-esxio_8.0.3-0.70.24674464 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.70.24674464

	• VMware_bootbank_bmcad-esxio_8.0.3-0.70.24674464 • VMware_bootbank_vcls-pod-crx_8.0.3-0.70.24674464 • VMware_bootbank_native-misc-drivers_8.0.3-0.70.24674464 • VMware_bootbank_loadesx_8.0.3-0.70.24674464 • VMware_bootbank_esx-update_8.0.3-0.70.24674464 • VMware_bootbank_loadesxio_8.0.3-0.70.24674464 • VMware_bootbank_esxio-update_8.0.3-0.70.24674464 • VMW_bootbank_ntg3_4.1.15.0-4vmw.803.0.70.24674464 • VMW_bootbank_lpf_14.4.0.40-35vmw.803.0.70.24674464 • VMW_bootbank_iavmd_3.0.0.1010-12vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-cc_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-cc-esxio_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-core_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-rdma_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-core-esxio_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nmlx5-rdma-esxio_4.23.6.5-1vmw.803.0.70.24674464 • VMW_bootbank_nvmetcp_1.0.1.32-1vmw.803.0.70.24674464 • VMW_bootbank_nvmetcp-esxio_1.0.1.32-1vmw.803.0.70.24674464 • VMW_bootbank_vmkusb-esxio_0.1-23vmw.803.0.70.24674464 • VMW_bootbank_vmkusb_0.1-23vmw.803.0.70.24674464
PRs Fixed	3498633, 3341273, 3410296, 3414552, 3415171, 3425316, 3427281, 3480453, 3431433, 3432906, 3433295, 3433436, 3433610, 3434171, 3434270, 3437835, 3437903, 3438291, 3439279, 3444039, 3444137, 3444408, 3445065, 3445276, 3446362, 3447505, 3447686, 3447687, 3447888, 3448872, 3449523, 3450290, 3450374, 3450501, 3451651, 3451669, 3451849, 3452081, 3452460, 3452487, 3452906, 3453391, 3453981, 3454108, 3454123, 3454926, 3454942, 3454983, 3455021, 3455315, 3456018, 3457039, 3457042, 3457050, 3457350, 3457744, 3457900, 3457959, 3458521, 3459100, 3459378, 3459526, 3459825, 3460312, 3460819, 3463099, 3463147, 3463364, 3465047, 3465070, 3465259, 3465275, 3465373, 3465624, 3465965, 3466011, 3466176, 3466283, 3466538, 3467137, 3467521, 3467708, 3468136, 3468139, 3468215, 3469014, 3469036, 3469145, 3469467, 3469652, 3470385, 3470427, 3470500, 3470714, 3470732, 3472623, 3473154, 3473626, 3473692, 3473767, 3473768, 3473860, 3474427, 3474443, 3474967, 3475015, 3475152, 3475393, 3476179, 3476630, 3476631, 3476632, 3476634, 3476963, 3477021, 3477062, 3477385, 3477409, 3477725, 3478001, 3478002, 3478236, 3478276, 3478486, 3479338, 3479340, 3479464, 3480154, 3480456, 3480968, 3481216, 3481241, 3481310, 3481386, 3481475, 3481507, 3482259, 3482802,

	3482963, 3483589, 3483610, 3483714, 3484669, 3484787, 3484836, 3485671, 3485971, 3486600, 3486647, 3486747, 3486955, 3487351, 3487380, 3488415, 3488480, 3494691, 3497983, 3462191, 3442016, 3478908, 3444303, 3454941, 3469074, 3473257, 3458162, 3474996
Related CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

This patch updates the following issues:

PR 3478908: On ConnectX-4 and later physical network adapters, you see slower TX transmission when packets in the same queue have different values for the Priority Code Point (PCP)

Packets with different VLAN (CoS PCP/801.2p) priorities might be transmitted slower on ConnectX-4 and later physical network adapters.

This issue is resolved in this release. The fix adds the `trust_dscp_same_priority` nmlx5 driver module parameter to map all DSCP priorities to the same traffic class to avoid slowness in the NIC hardware.

PR 3442191: Update of a vSAN environment with Original Storage Architecture (OSA) to version 8.0 Update 3 might be slower compared to previous updates

The first reboot of the vSphere system in a vSAN OSA environment following upgrade or update to 8.0 Update 3 might cause a delay in the update process, which depends on the workload running at the time of update. Nodes with heavy workloads running before the update are likely to cause higher delays. This is expected and for updates to 8.0 Update 3, environments with large install base must consider the possible delay in the update process.

This issue is resolved in this release.

PR 3506135: You see messages from the nmlx5 driver such as "Health: Miss counters detected" with "synd 0x0" and "IO was aborted"

The nmlx5 driver health logic might incorrectly identify the state of a NIC as faulty and fail firmware commands from the driver. In the vmkernel logs, you see similar messages at random times:

```
<NMLX_ERR> nmlx5_core: 0000:45:00.0: Health: Miss counters detected.
<NMLX_INF> synd 0x0: unrecognized error
<NMLX_INF> extSynd 0x0000
<NMLX_ERR> nmlx5_QueryNicVportContext:188 command failed: IO was aborted
<NMLX_ERR> nmlx5_QueryVportCounter:1851 command failed: IO was aborted
```

This issue is resolved in this release.

ESXi 8.0 Update 3e adds support for vSphere Quick Boot to:

- Intel vRAN Baseband Driver
- Intel Platform Monitoring Technology Driver
- Intel Data Center Graphics Driver
- AMD Instinct MI Series Driver

Support for Communication Device Class Network Control Model (CDC-NCM) in the ESXi USB driver:

Starting with ESX 8.0 Update 3e, the ESXi USB driver supports the CDC-NCM protocol for compatibility with HPE Gen12 iLO Virtual NIC and interoperability with HPE Agentless Management (AMS), Integrated Smart Update Tools (iSUT), the iLOrest config tool, Intelligent Provisioning, and DPUs.

PR 3479338: Delayed cleanup of a deleted component can lead to high space usage reports for vSAN datastores

The vSAN I/O scheduler manages in-flight I/Os for performance and fairness. Under certain conditions, the scheduler might get an incorrect counter for the in-flight resync I/Os, which blocks new resync I/Os. As a result, the cleanup of blocked resync I/Os is delayed and might lead to high space usage.

This issue is resolved in this release.

PR 3442016: You might see low throughput with North-South encapsulated traffic with NSX Edge VMs

In releases earlier than ESXi 8.0 Update 3e, the nmlx5 driver does not support decapsulation of GENEVE encapsulated packets with BF3 NICs. As a result, you might see low throughput with North-South encapsulated traffic with NSX Edge VMs.

This issue is resolved in this release. The fix adds support for decapsulation of GENEVE encapsulated packets with BF3 NICs for the nmlx5 driver.

PR 3476179: A memory leak might prevent the Object Storage File System Daemon (osfsd) from processing requests for vSphere Virtual Volumes

A rare memory leak condition in the osfsd service after prolonged uptime might prevent vSphere Virtual Volumes from provisioning virtual machines, including creation and cloning operations.

This issue is resolved in this release.

PR 3459100: Virtual machine migration in a vSphere Virtual Volumes datastore might cause an ESXi host to fail with a purple diagnostic screen

During virtual machine migration from a vSphere Virtual Volumes datastore, a rare race condition in Device Block allocation tracking might lead to an ESXi host failure, resulting in a purple diagnostic screen.

This issue is resolved in this release. The fix prevents the race condition.

PR 3484787: If you reconfigure the scratch location and reboot an ESXi host, all logging stops and the host becomes unresponsive

If you have audit recording to local storage enabled for an ESXi host on a directory configured at the default location `/scratch/auditLog` or at a subdirectory of `/scratch`, and then reconfigure `/scratch`, the syslog daemon fails after a host reboot, because it cannot find the record directory and files in the new location. As a result, all logging on such hosts stops and eventually the host becomes unresponsive.

This issue is resolved in this release. The fix handles such scenarios and deactivates audit records to local storage for security reasons, generating a VMkernel Observation (VOB) message such as `vSphere error: (esx.problem.vmsyslogd.auditrecord.local.disabled)` but preventing a failure of the vmsyslog daemon. You can re-enable audit recording in the new `/scratch` location after the host reboots. Old audit record files are available in the previously configured scratch path, if it is still accessible.

PR 3455315: The PowerSvc plug-in of the sensord service stops polling data by using Intelligent Platform Management Interface (IPMI) without a retry

Due to a change related to the Running Average Power Limit (RAPL) host power reporting in ESXi 8.0 Update 3, the sensord service might deactivate power monitoring by using IPMI if only one IPMI power data poll fails. The PowerSvc plug-in does not retry the poll and IPMI polling stops. PowerSvc goes offline and logs `failing to read from the power provider in the /var/log/sensord.log`.

This issue is resolved in this release. The fix makes sure the PowerSvc plug-in retries IPMI power data polls in case of failure and IPMI data collection continues.

PR 3460819: Virtual machines with UnixWare OS and a vlane virtual device might repeatedly reset when running on an ESXi 8.x host

An issue with vlane virtual devices might cause virtual machines with UnixWare OS on ESXi 8.x hosts to repeatedly fail due to a reboot of the guest OS with the error `Your guest operating system has accessed an I/O range where multiple devices overlap`.

This issue is resolved in this release.

PR 3483165: The Data Center Bridging group (DCB) state of nmlx5 driver NICs in ESXi hosts is always active in Hardware DCB mode

The DCB state of ESXi hosts is always active in VMNICs for the nmlx5 driver, irrespective of the DCB state of the physical switch.

This issue is resolved in this release. With the fix, the switch DCB state determines the DCB state of NICs in ESXi hosts when NICs are in Hardware DCB mode.

PR 3479340: You see checksum errors for some objects in a stretched cluster after a network outage

After a network outage or some other critical event on a stretched cluster, the vSAN Distributed Object Manager (DOM) might fail to get the correct checksums for objects that do not have sub-fault tolerance. As a result, you might see in tools such as VMware Skyline Health Diagnostics checksum errors for many objects in the stretched cluster.

This issue is resolved in this release.

PR 3461166: In the vSphere Client, you cannot change a VMNIC on the virtual switch of an ESXi host

The cache of the hostd service might not be up-to-date after you add a VMNIC to the virtual switch on an ESXi host due to a bug in the teaming policy of port groups. As a result, when you remove another VMNIC from the switch, both the newly added and the removed VMNICs are removed from the teaming policy of the port groups. When you work in the vSphere Client, the task seems to revert.

This issue is resolved in this release.

PR 3463099: Guest OS might mark file system read-only after repeated IO failure which can cause Kubernetes control plane nodes to go into Read-only mode

An issue with the check for max allowed memory address in memory mapping code during I/O processing at virtual SCSI might cause a guest OS on an ESXi host to mark guest file system as read-only. As a result, Kubernetes control plane nodes might go into Read-only mode. In the vmkernel logs, you see multiple error messages such as `failed to pin: I/O error`.

This issue is resolved in this release.

PR 3461321: Custom TLS settings of the Fault Domain Manager (FDM) might not work on ESXi hosts of earlier versions that you manage with a vCenter of version 8.0 Update 3

Custom TLS settings of the vmware-fdm service at port 8182 might not work on hosts of earlier versions in a cluster with active High Availability that you manage with a vCenter of version 8.0 Update 3.

This issue is resolved in this release. For more information on the workflows and the use of such custom settings, see KB [320798](#).

PR 3488415: ESXi host fails with a purple diagnostic screen with messages such as #PF Exception 14 in world XXXXXXXX:vmknvmeGener

In the discovery phase of NVMe over Fabrics (NVMe-oF), some targets might report to an ESXi host some subsystem NVMe Qualified Names (NQN) to which that ESXi host cannot connect due to missing permissions or unavailable route. The ESXi host continues attempting to connect to such NQNs and if a network issue occurs, in rare cases, the host might fail with a purple diagnostic screen with messages such as `#PF Exception 14 in world XXXXXXXX:vmknvmeGener`.

This issue is resolved in this release.

PR 3486647: Copying files with the cp command creates files with unexpected access timestamp (atime) and modification timestamp (mtime) values

If vSAN File Services NFS 4.1 is mounted to a RHEL 8 or later NFS client, when copying a file by using the `cp` command, you might see `atime` or `mtime` such as 1970 instead of the current date.

This issue is resolved in this release.

PR 3453042: An ESXi host fails with a purple diagnostic screen due to uplink watchdog resets on nmlx5 driver NICs

In rare cases, when ESXi on data processing units (DPU) fails with a kernel panic, the ESXi host also fails with a purple diagnostic screen. The issue occurs on ESXi hosts with virtual machines that have virtual NICs on a distributed switch backed by DPUs, when an uplink watchdog resets nmlx5 driver NICs.

This issue is resolved in this release.

PR 3486025: Virtual machines and interactive shells might become unresponsive due to an issue with the initialization of user worlds in ESXi

By default, the initialization of user worlds in ESXi adds an active group that requires the allocation of some CPU resources. If the initialization fails for some reason before the CPU resource is available, the creation of new user worlds starts failing. As a result, virtual machines and interactive shells might become unresponsive.

This issue is resolved in this release.

PR 3469652: Logging stops on all log files and configured remote host targets on ESXi hosts

A rare deadlock situation in the ESXi log daemon vmsyslogd might cause logging to stop on all log files and configured remote host targets on ESXi hosts. The issue might occur as the result of multiple syslog reloads.

This issue is resolved in this release.

PR 3466011: You see the datastore name in storage VMkernel Observation (VOB) reports as Unknown

Due to a missing function to grant access required to fetch datastore names, you might see the datastore name in `esx.problem.storage.vob` messages as unknown.

For example: 2024-08-27T04:31:29.150Z In(14) vobd[2097668]: [scsiCorrelator] 5815563033719us: [esx.problem.storage.redundancy.lost] Lost path redundancy to storage device naa.60060e80072b600000302b6000000026. Path vmhba1:C0:T1:L0 is down. Affected datastores: Unknown.

This issue is resolved in this release.

PR 3482259: NVMe over TCP controllers might stay offline and not recover even when the network and the target port have no issues

Due to issues with the network or the target, NVMe over TCP controllers in an ESXi host might enter an offline state. Usually, controllers recover and turn online as soon as the issue is resolved, but in some very rare scenarios, the controllers might stay offline and do not recover, which makes the NVMe/TCP namespaces inaccessible. The issue occurs due to a dead lock when many controllers are being reset and at the same time many I/O requests time out, which triggers multiple taskmgmt requests that compete for critical resource. The only way to exit this state is to reboot the ESXi host.

This issue is resolved in this release. The fix avoids such dead locks.

PR 3481216: "Route based on physical NIC load" policy does not work as expected

An issue with the permissions for the net-blt daemon might cause load balancing based on the **Routing based on PNIC load** teaming policy to not work as expected.

This issue is resolved in this release.

PR 3486955: Storage path unclaiming requests might be blocked due to a Fabric Performance Impact Notification (FPIN) issue

FPIN messages serve as an early warning system to relieve negative effects on the performance of Fibre Channel devices, such as optimizing I/O behavior and avoiding impaired paths by notifying devices on the current fabric conditions.

In some cases, when memory allocation fails, FPIN might keep a reference on a storage path. As a result, path scanning and claiming routine returns `BUSY`, as the reference no longer exists, which blocks path claiming and unclaiming, and FPIN-related information cannot go through to ESXi hosts. In the ESXCLI command output you do not see FPIN messages. When the issue occurs, you see a message in the VMkernel logs such as: `WARNING: StorageFPIN: 521: Failed to allocate memory.`

Another symptom is when after a SAN maintenance or in case of a storage path outage FPIN holds a reference count on the storage paths and the FC driver cannot allocate new paths or reestablish existing ones. In `/var/log/vmkernel.log`, you see the following sequence repeating:

```
WARNING: nfnic: <2>: fnic_handle_report_lun: 1467: lun add failure! in_remove: 0
ioAllowed: 1
```

```
WARNING: nfnic: <2>: fnic_tport_event_handler: 2130: lunmap update failed, retry ...
```

This issue is resolved in this release.

PR 3341273: Migration of virtual machines fails with an error "Failed to allocate migration heap" on the source ESXi host

In very rare cases, if a virtual machine powers off before a running migration completes, ESXi might fail to clean up the migration heap. If the memory leak continues over time, the migration heap might be exhausted, and virtual machines cannot migrate from such ESXi hosts.

This issue is resolved in this release.

PR 3467137: A race condition during heap cleanup might cause an ESXi host to fail with a purple diagnostic screen

A rare race condition when reading the VMX file descriptor table occurs at the time VMX shuts down might cause an ESXi host to fail with a purple diagnostic screen. The error message is similar to `PANIC bora/vmkernel/main/dlmalloc.c:4944 - Usage error in dlmalloc.`

This issue is resolved in this release.

PR 3483610: If you insert an invalid or blank DVD on an ESXi host and configure it as a host device on virtual machines, the VM performance degrades

If a blank or incompatible disc is in the DVD drive of an ESXi host, the ESXi HPP module checks for the SCSI sense key and additional sense keys to determine next steps. Due to a missing check to handle Active State Change (ASC) codes 0x30 and 0x3a, the check fails, causing a failover that keeps repeating. As a result, you see a degradation of performance on the virtual machines on such hosts.

This issue is resolved in this release. The fix handles the ASC.

PR 3475015: Duplicate traffic on multiple uplinks might cause traffic outage

During an upgrade of the physical switch infrastructure, if you use the `TeamPolicyUpDelay` setting and all uplinks on the distributed switch flap, a fallback mode during the time of delay might cause duplicate traffic on all uplinks of the switch, which leads to traffic outage.

This issue is resolved in this release. The fix adds the advanced config parameter `/Net/TeamingIgnoreShotgun` to prevent the fallback mode in such cases and avoid duplicate traffic.

PR 3451306: The usage scale in the CPU chart under Monitor > Performance Overview shows in thousands, and you cannot notice CPU usage spikes to 100%

The usage scale in the CPU chart under **Monitor > Performance Overview** for virtual machines might show in thousands, reaching up to 2400%, instead of hundreds to allow you to notice CPU usage spikes to 100%. The expanded scale also prevents realistic estimation of other key performance metrics on the graph, such as CPU Ready, as they are displayed at the bottom of the chart.

This issue is resolved in this release.

PR 3470732: Virtual machines using Raw Device Mapping (RDM) might become unresponsive due storage heap exhaustion

An issue in tracking SCSI Inquiry or Mode Sense requests that virtual machines issue to RDM LUNs might lead to heap exhaustion at the ESXi storage layer. As a result, such VMs become unresponsive. In the vSphere Client, you see series of **Lost access to volume** and **Successfully restored access** messages until you power off the VMs.

This issue is resolved in this release.

PR 3474443: A memory leak in the swapObj daemon (swapobjd) during object replication might cause a failure of the service due to stale metadata

An issue with stale metadata for swap objects might cause a memory leak in the swapobjd during object replication, leading to gradual memory exhaustion and ultimately to a failure of the service.

This issue is resolved in this release. The fix ensures stale swap object metadata is removed.

PR 3479464: ESXi hosts might fail with purple diagnostic screen while processing SCSI commands from the guest OS

An issue with the validation of SCSI commands from the guest OS at the virtual SCSI layer might cause an ESXi host to fail with a purple diagnostic screen.

This issue is resolved in this release. The fix adds enhanced command validation for SCSI commands at the virtual SCSI layer.

PR 3481507: Discovery of iSCSI dynamic targets, LUNs, and datastores fails on ESXi host reboot

In some cases, dynamic discovery of targets in systems that use dynamic discovery addresses for iSCSI configuration might fail. As a result, you do not see iSCSI targets, LUNs, and datastores after a reboot of ESXi hosts. The dynamic discovery of iSCSI targets might also fail when the Challenge-Handshake Authentication Protocol (CHAP) is active on the ESXi host.

This issue is resolved in this release.

PR 3454942: In the vSphere Client, you cannot reconfigure a vGPU virtual machine if more than 128 vCPUs are active

If you add more than 128 vCPUs on a VM with a vGPU, you cannot manage the configuration from the vSphere Client. The issue occurs because input-output memory management unit (IOMMU) support for a vGPU with more than 128 vCPUs is not available.

This issue is resolved in this release. With the fix, you can manage the vGPU VM configuration from the vSphere Client. However, the VM might not be able to power on, since IOMMU support is not available. To make sure the reconfiguration succeeds, add to the .vmx file the following 2 parameters: `pciPassthru.vmiop.allowViommu="TRUE"` and `pciPassthru.vmiop.enableViommu="TRUE"`.

PR 3477772: Encrypted virtual machines with active Change Block Tracking (CBT) might intermittently power off after a rekey operation

Due to a race condition between the VMX and hostd services for a specific VMcrypt-related reconfiguration, encrypted VMs with active CBT might unexpectedly power off during a rekey operation.

This issue is resolved in this release.

PR 3469014: ESXi host might fail with a purple diagnostic screen if the Multipath Plug-in (MPP) does not implement the event handler for Fabric Performance Impact Notification (FPIN)

Description: ESXi 8.0 Update 3 added APIs to notify MPPs in case of Fibre channel FPIN events. However, if the MPP is compiled against an API SDK of version earlier than 8.0 Update 3, or if the MPP does not implement the new event handler, then the ESXi host fails with a purple diagnostic screen on an FPIN event.

This issue is resolved in this release.

PR 3465259: NVMe Dell PowerStore clustered VMDK datastore fails to recover after an All-Paths-Down (APD) situation

A reservation conflict might prevent a NVMe Dell PowerStore clustered VMDK datastore to take a reservation and cause failure to recover after losing the reservation.

This issue is resolved in this release.

PR 3480154: High VMFS heap memory consumption might cause virtual machines to become unresponsive

In environments with many large datastores and many VMs on thin-provisioned disks, VMFS resource allocation might consume a high amount of heap memory, proportional to the size of the datastores. As a result, VMFS heap memory might be exhausted and lead to VMs becoming unresponsive or vSphere vMotion tasks failing with errors for VM invalid state.

This issue is resolved in this release.

PR 3460312: Virtual machines on NFSv3 datastores might fail during snapshot consolidation if a third-party backup is also active

Snapshot consolidation during backup with third-party software, such as NetBackup, might fail due to extra read-only locks on the VM snapshot file while the snapshot consolidation is in progress and trying to upgrade that same lock. As a result, a lock conflict might cause a VM to fail.

This issue is resolved in this release. The fix prevents snapshot consolidation tasks to upgrade the read-only lock when more than 1 holders for that lock exist.

PR 3457744: Migration of virtual machines to 8.0 Update 3 and later ESXi host might result in TPM null hierarchy issues

A rare issue related to how applications use TPM null hierarchy might lead to failure to authenticate virtual machines after migrating them from an ESXi host of version earlier than 8.0 Update 3 to 8.0 Update 3 and later. You can identify the issue if the command `Get-TpmSupportedFeature` in Windows returns NULL or the command `sudo tpm2_createprimary -C n -c context.out` in Linux fails.

This issue is resolved in this release. If you already face the issue, power cycle the virtual machines.

PR 3470385: Restoring the iSCSI configurations on an ESXi host might fail during boot due to stale entries in the ConfigStore database

If the ConfigStore database on an ESXi host has stale iSCSI configurations related to hardware iSCSI adapters, the iSCSI jumpstart might fail and restoring the iSCSI configuration on that host also fails during boot.

This issue is resolved in this release.

PR 3477062: vSAN hosts might become unresponsive due to high latencies and high CPU usage during peaks of guest unmap traffic

Internal threads in vSAN might cause a CPU busy loop during a peak of unmap traffic. As a result, some hosts might become unresponsive.

This issue is resolved in this release. The fix is relinquishing CPUs when unmap requests start to slow down.

PR 3466538: A race condition between the TCP control path and the keep alive timer might cause an ESXi host to fail with a purple diagnostic screen

A rare race condition between the TCP control path during a disconnect operation and the keep alive timer might cause ESXi hosts of version 8.0 Update 3 to fail with a purple diagnostic screen. The issue occurs under specific network conditions when the timer and disconnect logic overlap.

This issue is resolved in this release. The fix prevents the purple diagnostic screen failure.

PR 3473626: If a command issued to a vSAN disk group complete after 120 seconds, the vSAN host might fail with a purple diagnostic screen

In a vSAN configuration, if a command issued to a vSAN disk group takes more than 120 sec to complete, in a corner case the vSAN host might fail with a purple diagnostic screen. The issue occurs because after 120 seconds a red alert that indicates stuck I/O condition is triggered to vSAN and the resources for that command are released. As a result, if a given command completes after this alert, the host fails.

This issue is resolved in this release.

PR 3459526: ESXi boot might take longer on hosts with iSCSI configuration

On ESXi hosts with iSCSI configuration, the dynamic discovery of iSCSI targets might time out or fail because the firewall related to the iSCSI ruleset is not active during early boot. As a result, if many discovery addresses are configured, ESXi boot takes more time as discovery is done in a serial order.

This issue is resolved in this release.

PR 3447505: Nonuniform Sub-NUMA Clustering (SNC) sizes might cause ESXi hosts to fail with a purple diagnostic screen post upgrade to ESXi 8.0 Update 3

On ESXi hosts that use the SNC technology, if the CPU is not manufactured with equal number of cores per sub-NUMA cluster, or some cores are deactivated from the BIOS, the host fails to boot with a purple diagnostic screen post upgrade to ESXi 8.0 Update 3.

This issue is resolved in this release. Instead of replacing the hardware, ESXi 8.0 Update 3e adds a software fix of the issue when CPUs have nonuniform sub-NUMA cluster size. If the root cause is that CPUs have cores deactivated from the BIOS, just activate such cores.

PR 3486556: You cannot retrieve Active Directory domain accounts from ESXi hosts by using PowerCLI

Due to changes in the Open LDAP library, the `Get-VIAccount` PowerCLI script might not run correctly and produce the following error:

```
A general system error occurred: Error accessing directory: Can't set LDAP options.
```

As a result, you cannot retrieve information about Active Directory accounts from ESXi hosts joined to the domain.

This issue is resolved in this release.

PR 3484669: You see intermittent gaps in vSAN performance statistics

In certain cases, you can see intermittent gaps in vSAN performance statistics dashboard charts. The issue occurs mostly when the vSAN cluster is integrated with VMware Aria Operations.

This issue is resolved in this release.

PR 3469036: ESXi installation on an iSCSI LUN fails as the ESXi iSCSI Boot Firmware Table (iBFT) module fails to discover the LUN

In some cases, the iBFT control structure might have an optional expansion for which the ESXi `vmkibft` module doesn't have a proper check. As a result, iSCSI LUN discovery by using iBFT might fail and cause failure of ESXi installation on the iSCSI LUN.

This issue is resolved in this release.

PR 3450501: ESXi hosts working with both VMFS5 and VMFS6 datastores might fail with a purple diagnostic screen due to rare issue

On ESXi hosts that work with both VMFS5 and VMFS6 datastores, a rare issue with the I/O completion path might cause incorrect access to the data structure of a VMFS5 datastore instead of a VMFS6 datastore. As a result, the host fails with a purple diagnostic screen. The issue affects only hosts working with both VMFS5 and VMFS6 volumes and can occur without any user intervention.

This issue is resolved in this release.

PR 3433295: Virtual machines residing on a vSphere Virtual Volumes datastore might become inaccessible due to a size reporting issue

An issue in the VASA provider might report an inflated size for the home directory that contains metadata files for a virtual machine and VMFS automatically expands to accommodate the larger size. However, if the VASA provider later reports the correct size, VMFS cannot shrink, and the home directory cannot open. As a result, all VM files stored in the home folder become inaccessible and you cannot access the VMs.

This issue is resolved in this release.

PR 3431433: Storage tests might fail when using NVMe as Guest Controllers with SCSI vSphere Virtual Volumes on Windows Server Failover Clusters (WSFC)

ESXi 8.0 Update 3 added support for NVMe reservation on NVMe-backed disks on vSphere Virtual Volumes in WSFC. However, NVMe guest controllers with SCSI-backed vSphere Virtual Volumes disks are not supported and this might result in failure while running WSFC cluster validation for storage tests. The specific test that might fail is `Validate SCSI-3 Persistent Reservation`. The issue occurs when the storage array does not support implementation of the `SCSI Read_full_status Persistent Reservation` command.

This issue is resolved in this release.

PR 3483589: In very rare cases, ESXi hosts might fail with a purple diagnostic screen due to a SCSI-3 reservation failure

In a very corner case scenario, ESXi hosts with active support of SCSI-3 reservation type WEAR (write exclusive all registrant) for VMFS might fail with a purple diagnostic screen due to an issue with handling the SCSI-3 reservation failures.

This issue is resolved in this release.

PR 3441150: In the vSphere Client, you see unusually high storage utilization of vSphere Replication persistent state files (.psf) in a vSAN datastore

For vSAN datastores, the provision type of .psf depends on the configured namespace storage policy and if it is set to Thick Provisioning, the space reservation of such .psf files might seem unusually high, regardless of the actual size of vSAN objects. For example, you might see a .psf file reservation for tens of TB in the vSphere Client, while the actual vSAN storage is in tens of GB.

This issue is resolved in this release. The fix makes sure .psf files always use Thin Provisioning policy.

PR 3450290: You cannot manually deactivate the DVSSync firewall ruleset

The DVSSync firewall ruleset is part of the system-owned rulesets that are active by default to avoid possible break of service but in some cases, you might need to deactivate the rule for security reasons.

This issue is resolved in this release. You can manually deactivate the ruleset as required.

PR 3451651: ESXi host fails with a purple diagnostic screen after unsuccessful rebind operation of vSphere Virtual Volumes

In the rebind orchestration of vSphere Virtual Volumes, all ESXi hosts that have bindings to the volume must respond in a time window defined by the number of Read Block Zero (RBZ) calls that the hosts send. If a migration target fails to complete the RBZ in time, ESXi hosts might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3465047: You see no UNMAP commands sent to NVMe 2.0 targets or targets supporting TP4040

The `identify controller` command to fetch the `Unmap` attribute from a target might fail with a status `0xf` due to a buffer mismatch on the targets which support TP4040 or NVMe 2.0. As a result, you see no `UNMAP` commands sent to NVMe 2.0 targets or targets supporting TP4040. In the vmkernel log, you see errors such as:

```
command <cmd> failed: ctrlr <ctrlnum>, queue <queue num>, psaCmd <psa cmd>, status
0xf, opc 0x6, cid <cid>, nsid <nsid>Complete vmkNvmeCmd: <cmd>, vmkPsaCmd: <psa cmd>,
cmdId.initiator=<init>, CmdSN: 0x0, status: 0xfGet DSM constraints failed for device
<device name> Status Transient storage condition, suggest retryUnable to fetch UNMAP
Attributes for device <device name>
```

This issue is resolved in this release.

PR 3453046: You might see low throughput with IPv6 UDP traffic in Enhanced Datapath mode for the nmlx5 driver

RSS logic in the nmlx5 driver incorrectly calculates hash for the IPv6 traffic which results in utilizing just a single RX queue.

This issue is resolved in this release.

PR 3427281: You see a Skyline Health alert for NICs outside the vSAN configuration

When you have active Remote Direct Memory Access (RDMA) for vSAN clusters and any cluster host has NICs outside the vSAN configuration, the health check **Network (RDMA NIC: RoCE v2) is vSAN certified** also lists the result for such NICs. This is not expected, but you can ignore the issue as it doesn't indicate any data path issue.

This issue is resolved in this release.

PR 3433436: In the vSphere Client, you see multiple alarms "Cannot login user root@xxxxx: no permission" for ESXi hosts in lockdown mode

In the **Tasks and Events** tab, you see multiple alarms `Cannot login user root@xxxxx: no permission` for ESXi hosts in lockdown mode.

This issue is resolved in this release.

PR 3433610: ESXi hosts fail with purple diagnostic screen when connecting to the network I/O scheduler

In very rare cases, the network I/O scheduler might hold a lock for a certain grace period in case of memory allocation failure during the connection. This lock might cause ESXi hosts to fail with a purple diagnostic screen and an error such as `VERIFY bora/vmkernel/sched/cpushed.c`.

In the backtrace, you see errors such as:

```
2024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2a0:
[0x420006423c31]RCU_WaitForGracePeriod@vmkernel#nover+0x42 stack: 0x1000000, 0x300000100,
0xbad0014, 0x420006518bc2, 0x0

2024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2d0:
[0x420006518bc1]vmk_HashRelease@vmkernel#nover+0x66 stack: 0x420006ac6eb5, 0x420007175487,
0x516574616572436b, 0x4307194064d0, 0x3030306

2024-08-05T02:57:57.302Z cpu20:2097484)0x45390a61b2f0:
[0x420007175486]NetSchedHClkConnect@(netsched_hclk)#<None>+0x6ab stack: 0x30303062,
0x4300d3f27f48, 0x43126525a638, 0x0, 0x26
```

This issue is resolved in this release.

PR 3437835: AMD Turin processors with PCI Multi-segment Support preset in the BIOS might cause ESXi host failures

In AMD 5th generation EPYC 9005 series with code name Turin, **PCI Multi-segment Support** is active in the BIOS settings by default and as a result, you might see the following symptoms after installation:

- The ESXi host fails with a purple diagnostic screen.
- The ESXi host becomes unresponsive.
- Some I/O devices do not work.

The issue occurs because AMD Turin has preset PCIe multi-segment support, but ESXi does not provide PCIe multi-segment support for CPUs that have an input-output memory management unit (IOMMU). This issue is specific for AMD Turin CPUs.

This issue is resolved in this release.

PR 3451575: You don't see all paths to an IBM SAN Volume Controller (SVC) storage target

An issue in the lpfc driver might prevent the Emulex Fibre Channel Host Bus Adapter (FC HBA) from connecting to some IBM SVC storage targets. In a multipathing setup, you might see that not all storage paths are established.

This issue is resolved in this release.

PR 3469107: ESXi hosts intermittently disconnect from the Active Directory domain or from vCenter

Memory leaks might occur in Likewise during Active Directory operations or if smart card authentication is active on an ESXi host. As a result, Likewise processes might run out of memory and ESXi hosts intermittently disconnect from the Active Directory domain or from vCenter. In the `/var/log/vmkernel.log` file, you see messages similar to:

```
yyyy-mm-dd In(182) vmkernel: cpu72:2110465)uw.2110464 (44739) requires 1024 KB, asked 1024 KB from likewise (792) which has 93112 KB occupied and 72 KB available.
```

```
yyyy-mm-dd In(182) vmkernel: cpu72:2110465)Admission failure in path: host/vim/vmvisor/likewise:lwsmc.2110464:uw.2110464
```

In the 'hostd-probe.log' file you see a message such as: `hostd detected to be non-responsive.`

This issue is resolved in this release. The fix adds a system service that monitors Likewise memory usage and restarts relevant services in case memory consumption approaches Likewise limits.

PR 3463364: Storage policy change in a vSAN ESA cluster with less than six hosts might cause performance issues during vSAN resync

If you try to change the storage policy of a large object from RAID-1 to RAID-5 in a vSAN ESA cluster with less than six hosts, the change might never complete. As a result, you might see degraded performance of the cluster and host, or disk maintenance operations might be blocked. The issue occurs because the policy change of a large object in a small vSAN ESA cluster executes in steps in smaller regions of the object and might cause repeating attempts to change the policy to RAID-5 of a region even if the change is already complete. This issue might ultimately cause VMs to become unresponsive during vSAN resync.

This issue is resolved in this release. The fix makes sure vSAN tracks storage policy changes correctly to prevent repeating attempts. If you do not update to ESXi 8.0 Update 3e, avoid changing the object policy from RAID-1 to RAID-5 in small vSAN ESA clusters.

PR 3462191: You do not see all drivers plugged in behind Intel Volume Management Device (VMD) controllers

When you have multiple drives plugged in behind VMD controllers, not all drives might be enumerating due to insufficient NRM heap memory. In the vmkernel log, you see errors such as:

```
2024-04-10T21:05:26Z In(182) vmkernel: NRM: memory pool with TAG 16 creation failed
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library is going to be stopped!
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library stopped
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library stopped
2024-04-10T21:05:26Z In(182) vmkernel: NRM: library deinitialized started
```

This issue is resolved in this release.

PR 3474329: After deploying virtual machines, you see numerous log messages from the nmlx5 driver in the vmkernel.log

After the nmlx5 driver adds or removes the MAC address filter of a virtual machine or a VMkernel adapter, in the vmkernel.log, you might see multiple messages such as `nmlx5_en_L2TableIndexAdd` and

`nmlx5_en_L2TableIndexRemove` after deploying virtual machines. These are informational logs from the `nmlx5` driver indicating that filters are being applied or removed.

This issue is resolved in this release. With the fix, such messages are logged only in debug log level.

PR 3473257: NVMe/TCP controllers on an ESXi host might not recover connection when a target upgrades or restarts

When performing NVMe/TCP target upgrade or restart, controllers in the ESXi host might not recover their connection even after the target is online. As a result, you do not see the respective controllers from the vSphere Client. A rare race condition between the NVMe core layer and NVMe/TCP drivers causes the issue.

This issue is resolved in this release.

PR 3469074: ESXi NVMe/TCP controllers do not recover connectivity when a target upgrades or restarts

When performing NVMe/TCP target upgrade or restart, controllers in an ESXi host might not recover connectivity after the target is online. As a result, you do not see the respective controllers from the vSphere Client. A rare race condition between the NVMe core layer and NVMe/TCP drivers causes the issue.

This issue is resolved in this release.

PR 3477409: vSphere vMotion operations might fail due to a rare issue with no namespaces present in the `namespacemgr.db` file

When a virtual machine is configured with namespaces and for some reason the namespaces table in the `namespacemgr.db` file is empty in the ESXi host database, migration of such a VM by using vSphere vMotion might fail. In the `vmware.log` file, you see errors such as:

```
vmx - NamespaceMgrCheckpoint: No valid queue found while restoring the namespace events.
The migrate type is 1.
```

This issue is resolved in this release.

PR 3456018: ESXi hosts fail with purple diagnostic screen and message such as "PCI bus error, undiagnosed"

On ESXi hosts with active PCIe multi-segment support and hot-plug slots with active PCIe Downstream Port Containment (DPC), any PCIe Uncorrectable Error reported by NVMe drives plugged into the hot-plug slots can result in failure of the host with a purple diagnostics screen. The error message is `PCI bus error, undiagnosed`. This may be a hardware problem; please contact your hardware vendor.

This issue is resolved in this release.

PR 3454941: You cannot create datastores on namespaces backed by some kind of NVMe/TCP storage arrays in vSphere 8.0 Update 3b environments

An issue that affects only vSphere 8.0 Update 3b environments might cause tasks to create a datastore on namespaces backed by some kind of NVMe/TCP storage arrays to fail. Frequent messages in the `vmkernel` logs such as `nvmectp:nt_ReceiverR2TPdu:2667 [ctlr 259, queue 3] txPdu 0x433685d6eb40 (cid 169) expects PDUs in types 0x28 but received PDU in type 9: Failure` are indicative for the issue.

This issue is resolved in this release.

PR 3444408: Allocation failure in Physical Region Pages (PRP) during I/O split might cause stuck I/O

Starting with vSphere 8.0 Update 3, the handling of I/O for PRP, which most local NVMe devices use to describe user data location, was transferred from the driver to the Pluggable Storage Architecture (PSA). A bug in PSA PRP handling, when a command is split into multiple child commands to handle the PRP requirements, might cause indefinitely stuck I/O. The issue occurs because if a child PRP memory allocation command fails, the parent command cannot complete. As a result, you might see an error such as `BlueScreen: ASSERT bora/modules/vmkernel/vmfs/fs3Misc.c:4245`. The issue affects only local PCIe NVMe disks.

This issue is resolved in this release. The fix adds the frame back in the path queue to make sure all child commands are processed and fixes the PRP child count when I/O fails due to issues such as path loss or no connectivity.

PR 3444303: A rare race condition when connecting to a target NVMe over TCP controller might cause ESXi failure with a purple diagnostic screen

When connecting to a target NVMe over TCP controller, either for the first time or during a controller reset, the target might fail or not respond to the connect command for some reason. As a result, a race condition occurs in the ESXi host between the NVMe core layer, and the NVMe/TCP driver and the host fails with a purple diagnostic screen.

This issue is resolved in this release. This fix prevents the race condition.

PR 3450374: A rare race between rebind and lazy unbind tasks in vSphere Virtual Volume might cause an ESXi host to fail with a purple diagnostic screen

In rare cases, a lazy unbind of a vSphere Virtual Volume might start while a rebind operation is ongoing, or the other way around. As a result, the ESXi host fails with a purple diagnostic screen.

This issue is resolved in this release.

PR 3445276: Storage path unclaiming requests might be blocked due to a Fabric Performance Impact Notification (FPIN) issue

FPIN messages serve as an early warning system to relieve negative effects on the performance of Fibre Channel devices, such as optimizing I/O behavior and avoiding impaired paths by notifying devices on the current fabric conditions. In some cases, when memory allocation fails, FPIN might keep a reference on a storage path. As a result, path scanning and claiming routine returns BUSY, as the reference is no longer existing, which blocks path claiming and unclaiming. When the issue occurs, you see the similar logs in vmkernel.log: `WARNING: StorageFPIN: 521: Failed to allocate memory.`

This issue is resolved in this release.

PR 3437903: The intrusion detection and prevention system (IDS/IPS) might report false positives for dropped packets in vSphere reporting tools

In environments with active IDS/IPS, ESXi might report false positives for dropped packets to network monitoring solutions such as Aria Operations. If you deactivate IDS/IPS, the count of dropped packets falls. The issue affects all virtual machines connected to NSX segments.

This issue is resolved in this release.

PR 3438291: ESXi hosts on a vSAN cluster might fail with purple diagnostic screen due to a rare timing issue

A rare timing issue in vSAN might cause ESXi host on vSAN cluster to fail with purple diagnostic screen. In the vmkernel logs, you see an error such as `@BlueScreen: #PF Exception 14 in world 2099267:VSAN_0x432a4`. In the backtrace, you see errors such as:

```
0x453af7b9bcc0:[0x42000d8e1fdf]PREF_InsertEx@com.vmware.vsanutil#0.0.0.1+0xab stack: 0x9,
0x45bc0767b7c0, 0x1a0c1841bdb50, 0x432a49db7140, 0x45bc2bd84280
```

```
0x453af7b9bd10:
[0x42000d8dcb3c]VSANServer_SetOperationTimerTC@com.vmware.vsanutil#0.0.0.1+0x69 stack:
0x45bc0767b7c0, 0x67143756669a719a, 0x45bc0767b7c0, 0xeec7135183b47557, 0x45bc2bd84280
```

This issue is resolved in this release.

PR 3470500: After update to ESXi 8.0 Update 3b, you see an alarm requiring a hardware reset due to AMD erratum 1474 on all updated hosts

ESXi 8.0 Update 3b added a fix for vSphere systems becoming unresponsive in about 1044 days after the last system hardware reset due to AMD erratum 1474, related to the core-C6 (CC6) sleep state on EPYC processors of the series 7002, 7Fx2, 7Hx2, and 7001. The fix was to automatically deactivate CC6 in around 1000 days until the next hardware

reset, and send alarms to notify about the required reset. However, due to an integer overflow, the timer callback triggers in about 6 days and generates the alert, deactivating CC6 and suggesting that you do a hardware reset.

This issue is resolved in this release. The fix makes sure that the deactivation of CC6 and triggering the alarm occurs in the scheduled 1000 days.

PR 3410296: You see gaps in the collection of metrics for some virtual machines

Due to a time sync issue, collection of data for some VMs might fail and you do not see any performance data in platforms such as Aria Operations. In the `/var/run/log/vobd.log` on the ESXi host, you see errors such as: `1149:2024-06-11T05:01:11.250Z: [ClockCorrelator] 7826272852729us: [vob.clock.correction.adjtime.sync] system clock synchronized to upstream time servers`

This issue is resolved in this release. The fix caches results from Stats Registry queries with selected dynamic counters to make sure metrics are consistent.

PR 3434172: You cannot apply a reverse path filter for virtual machines in promiscuous mode even /Net/ReversePathFwdCheckPromisc is active

In very specific conditions, after an upgrade of a distributed virtual switch or uninstalling an NSX transport node on an ESXi host with converged distributed virtual switch that contains a port group with active promiscuous mode, you might see some networking issues, including:

1. You cannot apply a reverse path filter for virtual machines in promiscuous mode even `/Net/ReversePathFwdCheckPromisc` is active.
2. Intermittent packet connectivity issues when using promiscuous mode on a distributed virtual switch with multiple uplinks.

This issue is resolved in this release.

PR 3442088: vSphere vMotion tasks fail with an error NamespaceMgr could not lock the db file

When virtual machines are configured with namespaces, an issue with the namespace database might cause migration of such VMs with vSphere vMotion to fail. In the vSphere Client, you see errors such as:

Failed to receive migration. The source detected that the destination failed to resume.

An error occurred restoring the virtual machine state during migration. NamespaceMgr could not lock the db file.

The destination ESXi host `vmware.log` shows the following messages:

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.checkpoint.migration.failedReceive] Failed to receive migration.
```

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.namespaceMgr.noLock] NamespaceMgr could not lock the db file.
```

The issue is more likely to occur on VMs with hardware versions earlier than 19 and impacts only VMs provisioned on shared datastores other than NFS, such as VMFS, vSAN, and vSphere Virtual Volumes.

This issue was reported as known in KB [369767](#) and is resolved in this release.

PR 3390022: Capture of network packets by using the PacketCapture tool on ESXi does not work

Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in [Collecting network packets using the lightweight PacketCapture on ESXi](#).

This issue is resolved in this release.

ESXi-8.0U3se-24659227-standard

Profile Name	ESXi-8.0U3se-24659227-standard
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	April 10, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_drivervm-gpu-base_8.0.3-0.65.24659227 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.65.24659227 - VMware_bootbank_vsan_8.0.3-0.65.24659227 - VMware_bootbank_esxio-combiner_8.0.3-0.65.24659227 - VMware_bootbank_bmcad-esxio_8.0.3-0.65.24659227 - VMware_bootbank_trx_8.0.3-0.65.24659227 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.65.24659227 - VMware_bootbank_crx_8.0.3-0.65.24659227 - VMware_bootbank_cpu-microcode_8.0.3-0.65.24659227 - VMware_bootbank_clusterstore_8.0.3-0.65.24659227 - VMware_bootbank_esxio_8.0.3-0.65.24659227 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.65.24659227 - VMware_bootbank_vcls-pod-crx_8.0.3-0.65.24659227 - VMware_bootbank_gc_8.0.3-0.65.24659227 - VMware_bootbank_vds-vsip_8.0.3-0.65.24659227 - VMware_bootbank_gc-esxio_8.0.3-0.65.24659227 - VMware_bootbank_infravisor_8.0.3-0.65.24659227 - VMware_bootbank_vsanhealth_8.0.3-0.65.24659227 - VMware_bootbank_esx-xserver_8.0.3-0.65.24659227 - VMware_bootbank_native-misc-drivers_8.0.3-0.65.24659227 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.65.24659227 - VMware_bootbank_bmcad_8.0.3-0.65.24659227 - VMware_bootbank_vdfs_8.0.3-0.65.24659227 - VMware_bootbank_esxio-base_8.0.3-0.65.24659227 - VMware_bootbank_esx-base_8.0.3-0.65.24659227 - VMware_bootbank_loadesx_8.0.3-0.65.24659227 - VMware_bootbank_esx-update_8.0.3-0.65.24659227 - VMware_bootbank_loadesxio_8.0.3-0.65.24659227 - VMware_bootbank_esxio-update_8.0.3-0.65.24659227 - VMware_locker_tools-light_12.5.1.24649672-24659227
PRs Fixed	3404662, 3432571, 3432690, 3435458, 3458134, 3444037, 3444037, 3444626, 3447888, 3449523, 3452700, 3454671, 3454671, 3454671, 3457664, 3470207, 3442312, 3494209, 3451575
Related CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

This patch updates the following issues:

ESXi 8.0 Update 3e includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series

Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series

Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003801	8/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003801	8/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	baseline	0x7002a01	8/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series

Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd0003f5	8/2/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002c0	8/8/2024	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xba	8/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x3a	8/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x54	8/1/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	baseline	0x2c0003e0	7/30/2024	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b000620	7/30/2024	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xfa	7/28/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series

Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x102	7/31/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x63	8/1/2024	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x12c	9/25/2024	Intel Xeon E-2400 Series
Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x21000291	7/30/2024	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series

VMware Tools Bundling Changes in ESXi 8.0 Update 3e

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 3e:

- **windows.iso**: VMware Tools 12.5.1 supports 64-bit Windows versions - Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- **windows-x86.iso**: VMware Tools 12.4.6 supports 32-bit Windows versions.
- **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.5.1 Release Notes](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Upgrading VMware Tools](#)

ESXi-8.0U3se-24659227-no-tools

Profile Name	ESXi-8.0U3se-24659227-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	April 10, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	• VMware_bootbank_drivervm-gpu-base_8.0.3-0.65.24659227

	- VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.65.24659227 - VMware_bootbank_vsan_8.0.3-0.65.24659227 - VMware_bootbank_esxio-combiner_8.0.3-0.65.24659227 - VMware_bootbank_bmcad-esxio_8.0.3-0.65.24659227 - VMware_bootbank_trx_8.0.3-0.65.24659227 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.65.24659227 - VMware_bootbank_crx_8.0.3-0.65.24659227 - VMware_bootbank_cpu-microcode_8.0.3-0.65.24659227 - VMware_bootbank_clusterstore_8.0.3-0.65.24659227 - VMware_bootbank_esxio_8.0.3-0.65.24659227 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.65.24659227 - VMware_bootbank_vcls-pod-crx_8.0.3-0.65.24659227 - VMware_bootbank_gc_8.0.3-0.65.24659227 - VMware_bootbank_vds-vsip_8.0.3-0.65.24659227 - VMware_bootbank_gc-esxio_8.0.3-0.65.24659227 - VMware_bootbank_infravisor_8.0.3-0.65.24659227 - VMware_bootbank_vsanhealth_8.0.3-0.65.24659227 - VMware_bootbank_esx-xserver_8.0.3-0.65.24659227 - VMware_bootbank_native-misc-drivers_8.0.3-0.65.24659227 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.65.24659227 - VMware_bootbank_bmcad_8.0.3-0.65.24659227 - VMware_bootbank_vdfs_8.0.3-0.65.24659227 - VMware_bootbank_esxio-base_8.0.3-0.65.24659227 - VMware_bootbank_esx-base_8.0.3-0.65.24659227 - VMware_bootbank_loadesx_8.0.3-0.65.24659227 - VMware_bootbank_esx-update_8.0.3-0.65.24659227 - VMware_bootbank_loadesxio_8.0.3-0.65.24659227 - VMware_bootbank_esxio-update_8.0.3-0.65.24659227
PRs Fixed	3404662, 3432571, 3432690, 3435458, 3458134, 3444037, 3444037, 3444626, 3447888, 3449523, 3452700, 3454671, 3454671, 3454671, 3457664, 3470207, 3442312, 3494209, 3451575
Related CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

This patch updates the following issues:

ESXi 8.0 Update 3e includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series

Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655- LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600- QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100- C-v2 Series; Intel Pentium B925C

Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series

Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003801	8/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003801	8/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/ b)	0xbf	baseline	0x7002a01	8/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series

Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd0003f5	8/2/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002c0	8/8/2024	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xba	8/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x3a	8/1/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x54	8/1/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	baseline	0x2c0003e0	7/30/2024	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b000620	7/30/2024	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xfa	7/28/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x102	7/31/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x63	8/1/2024	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x12c	9/25/2024	Intel Xeon E-2400 Series

Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x21000291	7/30/2024	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series
-------------------	-------------------	------	----------	------------	-----------	--

VMware Tools Bundling Changes in ESXi 8.0 Update 3e

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 3e:

- **windows.iso**: VMware Tools 12.5.1 supports 64-bit Windows versions - Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- **windows-x86.iso**: VMware Tools 12.4.6 supports 32-bit Windows versions.
- **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.5.1 Release Notes](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Upgrading VMware Tools](#)

ESXi 8.0 U3e - 24674464

Name	ESXi
Version	ESXi 8.0 U3e - 24674464
Release Date	April 10, 2025
Category	Bugfix
Affected Components	• ESXi Component - core ESXi VIBs • ESXi Install/Upgrade Component • ESXi Install/Upgrade Component • Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers • Broadcom NetXtreme I ESX VMKAPI ethernet driver • Broadcom Emulex Connectivity Division FC and FCoE Driver • Intel NVME Driver with VMD Technology • VMware NVMe over TCP Driver • USB Native Driver for VMware • ESXi Tools Component

PRs Fixed	3498633, 3341273, 3410296, 3414552, 3415171, 3425316, 3427281, 3480453, 3431433, 3432906, 3433295, 3433436, 3433610, 3434171, 3434270, 3437835, 3437903, 3438291, 3439279, 3444039, 3444137, 3444408, 3445065, 3445276, 3446362, 3447505, 3447686, 3447687, 3447888, 3448872, 3449523, 3450290, 3450374, 3450501, 3451651, 3451669, 3451849, 3452081, 3452460, 3452487, 3452906, 3453391, 3453981, 3454108, 3454123, 3454926, 3454942, 3454983, 3455021, 3455315, 3456018, 3457039, 3457042, 3457050, 3457350, 3457744, 3457900, 3457959, 3458521, 3459100, 3459378, 3459526, 3459825, 3460312, 3460819, 3463099, 3463147, 3463364, 3465047, 3465070, 3465259, 3465275, 3465373, 3465624, 3465965, 3466011, 3466176, 3466283, 3466538, 3467137, 3467521, 3467708, 3468136, 3468139, 3468215, 3469014, 3469036, 3469145, 3469467, 3469652, 3470385, 3470427, 3470500, 3470714, 3470732, 3472623, 3473154, 3473626, 3473692, 3473767, 3473768, 3473860, 3474427, 3474443, 3474967, 3475015, 3475152, 3475393, 3476179, 3476630, 3476631, 3476632, 3476634, 3476963, 3477021, 3477062, 3477385, 3477409, 3477725, 3478001, 3478002, 3478236, 3478276, 3478486, 3479338, 3479340, 3479464, 3480154, 3480456, 3480968, 3481216, 3481241, 3481310, 3481386, 3481475, 3481507, 3482259, 3482802, 3482963, 3483589, 3483610, 3483714, 3484669, 3484787, 3484836, 3485671, 3485971, 3486600, 3486647, 3486747, 3486955, 3487351, 3487380, 3488415, 3488480, 3494691, 3497983, 3462191, 3442016, 3478908, 3444303, 3454941, 3469074, 3473257, 3458162, 3474996
Related CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

This patch resolves the issues listed in ESXi_8.0.3-0.70.24674464.

ESXi 8.0 U3se - 24659227

Name	ESXi
Version	ESXi 8.0 U3se - 24659227
Release Date	April 10, 2025
Category	Security
Affected Components	• ESXi Component - core ESXi VIBs • ESXi Install/Upgrade Component • ESXi Tools Component
PRs Fixed	3404662, 3432571, 3432690, 3435458, 3458134, 3444037, 3444037, 3444626, 3447888, 3449523, 3452700, 3454671, 3454671, 3454671, 3457664, 3470207, 3442312, 3494209, 3451575
Related CVE numbers	CVE-2025-41226, CVE-2025-41227, CVE-2025-41228

This patch resolves the issues listed in ESXi_8.0.3-0.65.24659227.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

You cannot update to ESXi 8.0 Update 2b by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

```
Please use a supported method to upgrade ESXi.
```

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see [Upgrade or Update a Host with Image Profiles](#).

The Cancel option in an interactive ESXi installation might not work as expected

Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

The `irdman` driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic

If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the `irdman` driver might fail. This issue is unlikely to occur, as the `irdman` driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

Workaround: Use ULPs with RC transport mode.

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1

The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out

vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD

In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See [Reconciling Discrepancies in the Managed Virtual Disk Catalog](#).

vCenter Server and vSphere Client Issues

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

VMware ESXi 8.0 Update 3d Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 3d | 04 MAR 2025 | Build 24585383
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

Patches Contained in This Release

VMware ESXi 8.0 Update 3d

Build Details

VMware vSphere Hypervisor (ESXi) Offline Bundle

Download Filename:	VMware-ESXi-8.0U3d-24585383-depot.zip
Build:	24585383
Download Size:	629.0 MB
SHA256 checksum:	71c09c47a35fd94bc726696a2b11d3df63a6f6fe71344864412f9fabca29251a
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.3-0.60.24585383	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.60.24585383	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.60.24585383	Security	Critical

Rollup Bulletins

These rollup bulletins contain the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U3d-24585383	Security	Critical	Security image

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U3d-24585383-standard
ESXi-8.0U3d-24585383-no-tools

ESXi Images

Name and Version	Release Date	Category	Detail
ESXi8.0U3d - 24585383	04 March 2025	Security	Security image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.3-0.60.24585383

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_gc-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.60.24585383 - VMware_bootbank_esxio-combiner_8.0.3-0.60.24585383 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.60.24585383 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.60.24585383 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.60.24585383 - VMware_bootbank_cpu-microcode_8.0.3-0.60.24585383

	- VMware_bootbank_vds-vsip_8.0.3-0.60.24585383 - VMware_bootbank_gc_8.0.3-0.60.24585383 - VMware_bootbank_infravisor_8.0.3-0.60.24585383 - VMware_bootbank_bmcals_8.0.3-0.60.24585383 - VMware_bootbank_esxio-base_8.0.3-0.60.24585383 - VMware_bootbank_vdfs_8.0.3-0.60.24585383 - VMware_bootbank_vsan_8.0.3-0.60.24585383 - VMware_bootbank_clusterstore_8.0.3-0.60.24585383 - VMware_bootbank_vsanhealth_8.0.3-0.60.24585383 - VMware_bootbank_native-misc-drivers_8.0.3-0.60.24585383 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esxio_8.0.3-0.60.24585383 - VMware_bootbank_bmcals-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esx-xserver_8.0.3-0.60.24585383 - VMware_bootbank_vcls-pod-crx_8.0.3-0.60.24585383 - VMware_bootbank_crx_8.0.3-0.60.24585383 - VMware_bootbank_esx-base_8.0.3-0.60.24585383 - VMware_bootbank_trx_8.0.3-0.60.24585383
PRs Fixed	N/A
CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch updates the `esx-base` VIB. Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `gc-esxio`, `esx-dvfilter-generic-fastpath`, `esxio-combiner`, `drivervm-gpu-base`, `esxio-dvfilter-generic-fastpath`, `esxio-combiner-esxio`, `cpu-microcode`, `vds-vsip`, `gc`, `infravisor`, `bmcals`, `esxio-base`, `vdfs`, `vsan`, `clusterstore`, `vsanhealth`, `native-misc-drivers`, `native-misc-drivers-esxio`, `esxio`, `bmcals-esxio`, `esx-xserver`, `vcls-pod-crx`, `crx`, `esx-base`, and `trx`. This patch resolves the following issue:

This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

esx-update_8.0.3-0.60.24585383

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esx-update_8.0.3-0.60.24585383 - VMware_bootbank_loadesx_8.0.3-0.60.24585383
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesx` and `esx-update`.

esxio-update_8.0.3-0.60.24585383

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esxio-update_8.0.3-0.60.24585383 - VMware_bootbank_loadesxio_8.0.3-0.60.24585383
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesxio` and `esxio-update`.

ESXi-8.0U3d-24585383-standard

Profile Name	ESXi-8.0U3d-24585383-standard
Build	For build information, see Patches Contained in This Release .
Vendor	Broadcom, Inc.
Release Date	March 4, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_gc-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.60.24585383 - VMware_bootbank_esxio-combiner_8.0.3-0.60.24585383 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.60.24585383 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.60.24585383 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.60.24585383 - VMware_bootbank_cpu-microcode_8.0.3-0.60.24585383 - VMware_bootbank_vds-vsip_8.0.3-0.60.24585383 - VMware_bootbank_gc_8.0.3-0.60.24585383 - VMware_bootbank_infravisor_8.0.3-0.60.24585383 - VMware_bootbank_bmcal_8.0.3-0.60.24585383 - VMware_bootbank_esxio-base_8.0.3-0.60.24585383 - VMware_bootbank_vdfs_8.0.3-0.60.24585383 - VMware_bootbank_vsan_8.0.3-0.60.24585383 - VMware_bootbank_clusterstore_8.0.3-0.60.24585383 - VMware_bootbank_vsanhealth_8.0.3-0.60.24585383 - VMware_bootbank_native-misc-drivers_8.0.3-0.60.24585383 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esxio_8.0.3-0.60.24585383

	- VMware_bootbank_bmcad-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esx-xserver_8.0.3-0.60.24585383 - VMware_bootbank_vcls-pod-crx_8.0.3-0.60.24585383 - VMware_bootbank_crx_8.0.3-0.60.24585383 - VMware_bootbank_esx-base_8.0.3-0.60.24585383 - VMware_bootbank_trx_8.0.3-0.60.24585383 - VMware_bootbank_esx-update_8.0.3-0.60.24585383 - VMware_bootbank_loadesx_8.0.3-0.60.24585383 - VMware_bootbank_esxio-update_8.0.3-0.60.24585383 - VMware_bootbank_loadesxio_8.0.3-0.60.24585383
PRs Fixed	N/A
Related CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch updates the following issue:

This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

ESXi-8.0U3d-24585383-no-tools

Profile Name	ESXi-8.0U3d-24585383-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	Broadcom, Inc.
Release Date	March 4, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_gc-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.60.24585383 - VMware_bootbank_esxio-combiner_8.0.3-0.60.24585383 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.60.24585383 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.60.24585383 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.60.24585383 - VMware_bootbank_cpu-microcode_8.0.3-0.60.24585383 - VMware_bootbank_vds-vsip_8.0.3-0.60.24585383 - VMware_bootbank_gc_8.0.3-0.60.24585383 - VMware_bootbank_infravisor_8.0.3-0.60.24585383 - VMware_bootbank_bmcad_8.0.3-0.60.24585383 - VMware_bootbank_esxio-base_8.0.3-0.60.24585383 - VMware_bootbank_vdfs_8.0.3-0.60.24585383 - VMware_bootbank_vsan_8.0.3-0.60.24585383 - VMware_bootbank_clusterstore_8.0.3-0.60.24585383 - VMware_bootbank_vsanhealth_8.0.3-0.60.24585383 - VMware_bootbank_native-misc-drivers_8.0.3-0.60.24585383 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.60.24585383

	- VMware_bootbank_esxio_8.0.3-0.60.24585383 - VMware_bootbank_bmcad-esxio_8.0.3-0.60.24585383 - VMware_bootbank_esx-xserver_8.0.3-0.60.24585383 - VMware_bootbank_vcls-pod-crx_8.0.3-0.60.24585383 - VMware_bootbank_crx_8.0.3-0.60.24585383 - VMware_bootbank_esx-base_8.0.3-0.60.24585383 - VMware_bootbank_trx_8.0.3-0.60.24585383 - VMware_bootbank_esx-update_8.0.3-0.60.24585383 - VMware_bootbank_loadesx_8.0.3-0.60.24585383 - VMware_bootbank_esxio-update_8.0.3-0.60.24585383 - VMware_bootbank_loadesxio_8.0.3-0.60.24585383
PRs Fixed	N/A
Related CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch updates the following issue:

This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

ESXi8.0U3d - 24585383

Name	ESXi
Version	ESXi8.0U3d - 24585383
Release Date	March 4, 2025
Category	Security
Affected Components	- ESXi Component - core ESXi VIBs - ESXi Install/Upgrade Component
PRs Fixed	N/A
Related CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch resolves the issues listed in [ESXi_8.0.3-0.60.24585383](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

You cannot update to ESXi 8.0 Update 2b by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see [Upgrade or Update a Host with Image Profiles](#).

The Cancel option in an interactive ESXi installation might not work as expected

Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings > <specific broadcom NIC> > Device Configuration Menu > Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

The irdman driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic

If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the irdman driver might fail. This issue is unlikely to occur, as the irdman driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

Workaround: Use ULPs with RC transport mode.

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1

The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

Capture of network packets by using the PacketCapture tool on ESXi does not work

Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in [Collecting network packets using the lightweight PacketCapture on ESXi](#).

Workaround: Use the pktcap-uw tool. For more information, see [Capture and Trace Network Packets by Using the pktcap-uw Utility](#).

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out

vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD

In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See [Reconciling Discrepancies in the Managed Virtual Disk Catalog](#).

vCenter Server and vSphere Client Issues

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

VMware ESXi 8.0 Update 3c Release Notes

This document contains the following sections

- [Introduction](#)

- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 3c | 12 DEC 2024 | Build 24414501
Check for additions and updates to these release notes.

What's New

- This release resolves an issue with vSphere vMotion tasks that fail with an error `NamespaceMgr could not lock the db file`.

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0d Release Notes](#)
- [VMware ESXi 8.0 Update 3b Release Notes](#)
- [VMware ESXi 8.0 Update 3 Release Notes](#)
- [VMware ESXi 8.0 Update 2c Release Notes](#)
- [VMware ESXi 8.0 Update 1d Release Notes](#)
- [VMware ESXi 8.0 Update 2b Release Notes](#)
- [VMware ESXi 8.0 Update 2 Release Notes](#)
- [VMware ESXi 8.0 Update 1c Release Notes](#)
- [VMware ESXi 8.0 Update 1a Release Notes](#)
- [VMware ESXi 8.0 Update 1 Release Notes](#)
- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Patches Contained in This Release

VMware ESXi 8.0 Update 3c

Build Details

VMware vSphere Hypervisor (ESXi) Offline Bundle

Download Filename:	VMware-ESXi-8.0U3c-24414501-depot.zip
Build:	24414501
Download Size:	629.0 MB
SHA256 checksum:	58d0632d3e51adf26ffacce57873e72ff41151d92a8378234f2daa476f2cfba5

Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.3-0.55.24414501	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.55.24414501	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.55.24414501	Bugfix	Critical

Rollup Bulletins

These rollup bulletins contain the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U3c-24414501	Bugfix	Critical	Bugfix image

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U3c-24414501-standard
ESXi-8.0U3c-24414501-no-tools

ESXi Images

Name and Version	Release Date	Category	Detail
ESXi8.0U3c - 24414501	12 DEC 2024	Bugfix	Bugfix image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.3-0.55.24414501

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-base_8.0.3-0.55.24414501 - VMware_bootbank_vcls-pod-crx_8.0.3-0.55.24414501 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.55.24414501 - VMware_bootbank_gc-esxio_8.0.3-0.55.24414501 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.55.24414501 - VMware_bootbank_trx_8.0.3-0.55.24414501 - VMware_bootbank_clusterstore_8.0.3-0.55.24414501 - VMware_bootbank_esx-xserver_8.0.3-0.55.24414501 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.55.24414501 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.55.24414501 - VMware_bootbank_infravisor_8.0.3-0.55.24414501 - VMware_bootbank_native-misc-drivers_8.0.3-0.55.24414501 - VMware_bootbank_bmc-al-esxio_8.0.3-0.55.24414501 - VMware_bootbank_vdfs_8.0.3-0.55.24414501 - VMware_bootbank_esx-base_8.0.3-0.55.24414501 - VMware_bootbank_vsan_8.0.3-0.55.24414501 - VMware_bootbank_vds-vsip_8.0.3-0.55.24414501 - VMware_bootbank_gc_8.0.3-0.55.24414501 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.55.24414501 - VMware_bootbank_bmc-al_8.0.3-0.55.24414501 - VMware_bootbank_esxio-combiner_8.0.3-0.55.24414501 - VMware_bootbank_crx_8.0.3-0.55.24414501 - VMware_bootbank_vsanhealth_8.0.3-0.55.24414501 - VMware_bootbank_cpu-microcode_8.0.3-0.55.24414501 - VMware_bootbank_esxio_8.0.3-0.55.24414501
PRs Fixed	3459675
CVE numbers	N/A

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

This patch updates the esxio-base, vcls-pod-crx, esxio-dvfilter-generic-fastpath, gc-esxio, drivervm-gpu-base, trx, clusterstore, esx-xserver, esx-dvfilter-generic-fastpath, native-misc-drivers-esxio, infravisor, native-misc-drivers, bmc-al-esxio, vdfs, esx-base, vsan, vds-vsip, gc, esxio-combiner-esxio, bmc-al, esxio-combiner, crx, vsanhealth, cpu-microcode, and esxio VIBs.

This patch resolves the following issue:

vSphere vMotion tasks fail with an error NamespaceMgr could not lock the db file

In rare cases, when virtual machines are configured with namespaces, an issue with the namespace database might cause migration of such VMs with vSphere vMotion to fail. In the vSphere Client, you see errors such as:

Failed to receive migration. The source detected that the destination failed to resume.

An error occurred restoring the virtual machine state during migration. NamespaceMgr could not lock the db file.

The destination ESXi host `vmware.log` shows the following messages:

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.checkpoint.migration.failedReceive] Failed to receive migration.
```

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.namespaceMgr.noLock] NamespaceMgr could not lock the db file.
```

The issue is more likely to occur on VMs with hardware version earlier than 19, and impacts only VMs provisioned on shared datastores other than NFS, such as VMFS, vSAN, and vSphere Virtual Volumes.

This issue was reported as known in KB [369767](#) and is resolved in this release.

esx-update_8.0.3-0.55.24414501

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esx-update_8.0.3-0.55.24414501 - VMware_bootbank_loadesx_8.0.3-0.55.24414501
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesx` and `esx-update`.

esxio-update_8.0.3-0.55.24414501

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esxio-update_8.0.3-0.55.24414501 - VMware_bootbank_loadesxio_8.0.3-0.55.24414501
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesxio` and `esxio-update`.

ESXi-8.0U3c-24414501-standard

Profile Name	ESXi-8.0U3c-24414501-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	December 12, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-base_8.0.3-0.55.24414501 - VMware_bootbank_vcls-pod-crx_8.0.3-0.55.24414501 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.55.24414501 - VMware_bootbank_gc-esxio_8.0.3-0.55.24414501 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.55.24414501 - VMware_bootbank_trx_8.0.3-0.55.24414501 - VMware_bootbank_clusterstore_8.0.3-0.55.24414501 - VMware_bootbank_esx-xserver_8.0.3-0.55.24414501 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.55.24414501 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.55.24414501 - VMware_bootbank_infravisor_8.0.3-0.55.24414501 - VMware_bootbank_native-misc-drivers_8.0.3-0.55.24414501 - VMware_bootbank_bmcad-esxio_8.0.3-0.55.24414501 - VMware_bootbank_vdfs_8.0.3-0.55.24414501 - VMware_bootbank_esx-base_8.0.3-0.55.24414501 - VMware_bootbank_vsan_8.0.3-0.55.24414501 - VMware_bootbank_vds-vsip_8.0.3-0.55.24414501 - VMware_bootbank_gc_8.0.3-0.55.24414501 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.55.24414501 - VMware_bootbank_bmcad_8.0.3-0.55.24414501 - VMware_bootbank_esxio-combiner_8.0.3-0.55.24414501 - VMware_bootbank_crx_8.0.3-0.55.24414501 - VMware_bootbank_vsanhealth_8.0.3-0.55.24414501 - VMware_bootbank_cpu-microcode_8.0.3-0.55.24414501 - VMware_bootbank_esxio_8.0.3-0.55.24414501 - VMware_bootbank_esx-update_8.0.3-0.55.24414501 - VMware_bootbank_loadesx_8.0.3-0.55.24414501 - VMware_bootbank_esxio-update_8.0.3-0.55.24414501 - VMware_bootbank_loadesxio_8.0.3-0.55.24414501
PRs Fixed	3459675
Related CVE numbers	N/A

This patch updates the following issue:

vSphere vMotion tasks fail with an error NamespaceMgr could not lock the db file

In rare cases, when virtual machines are configured with namespaces, an issue with the namespace database might cause migration of such VMs with vSphere vMotion to fail. In the vSphere Client, you see errors such as:

Failed to receive migration. The source detected that the destination failed to resume.

An error occurred restoring the virtual machine state during migration. NamespaceMgr could not lock the db file.

The destination ESXi host `vmware.log` shows the following messages:

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.checkpoint.migration.failedReceive] Failed to receive migration.
```

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.namespaceMgr.noLock] NamespaceMgr could not lock the db file.
```

The issue is more likely to occur on VMs with hardware version earlier than 19, and impacts only VMs provisioned on shared datastores other than NFS, such as VMFS, vSAN, and vSphere Virtual Volumes.

This issue was reported as known in KB [369767](#) and is resolved in this release.

ESXi-8.0U3c-24414501-no-tools

Profile Name	ESXi-8.0U3c-24414501-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	December 12, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-base_8.0.3-0.55.24414501 - VMware_bootbank_vcls-pod-crx_8.0.3-0.55.24414501 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.55.24414501 - VMware_bootbank_gc-esxio_8.0.3-0.55.24414501 - VMware_bootbank_drivervm-gpu-base_8.0.3-0.55.24414501 - VMware_bootbank_trx_8.0.3-0.55.24414501 - VMware_bootbank_clusterstore_8.0.3-0.55.24414501 - VMware_bootbank_esx-xserver_8.0.3-0.55.24414501 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.55.24414501 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.55.24414501 - VMware_bootbank_infravisor_8.0.3-0.55.24414501 - VMware_bootbank_native-misc-drivers_8.0.3-0.55.24414501 - VMware_bootbank_bmcad-esxio_8.0.3-0.55.24414501 - VMware_bootbank_vdfs_8.0.3-0.55.24414501 - VMware_bootbank_esx-base_8.0.3-0.55.24414501 - VMware_bootbank_vsan_8.0.3-0.55.24414501 - VMware_bootbank_vds-vsip_8.0.3-0.55.24414501 - VMware_bootbank_gc_8.0.3-0.55.24414501

	- VMware_bootbank_esxio-combiner-esxio_8.0.3-0.55.24414501 - VMware_bootbank_bmcad_8.0.3-0.55.24414501 - VMware_bootbank_esxio-combiner_8.0.3-0.55.24414501 - VMware_bootbank_crx_8.0.3-0.55.24414501 - VMware_bootbank_vsanhealth_8.0.3-0.55.24414501 - VMware_bootbank_cpu-microcode_8.0.3-0.55.24414501 - VMware_bootbank_esxio_8.0.3-0.55.24414501 - VMware_bootbank_esx-update_8.0.3-0.55.24414501 - VMware_bootbank_loadesx_8.0.3-0.55.24414501 - VMware_bootbank_esxio-update_8.0.3-0.55.24414501 - VMware_bootbank_loadesxio_8.0.3-0.55.24414501
PRs Fixed	3459675
Related CVE numbers	N/A

This patch updates the following issue:

vSphere vMotion tasks fail with an error NamespaceMgr could not lock the db file

In rare cases, when virtual machines are configured with namespaces, an issue with the namespace database might cause migration of such VMs with vSphere vMotion to fail. In the vSphere Client, you see errors such as:

Failed to receive migration. The source detected that the destination failed to resume.

An error occurred restoring the virtual machine state during migration. NamespaceMgr could not lock the db file.

The destination ESXi host `vmware.log` shows the following messages:

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.checkpoint.migration.failedReceive] Failed to receive migration.
```

```
[YYYY-MM-DDTHH:MM:SS] In(05) vmx - [msg.namespaceMgr.noLock] NamespaceMgr could not lock the db file.
```

The issue is more likely to occur on VMs with hardware version earlier than 19, and impacts only VMs provisioned on shared datastores other than NFS, such as VMFS, vSAN, and vSphere Virtual Volumes.

This issue was reported as known in KB [369767](#) and is resolved in this release.

ESXi-8.0U3c-24414501

Name	ESXi
Version	ESXi-8.0U3c-24414501
Release Date	December 12, 2024
Category	Bugfix
Affected Components	- ESXi Component - core ESXi VIBs - ESXi Install/Upgrade Component
PRs Fixed	3459675
Related CVE numbers	N/A

This patch resolves the issues listed in [ESXi_8.0.3-0.55.24414501](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

You cannot update to ESXi 8.0 Update 2b by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see [Upgrade or Update a Host with Image Profiles](#).

The Cancel option in an interactive ESXi installation might not work as expected

Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

The `irdman` driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic

If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the `irdman` driver might fail. This issue is unlikely to occur, as the `irdman` driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

Workaround: Use ULPs with RC transport mode.

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

~# reboot

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1

The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

Capture of network packets by using the PacketCapture tool on ESXi does not work

Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in [Collecting network packets using the lightweight PacketCapture on ESXi](#).

Workaround: Use the pktcap-uw tool. For more information, see [Capture and Trace Network Packets by Using the pktcap-uw Utility](#).

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out

vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD

In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See [Reconciling Discrepancies in the Managed Virtual Disk Catalog](#).

vCenter Server and vSphere Client Issues

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

VMware ESXi 8.0 Update 3b Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Product Support Notices](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 3b | 17 SEP 2024 | ISO Build 24280767
Check for additions and updates to these release notes.

What's New

This [release](#) provides bug and security fixes. For more information, see the [Resolved Issues](#) section.

DPU/SmartNIC

VMware vSphere Distributed Services Engine support with NVIDIA Bluefield-3 DPUs: Starting with ESXi 8.0 Update 3b, vSphere Distributed Services Engine adds support for NVIDIA Bluefield-3 DPUs. NVIDIA Bluefield-3 DPUs are supported in both single-DPU and dual-DPU configurations.

ESXi 8.0 Update 3b adds support to vSphere Quick Boot for multiple servers, including:

- HPE
 - ProLiant DL145 Gen11
 - ProLiant MicroServer Gen11

- Cisco Systems Inc.
 - UCSC-C245-M8SX
- Supermicro
 - AS-2025HS-TNR

For the full list of supported servers, see the [VMware Compatibility Guide](#).

Design improvements for cloud-init guestInfo variables: Starting with ESXi 8.0 Update 3b, setting `cloud-init` guestInfo variables by regular users in the guest results into an error. For more information, see KB [377267](#).

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0 Update 3 Release Notes](#)
- [VMware ESXi 8.0 Update 2c Release Notes](#)
- [VMware ESXi 8.0 Update 1d Release Notes](#)
- [VMware ESXi 8.0 Update 2b Release Notes](#)
- [VMware ESXi 8.0 Update 2 Release Notes](#)
- [VMware ESXi 8.0 Update 1c Release Notes](#)
- [VMware ESXi 8.0 Update 1a Release Notes](#)
- [VMware ESXi 8.0 Update 1 Release Notes](#)
- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Product Support Notices

For updates to the Product Support Notices section, see the [ESXi 8.0 Update 3](#) release notes.

Patches Contained in This Release

VMware ESXi 8.0 Update 3b

Build Details

VMware vSphere Hypervisor (ESXi ISO) image

Download Filename:	VMware-VMvisor-Installer-8.0U3b-24280767.x86_64.iso
Build:	24280767
Download Size:	607 MB
SHA256 checksum:	718c4fa3447dd138db64398604438ee0d700e2c2fae9276201c016e386c3812b
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

VMware vSphere Hypervisor (ESXi) Offline Bundle

Download Filename:	VMware-ESXi-8.0U3b-24280767-depot.zip
Build:	24280767
Download Size:	994.7 MB
SHA256 checksum:	1824dc62fb36e7e107bcf5526278436986287597c0c4f450646d8d524118782f
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.3-0.35.24280767	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.35.24280767	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.35.24280767	Bugfix	Critical
VMware BlueField RShim Driver	VMware-rshim_0.1-12vmw.803.0.35.24280767	Bugfix	Critical
Networking Device Driver for NVIDIA BlueField RShim devices	VMware-rshim-net_0.1.0-1vmw.803.0.35.24280767	Bugfix	Critical
NVIDIA BlueField boot control driver	VMware-mlnx-bfbootctl_0.1-7vmw.803.0.35.24280767	Bugfix	Critical
VMware NVMe over TCP Driver	VMware-NVMeoF-TCP_1.0.1.29-1vmw.803.0.35.24280767	Bugfix	Critical
ESXi Component - core ESXi VIBs	ESXi_8.0.3-0.30.24262298	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.3-0.30.24262298	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.3-0.30.24262298	Security	Critical
ESXi Tools Component	VMware-VM-Tools_12.4.5.23787635-24262298	Security	Critical

Rollup Bulletins

These rollup bulletins contain the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U3b-24280767	Bugfix	Critical	Security and Bugfix image
ESXi80U3sb-24262298	Security	Critical	Security only image

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name

ESXi-8.0U3b-24280767-standard
ESXi-8.0U3b-24280767-no-tools
ESXi-8.0U3sb-24262298-standard
ESXi-8.0U3sb-24262298-no-tools

ESXi Images

Name and Version	Release Date	Category	Detail
ESXi_8.0.3-0.35.24280767	09/17/2024	General	Security and Bugfix image
ESXi_8.0.3-0.30.24262298	09/17/2024	Security	Security only image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.3-0.35.24280767

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-base_8.0.3-0.35.24280767 - VMware_bootbank_esxio-base_8.0.3-0.35.24280767 - VMware_bootbank_bmcad_8.0.3-0.35.24280767 - VMware_bootbank_gc_8.0.3-0.35.24280767 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.35.24280767 - VMware_bootbank_gc-esxio_8.0.3-0.35.24280767 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.35.24280767 - VMware_bootbank_vdfs_8.0.3-0.35.24280767 - VMware_bootbank_cpu-microcode_8.0.3-0.35.24280767 - VMware_bootbank_vsan_8.0.3-0.35.24280767 - VMware_bootbank_trx_8.0.3-0.35.24280767 - VMware_bootbank_esx-xserver_8.0.3-0.35.24280767 - VMware_bootbank_clusterstore_8.0.3-0.35.24280767 - VMware_bootbank_esxio-combiner_8.0.3-0.35.24280767

	• VMware_bootbank_vcls-pod-crx_8.0.3-0.35.24280767 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.35.24280767 • VMware_bootbank_vds-vsip_8.0.3-0.35.24280767 • VMware_bootbank_vsanhealth_8.0.3-0.35.24280767 • VMware_bootbank_native-misc-drivers_8.0.3-0.35.24280767 • VMware_bootbank_crx_8.0.3-0.35.24280767 • VMware_bootbank_esxio_8.0.3-0.35.24280767 • VMware_bootbank_infravisor_8.0.3-0.35.24280767 • VMware_bootbank_drivervm-gpu-base_8.0.3-0.35.24280767 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.35.24280767 • VMware_bootbank_bmcal-esxio_8.0.3-0.35.24280767 • VMW_bootbank_pensandoatlas_1.46.0.E.41.1.326-2vmw.803.0.0.0.23797590
PRs Fixed	3426946, 3421084, 3421179, 3403683, 3388844, 3424051, 3422735, 3420488, 3421971, 3422005, 3417326, 3420850, 3420421, 3420586, 3410311, 3420907, 3419241, 3409528, 3404817, 3416221, 3421665, 3415365, 3415102, 3417329, 3421434, 3419074, 3417667, 3408802, 3417224, 3420702, 3385757, 3412010, 3418878, 3408477, 3406546, 3406968, 3406999, 3398549, 3402823, 3408281, 3409108, 3406627, 3407251, 3392225, 3412138, 3406037, 3406875, 3412536, 3407952, 3410114, 3386751, 3408145, 3403240, 3389766, 3283598, 3408302, 3403639, 3407532, 3394043, 3387721, 3408300, 3396479, 3401629, 3401392, 3409085, 3408640, 3405912, 3403680, 3405106, 3407204, 3407022, 3409176, 3408739, 3408740, 3390434, 3406053, 3400702, 3392173, 3403496, 3402955, 3433092
CVE numbers	N/A

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

This patch updates the esx-base, esxio-base, bmcal, gc, esxio-dvfilter-generic-fastpath, gc-esxio, native-misc-drivers-esxio, vdfs, cpu-microcode, vsan, trx, esx-xserver, clusterstore, esxio-combiner, vcls-pod-crx, esxio-combiner-esxio, vds-vsip, vsanhealth, native-misc-drivers, crx, esxio, infravisor, drivervm-gpu-base, esx-dvfilter-generic-fastpath, pensandoatlas, and bmcal-esxio VIBs.

This patch resolves the following issues:

PR 3392225: You see status Unknown for virtual machines after migration to an ESXi host with active host-local swap location

In rare occasions with very specific sequence of conditions, when you migrate to or power on a virtual machine on an ESXi host with active host-local swap location, the virtual machine might hang with a status `Unknown`.

This issue is resolved in this release.

PR 3421434: When you migrate virtual machines with snapshots from a vSAN ESA 8.x datastore, you might see errors at the destination datastore

In vSAN ESA 8.x environments, under certain conditions, migrating VMs with snapshots from a vSAN ESA datastore by using either a Storage vMotion, cross vMotion, cold relocate or clone operation might result in errors at the destination datastore, such as VMs failing to boot up. The issue is specific to vSAN ESA and is not applicable to vSAN OSA. It can affect both the VM snapshots and the running VM.

This issue is resolved in this release.

PR 3388844: You cannot activate Kernel Direct Memory Access (DMA) Protection for Windows guest OS on ESXi hosts with Intel CPU

If an input–output memory management unit (IOMMU) is active on a Windows guest OS, the **Kernel DMA Protection** option under **System Information** might be off for VMs running on ESXi hosts with Intel CPUs. As a result, you might not be able to fulfill some security requirements for your environment.

This issue is resolved in this release. The fix deactivates **Kernel DMA Protection** by default. ESXi 8.0 Update 3b adds the vmx parameter `acpi.dmar.enableDMAProtection` and to activate **Kernel DMA Protection** in a Windows guest OS, you must add `acpi.dmar.enableDMAProtection=TRUE` to the vmx file.

PR 3377863: You see "Hosts are remediated" message in the upgrade precheck results

When running a precheck before an ESXi update, you see a message such as `Hosts are remediated` in the precheck results, which is not clear and might be misleading.

This issue is resolved in this release. The new message is `Hosts are remediated sequentially`.

PR 3417329: Virtual machine tasks might intermittently fail due to a rare issue with the memory slab

Due to a rare issue with the vSAN DOM object, where a reference count on a component object might not decrement correctly, the in-memory object might never be released from the slab and can cause the component manager slab to reach its limit. As a result, you might not be able to create VMs, migrate VMs or might encounter VM power-on failures on vSAN clusters, either OSA or ESA.

This issue is resolved in this release.

PR 3406140: Extracting a vSphere Lifecycle Manager image from an existing ESXi host might fail after a kernel configuration change

Each update of the kernel configuration also triggers an update to the `/bootbank/useropts.gz` file, but due to a known issue, the `basemisc.tgz` might not contain the default `useropts` file after such an update. As a result, when attempting to extract a vSphere Lifecycle Manager image from an existing ESXi host, the absence of the default `useropts` file leads to failure to create the `esx-base.vib` file and the operation also fails.

This issue is resolved in this release.

PR 3408477: Some ESXi hosts might not have a locker directory after an upgrade from ESXi 6.x to 8.x

When you upgrade an ESXi host with a boot disk of less than 10 GB and not on USB from ESXi 6.x to 8.x, the upgrade process might not create a locker directory and the `/locker` symbolic link is not active.

This issue is resolved in this release. If you already face the issue, upgrading to ESXi 8.0 Update 3b creates a locker directory but does not automatically create a VMware Tools repository. As a result, clusters that host such ESXi hosts display as non-compliant. Remediate the cluster again to create a VMware Tools repository and to become compliant.

PR 3422005: ESXi hosts of version 8.0 Update 2 and later might fail to synchronize time with certain NTP servers

A change in the `ntp-4.2.8p17` package in ESXi 8.0 Update 2 might cause the NTP client to reject certain server packets as poorly formatted or invalid. For example, if a server sends packets with a `poll` value of 0, the NTP client on ESXi does not synchronize with the server.

This issue is resolved in this release.

PR 3408145: The vmx service might fail with a core dump due to rare issue with the vSphere Data Protection solution running out of resources

In rare cases, if a backup operation starts during high concurrent guest I/O load, for example a VM with high write I/O intensity and a high number of overwrites, the VAIO filter component of the vSphere Data Protection solution might run out of resources to handle the guest write I/Os. As a result, the vmx service might fail with a core dump and restart.

This issue is resolved in this release. Alternatively, run backups during periods of low guest I/O activity.

PR 3405912: In the vSphere Client, you do not see the correct total vSAN storage consumption for which you have a license

Due to a rare race condition in environments with many vSAN clusters managed by a single vCenter instance, in the vSphere Client under **Licensing > Licenses** you might see a discrepancy between the total claimed vSAN storage capacity and the reported value for the clusters.

This issue is resolved in this release.

PR 3414588: Snapshot tasks on virtual machines on NVMe/TCP datastores might take much longer than on VMs provisioned on NVMe/FC datastores

When creating or deleting a snapshot on a VM provisioned on datastores backed by NVMe/TCP namespaces, such tasks might take much longer than on VMs provisioned on datastores backed by NVMe/FC namespaces. The issue occurs because the nvmetcp driver handles some specific NVMe commands not in the way the NVMe/TCP target systems expect.

This issue is resolved in this release.

PR 3392173: A rare issue with the Virsto vSAN component might cause failure to create vSAN objects, or unmount disk groups, or reboot ESXi hosts

In very rare cases, if a Virsto component creation task fails, it might not be properly handled and cause background deletion of virtual disks to stop. As a result, deleting virtual disks in tasks such as creating vSAN objects, or unmounting of disk groups, or rebooting ESXi hosts does not occur as expected and might cause such tasks to fail.

This issue is resolved in this release.

PR 3415365: ESXi upgrade to 8.0 Update 3 fails with an error in the vFAT bootbank partitions

ESXi 8.0 Update 3 adds a precheck in the upgrade workflow that uses the `dosfsck` tool to catch vFAT corruptions. One of the errors that `dosfsck` flags is the `dirty bit set` but ESXi does not use that concept and such errors are false positive.

In the vSphere Client, you see an error such as `A problem with one or more vFAT bootbank partitions was detected`. Please refer to KB 91136 and run `dosfsck` on bootbank partitions.

In the remediation logs on ESXi hosts, you see logs such as:

```
2024-07-02T16:01:16Z In(14) lifecycle[122416262]: runcommand:199 runcommand called with:
args = ['/bin/dosfsck', '-V', '-n', '/dev/disks/naa.600508b1001c7d25f5336a7220b5afc1:6'],
outfile = None, returnoutput= True, timeout = 10.
```

```
2024-07-02T16:01:16Z In(14) lifecycle[122416262]: upgrade_precheck:1836 dosfsck
output: b'CP850//TRANSLIT: Invalid argument\nCP850: Invalid argument\nfsck.fat
4.1+git (2017-01-24)\n0x25: Dirty bit is set. Fswas not properly unmounted and some
data may be corrupt.\n Automatically removing dirty bit.\nStarting check/repair
pass.\nStarting verification pass.\n\nLeaving filesystem unchanged.\n/dev/disks/
naa.600508b1001c7d25f5336a7220b5afc1:6: 121 files, 5665/65515 clusters\n'
```

This issue is resolved in this release.

PR 3421084: An ESXi host might become temporarily inaccessible from vCenter if an NFSv3 datastore fails to mount during reboot or bring up

During an ESXi host reboot, if an NFSv3 datastore fails to mount during the reboot or bring up in VMware Cloud Foundation environments, retries to mount the datastore continue in the background. However, while the datastore is still not available, the `hostd` daemon might fail with a core dump when trying to access it and cause the host to lose connectivity to the vCenter system for a short period.

This issue is resolved in this release.

PR 3407251: ESXi host fails with a purple diagnostic screen due to a rare physical CPU (PCPU) lockup

In the vSphere Client, when you use the **Delete from Disk** option to remove a virtual machine from a vCenter system and delete all VM files from the datastore, including the configuration file and virtual disk files, if any of the files is corrupted, a rare issue with handling corrupted files in the delete path might lead to a PCPU lockup. As a result, the ESXi host fails with a purple diagnostic screen and a message such as `NMI IPI: Panic requested by another PCPU`.

This issue is resolved in this release.

PR 3406627: VMFS6 automatic UNMAP feature might fail to reclaim filesystem space beyond 250GB

In certain cases, when you delete a filesystem space of more than 250GB on a VMFS6 volume, for example 1 TB, if the volume has no active references such as active VMs, then the VMFS6 automatic UNMAP feature might fail to reclaim the filesystem space beyond 250GB.

This issue is resolved in this release.

PR 3419241: After deleting a snapshot or snapshot consolidation, some virtual machines intermittently fail

When a VM on a NFSv3 datastore has multiple snapshots, such as s1, s2, s3, s4, if the VM reverts to one of the snapshots, for example s2, then powers on, and then one of the other snapshots, such as s3, is deleted, the vmx service might fail. The issue occurs because the code tries to consolidate links of a disk that is not part of the VM current state and gets a null pointer. As a result, snapshot consolidation might also fail and cause the vmx service to fail as well.

This issue is resolved in this release. If you already face the issue, power off the VM, edit its `.vmx` file to add the following setting: `consolidate.upgradeNFS3Locks = "FALSE"`, and power on the VM.

PR 3407532: VMs with snapshots and active encryption experience higher I/O latency

In rare cases, encrypted VMs with snapshots might experience higher than expected latency. This issue occurs due to unaligned I/O operations that generate excessive metadata requests to the underlying storage and lead to increased latency.

This issue is resolved in this release. To optimize performance, the VMcrypt I/O filter is enhanced to allocate memory in 4K-aligned blocks for both read and write operations. This reduction in metadata overhead significantly improves overall I/O performance.

PR 3410311: You cannot log in to the Direct Console User Interface (DCUI) with regular Active Directory credentials

When you try to log with a regular Active Directory account in to the DCUI of an ESXi host by using either a remote management application such as HP Integrated Lights-Out (iLO) or Dell Remote Access Card (DRAC), or a server management system such as Lenovo XCC or Huawei iBMC, login might fail. In the DCUI, you see an error such as `Wrong user name or password`. In the `vmkernel.log` file, you see logs such as:

```
2024-07-01T10:40:53.007Z In(182) vmkernel: cpu1:264954)VmkAccess: 106: dcui: running in
dcuiDom(7): socket = /etc/likewise/lib/.lsassd (unix_stream_socket_connect): Access denied
by vmkernel access control policy.
```

The issue occurs due to a restriction of ESXi processes to access certain resources, such as Likewise.

This issue is resolved in this release.

PR 3403706: Hot extending a non-shared disk in a Windows Server Failover Cluster might result in lost reservations on shared disks

In a WSFC cluster, due to an issue with releasing SCSI reservations, in some cases hot extending a non-shared disk might result in lost reservations on shared disks and failover of the disk resource.

This issue is resolved in this release. The fix makes sure releasing SCSI reservations are properly handled for all type of shared disks.

PR 3394043: Creating a vSAN File Service fails when you use IPs within the 172.17.0.0/16 range as mount points

Prior to vSphere 8.0 Update 3b, you must change your network configuration in cases when the specified file service network overlaps with the Docker default internal network 172.17.0.0/16. As a result, you see Skyline Health warnings for DNS lookup failures and you cannot create vSAN File Services.

This issue is resolved in this release. The fix routes traffic to the correct endpoint to avoid possible conflicts.

PR 3389766: During a vSphere vMotion migration of a fault tolerant Primary VM with encryption, the migration task might fail and vSphere FT failover occurs

In rare cases, the encryption key package might not be sent or received correctly during a vSphere vMotion migration of a fault tolerant Primary VM with encryption, and as a result the migration task fails and vSphere FT failover occurs.

This issue is resolved in this release. The fix sends the encryption key at the end of the vSphere FT checkpoint to avoid errors.

PR 3403680: Mounting of vSphere Virtual Volumes stretched storage container fails with an undeclared fault

In the vSphere Client, you might see the error `Undeclared fault` while mounting a newly created vSphere Virtual Volumes stretched storage container from an existing storage array. The issue occurs due to a rare race condition. vSphere Virtual Volumes generates a core dump and restarts after the failure.

This issue is resolved in this release.

PR 3408300: You cannot remove or delete a VMFS partition on a 4K native (4Kn) Software Emulation (SWE) disk

When you attempt to remove or delete a VMFS partition on a 4Kn SWE disk, in the vSphere Client you see an error such as `Read-only file system during write on /dev/disks/<device name>` and the operation fails. In the `vmkernel.log`, you see entries such as `in-use partition <part num>, modification is not supported`.

This issue is resolved in this release.

PR 3402823: Fresh installation or creating VMFS partitions on a Micron 7500 or Intel D5-P5336 NVMe drives might fail with a purple diagnostic screen

UNMAP commands enable ESXi hosts to release storage space that is mapped to data deleted from the host. In NVMe, the equivalent of UNMAP commands is a deallocate DSM request. Micron 7500 and Intel D5-P5336 devices advertise a very large value in one of the deallocate limit attributes, DMSRL, which is the maximum number of logical blocks in a single range for a Dataset Management command. This leads to an integer overflow when the ESXi unmap split code converts number of blocks to number of bytes, which in turn might cause a failure of either installation or VMFS creation. You see a purple diagnostics screen with an error such as `Exception 14 or corruption in dlmalloc`. The issue affects ESXi 8.0 Update 2 and later.

This issue is resolved in this release.

PR 3396479: Standard image profiles for ESXi 8.0 Update 3 show last modified date as release date

The **Release Date** field of the standard image profile for ESXi 8.0 Update 3 shows the **Last Modified Date** value. The issue is only applicable to the image profiles used in Auto Deploy or ESXCLI. Base images used in vSphere Lifecycle Manager workflows display the release date correctly. This issue has no functional impact. The side effect is that if you search for profiles by release date, the profile does not show with the actual release date.

This issue is resolved in this release.

esx-update_8.0.3-0.35.24280767

Patch Category	Bugfix
Patch Severity	Critical

Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesx_8.0.3-0.35.24280767 - VMware_bootbank_esx-update_8.0.3-0.35.224280767
PRs Fixed	3412448, 3389057, 3397914, 3421179, 3426946
CVE numbers	N/A

Updates the `esx-update` and `loadesx` VIBs to resolve the following issue:

PR 3412448: Remediation of ESXi hosts might fail due to a timeout in the update of the VMware-VM-Tools component

While upgrading ESXi hosts by using vSphere Lifecycle Manager, remediation of the VMware-VM-Tools component might occasionally fail due to a timeout. The issue occurs when the update process takes longer than the timeout setting of 30 sec for the VMware-VM-Tools component.

This issue is resolved in this release. The fix increases the timeout limit to 120 sec.

esxio-update_8.0.3-0.35.24280767

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesxio_8.0.3-0.35.24280767 - VMware_bootbank_esxio-update_8.0.3-0.35.24280767
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

VMware-rshim_0.1-12vmw.803.0.35.24280767

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMW_bootbank_rshim_0.1-12vmw.803.0.35.24280767
PRs Fixed	3421179, 3426946
CVE numbers	N/A

Updates the `rshim` VIB.

VMware-rshim-net_0.1.0-1vmw.803.0.35.24280767

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMW_bootbank_rshim-net_0.1.0-1vmw.803.0.35.24280767
PRs Fixed	3421179
CVE numbers	N/A

Updates the `rshim-net` VIB.

VMware-mlnx-bfbootctl_0.1-7vmw.803.0.35.24280767

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMW_bootbank_mlnx-bfbootctl-esxio_0.1-7vmw.803.0.35.24280767
PRs Fixed	3421179
CVE numbers	N/A

Updates the `mlnx-bfbootctl-esxio` VIB.

VMware-NVMeoF-TCP_1.0.1.29-1vmw.803.0.35.24280767

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMW_bootbank_nvmetcp_1.0.1.29-1vmw.803.0.35.24280767 - VMW_bootbank_nvmetcp-esxio_1.0.1.29-1vmw.803.0.35.24280767
PRs Fixed	3414588
CVE numbers	N/A

Updates the `nvmetcp` and `nvmetcp-esxio` VIBs.

ESXi_8.0.3-0.30.24262298

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_drivervm-gpu-base_8.0.3-0.30.24262298 - VMware_bootbank_esxio-combiner_8.0.3-0.30.24262298 - VMware_bootbank_esx-xserver_8.0.3-0.30.24262298 - VMware_bootbank_vdfs_8.0.3-0.30.24262298 - VMware_bootbank_cpu-microcode_8.0.3-0.30.24262298 - VMware_bootbank_vds-vsip_8.0.3-0.30.24262298 - VMware_bootbank_vsan_8.0.3-0.30.24262298 - VMware_bootbank_bmcad_8.0.3-0.30.24262298 - VMware_bootbank_clusterstore_8.0.3-0.30.24262298 - VMware_bootbank_native-misc-drivers_8.0.3-0.30.24262298 - VMware_bootbank_infravisor_8.0.3-0.30.24262298 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.30.24262298 - VMware_bootbank_crx_8.0.3-0.30.24262298 - VMware_bootbank_gc_8.0.3-0.30.24262298 - VMware_bootbank_vcls-pod-crx_8.0.3-0.30.24262298 - VMware_bootbank_esxio-base_8.0.3-0.30.24262298 - VMware_bootbank_esx-base_8.0.3-0.30.24262298 - VMware_bootbank_esxio_8.0.3-0.30.24262298 - VMware_bootbank_gc-esxio_8.0.3-0.30.24262298 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.30.24262298 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.30.24262298 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.30.24262298 - VMware_bootbank_vsanhealth_8.0.3-0.30.24262298 - VMware_bootbank_trx_8.0.3-0.30.24262298 - VMware_bootbank_bmcad-esxio_8.0.3-0.30.24262298 - VMW_bootbank_pensandoatlas_1.46.0.E.41.1.326-2vmw.803.0.0.0.23797590
PRs Fixed	3425039, 3423080, 3408352, 3415908, 3380359, 3404366, 3421357, 3421359, 3390663, 3404362, 3396479, 3404367, 3398130, 3398132, 3316536, 3410159, 3412279, 3395162, 3404369, 3404657, 3390662
CVE numbers	N/A

The `ESXi` and `esx-update` bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

This patch updates the `drivervm-gpu-base`, `esxio-combiner`, `esx-xserver`, `vdfs`, `cpu-microcode`, `vds-vsip`, `vsan`, `bmc`, `clusterstore`, `native-misc-drivers`, `infravisor`, `esxio-dvfilter-generic-fastpath`, `crx`, `gc`, `vcls-pod-crx`, `esxio-base`, `esx-base`, `esxio`, `gc-esxio`, `native-misc-drivers-esxio`, `esxio-combiner-esxio`, `esx-dvfilter-generic-fastpath`, `vsanhealth`, `trx`, `pensandoatlas`, and `bmc`-`esxio` VIBs.

This patch resolves the following issues:

ESXi 8.0 Update 3b includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007206	4/15/2024	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003707	3/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003707	3/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Cooper Lake	0x5065b (06/55/b)	0xbf	baseline	0x7002904	4/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd0003e7	4/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002b0	4/3/2024	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xb8	2/15/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x38	2/15/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x52	2/15/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b0005c0	2/5/2024	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x100	2/5/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x62	3/7/2024	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x125	4/16/2024	Intel Xeon E-2400 Series
Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x21000200	11/20/2023	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series

ESX 8.0 Update 3b contains the following security updates:

- OpenSSH is updated to version 9.8p1.

esx-update_8.0.3-0.30.24262298

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	• VMware_bootbank_esx-update_8.0.3-0.30.24262298 • VMware_bootbank_loadesx_8.0.3-0.30.24262298
PRs Fixed	N/A
CVE numbers	N/A

Updates the `esx-update` and `loadesx` VIBs.

esxio-update_8.0.3-0.30.24262298

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esxio-update_8.0.3-0.30.24262298 - VMware_bootbank_loadesxio_8.0.3-0.30.24262298
PRs Fixed	N/A
CVE numbers	N/A

Updates the `esxio-update` and `loadesxio` VIBs.

VMware-VM-Tools_12.4.5.23787635-24262298

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	No
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMware_locker_tools-light_12.4.5.23787635-24262298
PRs Fixed	3409328
CVE numbers	N/A

Updates the `tools-light` VIB.

VMware Tools Bundling Changes in ESXi 8.0 Update 3b

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 2b:

- **windows.iso**: VMware Tools 12.4.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.4.5 Release Notes](#)

- [Earlier versions of VMware Tools](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Updating VMware Tools](#)

ESXi-8.0U3b-24280767-standard

Profile Name	ESXi-8.0U3b-24280767-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	September 17, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	• VMware_bootbank_esxio-base_8.0.3-0.35.24280767 • VMware_bootbank_bmcad_8.0.3-0.35.24280767 • VMware_bootbank_gc_8.0.3-0.35.24280767 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.35.24280767 • VMware_bootbank_gc-esxio_8.0.3-0.35.24280767 • VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.35.24280767 • VMware_bootbank_vdfs_8.0.3-0.35.24280767 • VMware_bootbank_cpu-microcode_8.0.3-0.35.24280767 • VMware_bootbank_vsan_8.0.3-0.35.24280767 • VMware_bootbank_trx_8.0.3-0.35.24280767 • VMware_bootbank_esx-xserver_8.0.3-0.35.24280767 • VMware_bootbank_clusterstore_8.0.3-0.35.24280767 • VMware_bootbank_esxio-combiner_8.0.3-0.35.24280767 • VMware_bootbank_vcls-pod-crx_8.0.3-0.35.24280767 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.35.24280767 • VMware_bootbank_vds-vsip_8.0.3-0.35.24280767 • VMware_bootbank_vsanhealth_8.0.3-0.35.24280767 • VMware_bootbank_native-misc-drivers_8.0.3-0.35.24280767 • VMware_bootbank_crx_8.0.3-0.35.24280767 • VMware_bootbank_esxio_8.0.3-0.35.24280767 • VMware_bootbank_infravisor_8.0.3-0.35.24280767 • VMware_bootbank_drivervm-gpu-base_8.0.3-0.35.24280767 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.35.24280767 • VMware_bootbank_bmcad-esxio_8.0.3-0.35.24280767 • VMware_bootbank_esx-base_8.0.3-0.35.24280767 • VMware_bootbank_loadesx_8.0.3-0.35.24280767 • VMware_bootbank_esx-update_8.0.3-0.35.24280767 • VMware_bootbank_loadesxio_8.0.3-0.35.24280767 • VMware_bootbank_esxio-update_8.0.3-0.35.24280767 • VMW_bootbank_rshim_0.1-12vmw.803.0.35.24280767 • VMW_bootbank_rshim-net_0.1.0-1vmw.803.0.35.24280767

	• VMW_bootbank_mlnx-bfbootctl-esxio_0.1-7vmw.803.0.35.242.80767 • VMW_bootbank_pensandoatlas_1.46.0.E.41.1.326-2vmw.803.0.0.0.23797590
PRs Fixed	3426946, 3421084, 3421179, 3403683, 3388844, 3424051, 3422735, 3420488, 3421971, 3422005, 3417326, 3420850, 3420421, 3420586, 3410311, 3420907, 3419241, 3409528, 3404817, 3416221, 3421665, 3415365, 3415102, 3417329, 3421434, 3419074, 3417667, 3408802, 3417224, 3420702, 3385757, 3412010, 3418878, 3408477, 3406546, 3406968, 3406999, 3398549, 3402823, 3408281, 3409108, 3406627, 3407251, 3392225, 3412138, 3406037, 3406875, 3412536, 3407952, 3410114, 3386751, 3408145, 3403240, 3389766, 3283598, 3408302, 3403639, 3407532, 3394043, 3387721, 3408300, 3396479, 3401629, 3401392, 3409085, 3408640, 3405912, 3403680, 3405106, 3407204, 3407022, 3409176, 3408739, 3408740, 3390434, 3406053, 3400702, 3392173, 3403496, 3402955, 3433092, 3421179, 3414588
Related CVE numbers	N/A

This patch updates the following issues:

PR 3392225: You see status Unknown for virtual machines after migration to an ESXi host with active host-local swap location

In rare occasions with very specific sequence of conditions, when you migrate to or power on a virtual machine on an ESXi host with active host-local swap location, the virtual machine might hang with a status `Unknown`.

This issue is resolved in this release.

PR 3388844: You cannot activate Kernel Direct Memory Access (DMA) Protection for Windows guest OS on ESXi hosts with Intel CPU

If an input–output memory management unit (IOMMU) is active on a Windows guest OS, the **Kernel DMA Protection** option under **System Information** might be off for VMs running on ESXi hosts with Intel CPUs. As a result, you might not be able to fulfill some security requirements for your environment.

This issue is resolved in this release. The fix activates **Kernel DMA Protection** by default. The fix also adds the `vmx` parameter `acpi.dmar.enabledDMAProtection` to activate the **Kernel DMA Protection** in Windows OS. The default value of the parameter is `FALSE` and to activate **Kernel DMA Protection** in a Windows guest OS, you must add `acpi.dmar.enabledDMAProtection=TRUE` to the `vmx` file.

PR 3412448: Remediation of ESXi hosts might fail due to a timeout in the update of the VMware-VM-Tools component

While upgrading ESXi hosts by using vSphere Lifecycle Manager, remediation of the VMware-VM-Tools component might occasionally fail due to a timeout. The issue occurs when the update process takes longer than the timeout setting of 30 sec for the VMware-VM-Tools component.

This issue is resolved in this release. The fix increases the timeout limit to 120 sec.

PR 3377863: You see "Hosts are remediated" message in the upgrade precheck results

When running a precheck before an ESXi update, you see a message such as `Hosts are remediated` in the precheck results, which is not clear and might be misleading.

This issue is resolved in this release. The new message is `Hosts are remediated sequentially`.

PR 3417329: Virtual machine tasks might intermittently fail due to a rare issue with the memory slab

Due to a rare issue with the vSAN DOM object, where a reference count on a component object might not decrement correctly, the in-memory object might never be released from the slab and can cause the component manager slab to reach its limit. As a result, you might not be able to create VMs, migrate VMs or might encounter VM power-on failures on vSAN clusters, either OSA or ESA.

This issue is resolved in this release.

PR 3406140: Extracting a vSphere Lifecycle Manager image from an existing ESXi host might fail after a kernel configuration change

Each update of the kernel configuration also triggers an update to the `/bootbank/useropts.gz` file, but due to a known issue, the `basemisc.tgz` might not contain the default `useropts` file after such an update. As a result, when attempting to extract a vSphere Lifecycle Manager image from an existing ESXi host, the absence of the default `useropts` file leads to failure to create the `esx-base.vib` file and the operation also fails.

This issue is resolved in this release.

PR 3408477: Some ESXi hosts might not have a locker directory after an upgrade from ESXi 6.x to 8.x

When you upgrade an ESXi host with a boot disk of less than 10 GB and not on USB from ESXi 6.x to 8.x, the upgrade process might not create a locker directory and the `/locker` symbolic link is not active.

This issue is resolved in this release. If you already face the issue, upgrading to ESXi 8.0 Update 3b creates a locker directory but does not automatically create a VMware Tools repository. As a result, clusters that host such ESXi hosts display as non-compliant. Remediate the cluster again to create a VMware Tools repository and to become compliant.

PR 3422005: ESXi hosts of version 8.0 Update 2 and later might fail to synchronize time with certain NTP servers

A change in the `ntp-4.2.8p17` package in ESXi 8.0 Update 2 might cause the NTP client to reject certain server packets as poorly formatted or invalid. For example, if a server sends packets with a `poll` value of 0, the NTP client on ESXi does not synchronize with the server.

This issue is resolved in this release.

PR 3408145: The vmx service might fail with a core dump due to rare issue with the vSphere Data Protection solution running out of resources

In rare cases, if a backup operation starts during high concurrent guest I/O load, for example a VM with high write I/O intensity and a high number of overwrites, the VAIO filter component of the vSphere Data Protection solution might run out of resources to handle the guest write I/Os. As a result, the `vmx` service might fail with a core dump and restart.

This issue is resolved in this release. Alternatively, run backups during periods of low guest I/O activity.

PR 3405912: In the vSphere Client, you do not see the correct total vSAN storage consumption for which you have a license

Due to a rare race condition in environments with many vSAN clusters managed by a single vCenter instance, in the vSphere Client under **Licensing > Licenses** you might see a discrepancy between the total claimed vSAN storage capacity and the reported value for the clusters.

This issue is resolved in this release.

PR 3414588: Snapshot tasks on virtual machines on NVMe/TCP datastores might take much longer than on VMs provisioned on NVMe/FC datastores

When creating or deleting a snapshot on a VM provisioned on datastores backed by NVMe/TCP namespaces, such tasks might take much longer than on VMs provisioned on datastores backed by NVMe/FC namespaces. The issue occurs because the `nvmetcp` driver handles some specific NVMe commands not in the way the NVMe/TCP target systems expect.

This issue is resolved in this release.

PR 3392173: A rare issue with the Virsto vSAN component might cause failure to create vSAN objects, or unmount disk groups, or reboot ESXi hosts

In very rare cases, if a Virsto component creation task fails, it might not be properly handled and cause background deletion of virtual disks to stop. As a result, deleting virtual disks in tasks such as creating vSAN objects, or unmounting of disk groups, or rebooting ESXi hosts does not occur as expected and might cause such tasks to fail.

This issue is resolved in this release.

PR 3415365: ESXi upgrade to 8.0 Update 3 fails with an error in the vFAT bootbank partitions

ESXi 8.0 Update 3 adds a precheck in the upgrade workflow that uses the `dosfsck` tool to catch vFAT corruptions. One of the errors that `dosfsck` flags is the `dirty bit set` but ESXi does not use that concept and such errors are false positive.

In the vSphere Client, you see an error such as `A problem with one or more vFAT bootbank partitions was detected`. Please refer to KB 91136 and run `dosfsck` on bootbank partitions.

In the remediation logs on ESXi hosts, you see logs such as:

```
2024-07-02T16:01:16Z In(14) lifecycle[122416262]: runcommand:199 runcommand called with:
args = ['/bin/dosfsck', '-V', '-n', '/dev/disks/naa.600508b1001c7d25f5336a7220b5afc1:6'],
outfile = None, returnoutput= True, timeout = 10.
```

```
2024-07-02T16:01:16Z In(14) lifecycle[122416262]: upgrade_precheck:1836 dosfsck
output: b'CP850//TRANSLIT: Invalid argument\nCP850: Invalid argument\nfsck.fat
4.1+git (2017-01-24)\n0x25: Dirty bit is set. Fswas not properly unmounted and some
data may be corrupt.\n Automatically removing dirty bit.\nStarting check/repair
pass.\nStarting verification pass.\n\nLeaving filesystem unchanged.\n/dev/disks/
naa.600508b1001c7d25f5336a7220b5afc1:6: 121 files, 5665/65515 clusters\n'
```

This issue is resolved in this release.

PR 3421084: An ESXi host might become temporarily inaccessible from vCenter if an NFSv3 datastore fails to mount during reboot or bring up

During an ESXi host reboot, if an NFSv3 datastore fails to mount during the reboot or bring up in VMware Cloud Foundation environments, retries to mount the datastore continue in the background. However, while the datastore is still not available, the `hostd` daemon might fail with a core dump when trying to access it and cause the host to lose connectivity to the vCenter system for a short period.

This issue is resolved in this release.

PR 3407251: ESXi host fails with a purple diagnostic screen due to a rare physical CPU (PCPU) lockup

In the vSphere Client, when you use the **Delete from Disk** option to remove a virtual machine from a vCenter system and delete all VM files from the datastore, including the configuration file and virtual disk files, if any of the files is corrupted, a rare issue with handling corrupted files in the delete path might lead to a PCPU lockup. As a result, the ESXi host fails with a purple diagnostic screen and a message such as `NMI IPI: Panic requested by another PCPU`.

This issue is resolved in this release.

PR 3406627: VMFS6 automatic UNMAP feature might fail to reclaim filesystem space beyond 250GB

In certain cases, when you delete a filesystem space of more than 250GB on a VMFS6 volume, for example 1 TB, if the volume has no active references such as active VMs, then the VMFS6 automatic UNMAP feature might fail to reclaim the filesystem space beyond 250GB.

This issue is resolved in this release.

PR 3419241: After deleting a snapshot or snapshot consolidation, some virtual machines intermittently fail

When a VM on a NFSv3 datastore has multiple snapshots, such as s1, s2, s3, s4, if the VM reverts to one of the snapshots, for example s2, then powers on, and then one of the other snapshots, such as s3, is deleted, the vmx service might fail. The issue occurs because the code tries to consolidate links of a disk that is not part of the VM current state and gets a null pointer. As a result, snapshot consolidation might also fail and cause the vmx service to fail as well.

This issue is resolved in this release. If you already face the issue, power off the VM, edit its `.vmx` file to add the following setting: `consolidate.upgradeNFS3Locks = "FALSE"`, and power on the VM.

PR 3407532: VMs with snapshots and active encryption experience higher I/O latency

In rare cases, encrypted VMs with snapshots might experience higher than expected latency. This issue occurs due to unaligned I/O operations that generate excessive metadata requests to the underlying storage and lead to increased latency.

This issue is resolved in this release. To optimize performance, the VMcrypt I/O filter is enhanced to allocate memory in 4K-aligned blocks for both read and write operations. This reduction in metadata overhead significantly improves overall I/O performance.

PR 3410311: You cannot log in to the Direct Console User Interface (DCUI) with regular Active Directory credentials

When you try to log with a regular Active Directory account in to the DCUI of an ESXi host by using either a remote management application such as HP Integrated Lights-Out (iLO) or Dell Remote Access Card (DRAC), or a server management system such as Lenovo XCC or Huawei iBMC, login might fail. In the DCUI, you see an error such as `Wrong user name or password`. In the `vmkernel.log` file, you see logs such as:

```
2024-07-01T10:40:53.007Z In(182) vmkernel: cpu1:264954)VmkAccess: 106: dcui: running in
dcuiDom(7): socket = /etc/likewise/lib/.lsassd (unix_stream_socket_connect): Access denied
by vmkernel access control policy.
```

The issue occurs due to a restriction of ESXi processes to access certain resources, such as Likewise.

This issue is resolved in this release.

PR 3403706: Hot extending a non-shared disk in a Windows Server Failover Cluster might result in lost reservations on shared disks

In a WSFC cluster, due to an issue with releasing SCSI reservations, in some cases hot extending a non-shared disk might result in lost reservations on shared disks and failover of the disk resource.

This issue is resolved in this release. The fix makes sure releasing SCSI reservations are properly handled for all type of shared disks.

PR 3394043: Creating a vSAN File Service fails when you use IPs within the 172.17.0.0/16 range as mount points

Prior to vSphere 8.0 Update 3b, you must change your network configuration in cases when the specified file service network overlaps with the Docker default internal network 172.17.0.0/16. As a result, you see Skyline Health warnings for DNS lookup failures and you cannot create vSAN File Services.

This issue is resolved in this release. The fix routes traffic to the correct endpoint to avoid possible conflicts.

PR 3389766: During a vSphere vMotion migration of a fault tolerant Primary VM with encryption, the migration task might fail and vSphere FT failover occurs

In rare cases, the encryption key package might not be sent or received correctly during a vSphere vMotion migration of a fault tolerant Primary VM with encryption, and as a result the migration task fails and vSphere FT failover occurs.

This issue is resolved in this release. The fix sends the encryption key at the end of the vSphere FT checkpoint to avoid errors.

PR 3403680: Mounting of vSphere Virtual Volumes stretched storage container fails with an undeclared fault

In the vSphere Client, you might see the error `Undeclared fault` while mounting a newly created vSphere Virtual Volumes stretched storage container from an existing storage array. The issue occurs due to a rare race condition. vSphere Virtual Volumes generates a core dump and restarts after the failure.

This issue is resolved in this release.

PR 3408300: You cannot remove or delete a VMFS partition on a 4K native (4Kn) Software Emulation (SWE) disk

When you attempt to remove or delete a VMFS partition on a 4Kn SWE disk, in the vSphere Client you see an error such as `Read-only file system during write on /dev/disks/<device name>` and the operation fails. In the `vmkernel.log`, you see entries such as `in-use partition <part num>, modification is not supported`.

This issue is resolved in this release.

PR 3402823: Fresh installation or creating VMFS partitions on a Micron 7500 or Intel D5-P5336 NVMe drives might fail with a purple diagnostic screen

UNMAP commands enable ESXi hosts to release storage space that is mapped to data deleted from the host. In NVMe, the equivalent of UNMAP commands is a deallocate DSM request. Micron 7500 and Intel D5-P5336 devices advertise a very large value in one of the deallocate limit attributes, DMSRL, which is the maximum number of logical blocks in a single range for a Dataset Management command. This leads to an integer overflow when the ESXi unmap split code converts number of blocks to number of bytes, which in turn might cause a failure of either installation or VMFS creation. You see a purple diagnostics screen with an error such as `Exception 14 or corruption in dlmalloc`. The issue affects ESXi 8.0 Update 2 and later.

This issue is resolved in this release.

PR 3396479: Standard image profiles for ESXi 8.0 Update 3 show last modified date as release date

The **Release Date** field of the standard image profile for ESXi 8.0 Update 3 shows the **Last Modified Date** value. The issue is only applicable to the image profiles used in Auto Deploy or ESXCLI. Base images used in vSphere Lifecycle Manager workflows display the release date correctly. This issue has no functional impact. The side effect is that if you search for profiles by release date, the profile does not show with the actual release date.

This issue is resolved in this release.

ESXi-8.0U3b-24280767-no-tools

Profile Name	ESXi-8.0U3b-24280767-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	September 17, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-base_8.0.3-0.35.24280767 - VMware_bootbank_bmcad_8.0.3-0.35.24280767 - VMware_bootbank_gc_8.0.3-0.35.24280767 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.35.24280767 - VMware_bootbank_gc-esxio_8.0.3-0.35.24280767 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.35.24280767 - VMware_bootbank_vdfs_8.0.3-0.35.24280767

	• VMware_bootbank_cpu-microcode_8.0.3-0.35.24280767 • VMware_bootbank_vsan_8.0.3-0.35.24280767 • VMware_bootbank_trx_8.0.3-0.35.24280767 • VMware_bootbank_esx-xserver_8.0.3-0.35.24280767 • VMware_bootbank_clusterstore_8.0.3-0.35.24280767 • VMware_bootbank_esxio-combiner_8.0.3-0.35.24280767 • VMware_bootbank_vcls-pod-crx_8.0.3-0.35.24280767 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.35.24280767 • VMware_bootbank_vds-vsip_8.0.3-0.35.24280767 • VMware_bootbank_vsanhealth_8.0.3-0.35.24280767 • VMware_bootbank_native-misc-drivers_8.0.3-0.35.24280767 • VMware_bootbank_crx_8.0.3-0.35.24280767 • VMware_bootbank_esxio_8.0.3-0.35.24280767 • VMware_bootbank_infravisor_8.0.3-0.35.24280767 • VMware_bootbank_drivervm-gpu-base_8.0.3-0.35.24280767 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.35.24280767 • VMware_bootbank_bmc-al-esxio_8.0.3-0.35.24280767 • VMware_bootbank_esx-base_8.0.3-0.35.24280767 • VMware_bootbank_loadesx_8.0.3-0.35.24280767 • VMware_bootbank_esx-update_8.0.3-0.35.24280767 • VMware_bootbank_loadesxio_8.0.3-0.35.24280767 • VMware_bootbank_esxio-update_8.0.3-0.35.24280767 • VMW_bootbank_rshim_0.1-12vmw.803.0.35.24280767 • VMW_bootbank_rshim-net_0.1.0-1vmw.803.0.35.24280767 • VMW_bootbank_mlnx-bfbootctl-esxio_0.1-7vmw.803.0.35.24280767 • VMW_bootbank_pensandoatlas_1.46.0.E.41.1.326-2vmw.803.0.0.0.23797590
PRs Fixed	3426946, 3421084, 3421179, 3403683, 3388844, 3424051, 3422735, 3420488, 3421971, 3422005, 3417326, 3420850, 3420421, 3420586, 3410311, 3420907, 3419241, 3409528, 3404817, 3416221, 3421665, 3415365, 3415102, 3417329, 3421434, 3419074, 3417667, 3408802, 3417224, 3420702, 3385757, 3412010, 3418878, 3408477, 3406546, 3406968, 3406999, 3398549, 3402823, 3408281, 3409108, 3406627, 3407251, 3392225, 3412138, 3406037, 3406875, 3412536, 3407952, 3410114, 3386751, 3408145, 3403240, 3389766, 3283598, 3408302, 3403639, 3407532, 3394043, 3387721, 3408300, 3396479, 3401629, 3401392, 3409085, 3408640, 3405912, 3403680, 3405106, 3407204, 3407022, 3409176, 3408739, 3408740, 3390434, 3406053, 3400702, 3392173, 3403496, 3402955, 3433092, 3421179, 3414588
Related CVE numbers	N/A

This patch updates the following issues:

PR 3388844: You cannot activate Kernel Direct Memory Access (DMA) Protection for Windows guest OS on ESXi hosts with Intel CPU

If an input–output memory management unit (IOMMU) is active on a Windows guest OS, the **Kernel DMA Protection** option under **System Information** might be off for VMs running on ESXi hosts with Intel CPUs. As a result, you might not be able to fulfill some security requirements for your environment.

This issue is resolved in this release. The fix activates **Kernel DMA Protection** by default. The fix also adds the `vmx` parameter `acpi.dmar.enableDMAProtection` to activate the **Kernel DMA Protection** in Windows OS. The default value of the parameter is `FALSE` and to activate **Kernel DMA Protection** in a Windows guest OS, you must add `acpi.dmar.enableDMAProtection=TRUE` to the `vmx` file.

PR 3412448: Remediation of ESXi hosts might fail due to a timeout in the update of the VMware-VM-Tools component

While upgrading ESXi hosts by using vSphere Lifecycle Manager, remediation of the VMware-VM-Tools component might occasionally fail due to a timeout. The issue occurs when the update process takes longer than the timeout setting of 30 sec for the VMware-VM-Tools component.

This issue is resolved in this release. The fix increases the timeout limit to 120 sec.

PR 3377863: You see "Hosts are remediated" message in the upgrade precheck results

When running a precheck before an ESXi update, you see a message such as `Hosts are remediated` in the precheck results, which is not clear and might be misleading.

This issue is resolved in this release. The new message is `Hosts are remediated sequentially`.

PR 3417329: Virtual machine tasks might intermittently fail due to a rare issue with the memory slab

Due to a rare issue with the vSAN DOM object, where a reference count on a component object might not decrement correctly, the in-memory object might never be released from the slab and can cause the component manager slab to reach its limit. As a result, you might not be able to create VMs, migrate VMs or might encounter VM power-on failures on vSAN clusters, either OSA or ESA.

This issue is resolved in this release.

PR 3406140: Extracting a vSphere Lifecycle Manager image from an existing ESXi host might fail after a kernel configuration change

Each update of the kernel configuration also triggers an update to the `/bootbank/useropts.gz` file, but due to a known issue, the `basemisc.tgz` might not contain the default `useropts` file after such an update. As a result, when attempting to extract a vSphere Lifecycle Manager image from an existing ESXi host, the absence of the default `useropts` file leads to failure to create the `esx-base.vib` file and the operation also fails.

This issue is resolved in this release.

PR 3408477: Some ESXi hosts might not have a locker directory after an upgrade from ESXi 6.x to 8.x

When you upgrade an ESXi host with a boot disk of less than 10 GB and not on USB from ESXi 6.x to 8.x, the upgrade process might not create a locker directory and the `/locker` symbolic link is not active.

This issue is resolved in this release. If you already face the issue, upgrading to ESXi 8.0 Update 3b creates a locker directory but does not automatically create a VMware Tools repository. As a result, clusters that host such ESXi hosts display as non-compliant. Remediate the cluster again to create a VMware Tools repository and to become compliant.

PR 3422005: ESXi hosts of version 8.0 Update 2 and later might fail to synchronize time with certain NTP servers

A change in the `ntp-4.2.8p17` package in ESXi 8.0 Update 2 might cause the NTP client to reject certain server packets as poorly formatted or invalid. For example, if a server sends packets with a `poll` value of 0, the NTP client on ESXi does not synchronize with the server.

This issue is resolved in this release.

PR 3408145: The vmx service might fail with a core dump due to rare issue with the vSphere Data Protection solution running out of resources

In rare cases, if a backup operation starts during high concurrent guest I/O load, for example a VM with high write I/O intensity and a high number of overwrites, the VAIO filter component of the vSphere Data Protection solution might run out of resources to handle the guest write I/Os. As a result, the `vmx` service might fail with a core dump and restart.

This issue is resolved in this release. Alternatively, run backups during periods of low guest I/O activity.

PR 3405912: In the vSphere Client, you do not see the correct total vSAN storage consumption for which you have a license

Due to a rare race condition in environments with many vSAN clusters managed by a single vCenter instance, in the vSphere Client under **Licensing > Licenses** you might see a discrepancy between the total claimed vSAN storage capacity and the reported value for the clusters.

This issue is resolved in this release.

PR 3414588: Snapshot tasks on virtual machines on NVMe/TCP datastores might take much longer than on VMs provisioned on NVMe/FC datastores

When creating or deleting a snapshot on a VM provisioned on datastores backed by NVMe/TCP namespaces, such tasks might take much longer than on VMs provisioned on datastores backed by NVMe/FC namespaces. The issue occurs because the nvmetcp driver handles some specific NVMe commands not in the way the NVMe/TCP target systems expect.

This issue is resolved in this release.

PR 3392173: A rare issue with the Virsto vSAN component might cause failure to create vSAN objects, or unmount disk groups, or reboot ESXi hosts

In very rare cases, if a Virsto component creation task fails, it might not be properly handled and cause background deletion of virtual disks to stop. As a result, deleting virtual disks in tasks such as creating vSAN objects, or unmounting of disk groups, or rebooting ESXi hosts does not occur as expected and might cause such tasks to fail.

This issue is resolved in this release.

PR 3415365: ESXi upgrade to 8.0 Update 3 fails with an error in the vFAT bootbank partitions

ESXi 8.0 Update 3 adds a precheck in the upgrade workflow that uses the `dosfsck` tool to catch vFAT corruptions. One of the errors that `dosfsck` flags is the `dirty bit set` but ESXi does not use that concept and such errors are false positive.

In the vSphere Client, you see an error such as `A problem with one or more vFAT bootbank partitions was detected`. Please refer to KB 91136 and run `dosfsck` on bootbank partitions.

In the remediation logs on ESXi hosts, you see logs such as:

```
2024-07-02T16:01:16Z In(14) lifecycle[122416262]: runcommand:199 runcommand called with:
args = ['/bin/dosfsck', '-V', '-n', '/dev/disks/naa.600508b1001c7d25f5336a7220b5afc1:6'],
outfile = None, returnoutput= True, timeout = 10.
```

```
2024-07-02T16:01:16Z In(14) lifecycle[122416262]: upgrade_precheck:1836 dosfsck
output: b'CP850//TRANSLIT: Invalid argument\nCP850: Invalid argument\nfsck.fat
4.1+git (2017-01-24)\n0x25: Dirty bit is set. Fswas not properly unmounted and some
data may be corrupt.\n Automatically removing dirty bit.\nStarting check/repair
pass.\nStarting verification pass.\n\nLeaving filesystem unchanged.\n/dev/disks/
naa.600508b1001c7d25f5336a7220b5afc1:6: 121 files, 5665/65515 clusters\n'
```

This issue is resolved in this release.

PR 3421084: An ESXi host might become temporarily inaccessible from vCenter if an NFSv3 datastore fails to mount during reboot or bring up

During an ESXi host reboot, if an NFSv3 datastore fails to mount during the reboot or bring up in VMware Cloud Foundation environments, retries to mount the datastore continue in the background. However, while the datastore is still not available, the `hostd` daemon might fail with a core dump when trying to access it and cause the host to lose connectivity to the vCenter system for a short period.

This issue is resolved in this release.

PR 3407251: ESXi host fails with a purple diagnostic screen due to a rare physical CPU (PCPU) lockup

In the vSphere Client, when you use the **Delete from Disk** option to remove a virtual machine from a vCenter system and delete all VM files from the datastore, including the configuration file and virtual disk files, if any of the files is corrupted, a rare issue with handling corrupted files in the delete path might lead to a PCPU lockup. As a result, the ESXi host fails with a purple diagnostic screen and a message such as `NMI IPI: Panic requested by another PCPU`.

This issue is resolved in this release.

PR 3406627: VMFS6 automatic UNMAP feature might fail to reclaim filesystem space beyond 250GB

In certain cases, when you delete a filesystem space of more than 250GB on a VMFS6 volume, for example 1 TB, if the volume has no active references such as active VMs, then the VMFS6 automatic UNMAP feature might fail to reclaim the filesystem space beyond 250GB.

This issue is resolved in this release.

PR 3419241: After deleting a snapshot or snapshot consolidation, some virtual machines intermittently fail

When a VM on a NFSv3 datastore has multiple snapshots, such as s1, s2, s3, s4, if the VM reverts to one of the snapshots, for example s2, then powers on, and then one of the other snapshots, such as s3, is deleted, the vmx service might fail. The issue occurs because the code tries to consolidate links of a disk that is not part of the VM current state and gets a null pointer. As a result, snapshot consolidation might also fail and cause the vmx service to fail as well.

This issue is resolved in this release. If you already face the issue, power off the VM, edit its `.vmx` file to add the following setting: `consolidate.upgradeNFS3Locks = "FALSE"`, and power on the VM.

PR 3407532: VMs with snapshots and active encryption experience higher I/O latency

In rare cases, encrypted VMs with snapshots might experience higher than expected latency. This issue occurs due to unaligned I/O operations that generate excessive metadata requests to the underlying storage and lead to increased latency.

This issue is resolved in this release. To optimize performance, the VMcrypt I/O filter is enhanced to allocate memory in 4K-aligned blocks for both read and write operations. This reduction in metadata overhead significantly improves overall I/O performance.

PR 3410311: You cannot log in to the Direct Console User Interface (DCUI) with regular Active Directory credentials

When you try to log with a regular Active Directory account in to the DCUI of an ESXi host by using either a remote management application such as HP Integrated Lights-Out (iLO) or Dell Remote Access Card (DRAC), or a server management system such as Lenovo XCC or Huawei iBMC, login might fail. In the DCUI, you see an error such as `Wrong user name or password`. In the `vmkernel.log` file, you see logs such as:

```
2024-07-01T10:40:53.007Z In(182) vmkernel: cpu1:264954)VmkAccess: 106: dcui: running in
dcuiDom(7): socket = /etc/likewise/lib/.lsassd (unix_stream_socket_connect): Access denied
by vmkernel access control policy.
```

The issue occurs due to a restriction of ESXi processes to access certain resources, such as Likewise.

This issue is resolved in this release.

PR 3403706: Hot extending a non-shared disk in a Windows Server Failover Cluster might result in lost reservations on shared disks

In a WSFC cluster, due to an issue with releasing SCSI reservations, in some cases hot extending a non-shared disk might result in lost reservations on shared disks and failover of the disk resource.

This issue is resolved in this release. The fix makes sure releasing SCSI reservations are properly handled for all type of shared disks.

PR 3394043: Creating a vSAN File Service fails when you use IPs within the 172.17.0.0/16 range as mount points

Prior to vSphere 8.0 Update 3b, you must change your network configuration in cases when the specified file service network overlaps with the Docker default internal network 172.17.0.0/16. As a result, you see Skyline Health warnings for DNS lookup failures and you cannot create vSAN File Services.

This issue is resolved in this release. The fix routes traffic to the correct endpoint to avoid possible conflicts.

PR 3389766: During a vSphere vMotion migration of a fault tolerant Primary VM with encryption, the migration task might fail and vSphere FT failover occurs

In rare cases, the encryption key package might not be sent or received correctly during a vSphere vMotion migration of a fault tolerant Primary VM with encryption, and as a result the migration task fails and vSphere FT failover occurs.

This issue is resolved in this release. The fix sends the encryption key at the end of the vSphere FT checkpoint to avoid errors.

PR 3403680: Mounting of vSphere Virtual Volumes stretched storage container fails with an undeclared fault

In the vSphere Client, you might see the error `Undeclared fault` while mounting a newly created vSphere Virtual Volumes stretched storage container from an existing storage array. The issue occurs due to a rare race condition. vSphere Virtual Volumes generates a core dump and restarts after the failure.

This issue is resolved in this release.

PR 3408300: You cannot remove or delete a VMFS partition on a 4K native (4Kn) Software Emulation (SWE) disk

When you attempt to remove or delete a VMFS partition on a 4Kn SWE disk, in the vSphere Client you see an error such as `Read-only file system during write on /dev/disks/<device name>` and the operation fails. In the `vmkernel.log`, you see entries such as `in-use partition <part num>, modification is not supported`.

This issue is resolved in this release.

PR 3402823: Fresh installation or creating VMFS partitions on a Micron 7500 or Intel D5-P5336 NVMe drives might fail with a purple diagnostic screen

UNMAP commands enable ESXi hosts to release storage space that is mapped to data deleted from the host. In NVMe, the equivalent of UNMAP commands is a deallocate DSM request. Micron 7500 and Intel D5-P5336 devices advertise a very large value in one of the deallocate limit attributes, DMSRL, which is the maximum number of logical blocks in a single range for a Dataset Management command. This leads to an integer overflow when the ESXi unmap split code converts number of blocks to number of bytes, which in turn might cause a failure of either installation or VMFS creation. You see a purple diagnostics screen with an error such as `Exception 14 or corruption in dlmalloc`. The issue affects ESXi 8.0 Update 2 and later.

This issue is resolved in this release.

PR 3396479: Standard image profiles for ESXi 8.0 Update 3 show last modified date as release date

The **Release Date** field of the standard image profile for ESXi 8.0 Update 3 shows the **Last Modified Date** value. The issue is only applicable to the image profiles used in Auto Deploy or ESXCLI. Base images used in vSphere Lifecycle Manager workflows display the release date correctly. This issue has no functional impact. The side effect is that if you search for profiles by release date, the profile does not show with the actual release date.

This issue is resolved in this release.

ESXi-8.0U3sb-24262298-standard

Profile Name	ESXi-8.0U3sb-24262298-standard
Build	For build information, see Patches Contained in This Release .

Vendor	VMware by Broadcom, Inc.
Release Date	September 17, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_drivervm-gpu-base_8.0.3-0.30.24262298 - VMware_bootbank_esxio-combiner_8.0.3-0.30.24262298 - VMware_bootbank_esx-xserver_8.0.3-0.30.24262298 - VMware_bootbank_vdfs_8.0.3-0.30.24262298 - VMware_bootbank_cpu-microcode_8.0.3-0.30.24262298 - VMware_bootbank_vds-vsip_8.0.3-0.30.24262298 - VMware_bootbank_vsan_8.0.3-0.30.24262298 - VMware_bootbank_bmcals_8.0.3-0.30.24262298 - VMware_bootbank_clusterstore_8.0.3-0.30.24262298 - VMware_bootbank_native-misc-drivers_8.0.3-0.30.24262298 - VMware_bootbank_infravisor_8.0.3-0.30.24262298 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.30.24262298 - VMware_bootbank_crx_8.0.3-0.30.24262298 - VMware_bootbank_gc_8.0.3-0.30.24262298 - VMware_bootbank_vcls-pod-crx_8.0.3-0.30.24262298 - VMware_bootbank_esxio-base_8.0.3-0.30.24262298 - VMware_bootbank_esx-base_8.0.3-0.30.24262298 - VMware_bootbank_esxio_8.0.3-0.30.24262298 - VMware_bootbank_gc-esxio_8.0.3-0.30.24262298 - VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.30.24262298 - VMware_bootbank_esxio-combiner-esxio_8.0.3-0.30.24262298 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.30.24262298 - VMware_bootbank_vsanhealth_8.0.3-0.30.24262298 - VMware_bootbank_trx_8.0.3-0.30.24262298 - VMware_bootbank_bmcals-esxio_8.0.3-0.30.24262298 - VMware_bootbank_esx-update_8.0.3-0.30.24262298 - VMware_bootbank_loadesx_8.0.3-0.30.24262298 - VMware_bootbank_esxio-update_8.0.3-0.30.24262298 - VMware_bootbank_loadesxio_8.0.3-0.30.24262298 - VMware_locker_tools-light_12.4.5.23787635-24262298 - VMW_bootbank_pensandoatlas_1.46.0.E.41.1.326-2vmw.803.0.0.0.23797590
PRs Fixed	3425039, 3423080, 3408352, 3415908, 3380359, 3404366, 3421357, 3421359, 3390663, 3404362, 3396479, 3404367, 3398130, 3398132, 3316536, 3410159, 3412279, 3395162, 3404369, 3404657, 3390662, 3409328
Related CVE numbers	N/A

This patch updates the following issues:

PR 3392225: You see status Unknown for virtual machines after migration to an ESXi host with active host-local swap location

In rare occasions with very specific sequence of conditions, when you migrate to or power on a virtual machine on an ESXi host with active host-local swap location, the virtual machine might hang with a status `Unknown`.

This issue is resolved in this release.

VMware Tools Bundling Changes in ESXi 8.0 Update 3b

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 2b:

- **windows.iso**: VMware Tools 12.4.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.4.5 Release Notes](#)
- [Earlier versions of VMware Tools](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Updating VMware Tools](#)

ESXi 8.0 Update 3b includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
						Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007206	4/15/2024	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
						Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003707	3/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003707	3/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	baseline	0x7002904	4/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd0003e7	4/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002b0	4/3/2024	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xb8	2/15/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x38	2/15/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x52	2/15/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b0005c0	2/5/2024	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x100	2/5/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x62	3/7/2024	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x125	4/16/2024	Intel Xeon E-2400 Series
Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x21000200	11/20/2023	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series

ESX 8.0 Update 3b contains the following security updates:

- OpenSSH is updated to version 9.8p1.

ESXi-8.0U3sb-24262298-no-tools

Profile Name	ESXi-8.0U3sb-24262298-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	September 17, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	• VMware_bootbank_drivervm-gpu-base_8.0.3-0.30.24262298 • VMware_bootbank_esxio-combiner_8.0.3-0.30.24262298 • VMware_bootbank_esx-xserver_8.0.3-0.30.24262298 • VMware_bootbank_vdfs_8.0.3-0.30.24262298 • VMware_bootbank_cpu-microcode_8.0.3-0.30.24262298 • VMware_bootbank_vds-vsip_8.0.3-0.30.24262298 • VMware_bootbank_vsan_8.0.3-0.30.24262298 • VMware_bootbank_bmcad_8.0.3-0.30.24262298 • VMware_bootbank_clusterstore_8.0.3-0.30.24262298 • VMware_bootbank_native-misc-drivers_8.0.3-0.30.24262298 • VMware_bootbank_infravisor_8.0.3-0.30.24262298 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.3-0.30.24262298 • VMware_bootbank_crx_8.0.3-0.30.24262298 • VMware_bootbank_gc_8.0.3-0.30.24262298 • VMware_bootbank_vcls-pod-crx_8.0.3-0.30.24262298 • VMware_bootbank_esxio-base_8.0.3-0.30.24262298 • VMware_bootbank_esx-base_8.0.3-0.30.24262298 • VMware_bootbank_esxio_8.0.3-0.30.24262298 • VMware_bootbank_gc-esxio_8.0.3-0.30.24262298 • VMware_bootbank_native-misc-drivers-esxio_8.0.3-0.30.24262298 • VMware_bootbank_esxio-combiner-esxio_8.0.3-0.30.24262298 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.3-0.30.24262298 • VMware_bootbank_vsanhealth_8.0.3-0.30.24262298 • VMware_bootbank_trx_8.0.3-0.30.24262298 • VMware_bootbank_bmcad-esxio_8.0.3-0.30.24262298 • VMware_bootbank_esx-update_8.0.3-0.30.24262298 • VMware_bootbank_loadesx_8.0.3-0.30.24262298 • VMware_bootbank_esxio-update_8.0.3-0.30.24262298 • VMware_bootbank_loadesxio_8.0.3-0.30.24262298 • VMW_bootbank_pensandoatlas_1.46.0.E.41.1.326-2vmw.803.0.0.0.23797590

PRs Fixed	3425039, 3423080, 3408352, 3415908, 3380359, 3404366, 3421357, 3421359, 3390663, 3404362, 3396479, 3404367, 3398130, 3398132, 3316536, 3410159, 3412279, 3395162, 3404369, 3404657, 3390662, 3409328
Related CVE numbers	N/A

This patch updates the following issues:

VMware Tools Bundling Changes in ESXi 8.0 Update 3b

The following VMware Tools ISO images are bundled with ESXi 8.0 Update 2b:

- **windows.iso**: VMware Tools 12.4.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.4.5 Release Notes](#)
- [Earlier versions of VMware Tools](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Updating VMware Tools](#)

ESXi 8.0 Update 3b includes the following Intel microcode:

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	baseline	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	baseline	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	baseline	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	baseline	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Westmere EP	0x206c2 (06/2c/2)	0x03	baseline	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	baseline	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	baseline	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	baseline	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	baseline	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	baseline	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	baseline	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	baseline	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	baseline	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	baseline	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	baseline	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	baseline	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	baseline	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	baseline	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	baseline	0x2007206	4/15/2024	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
						Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	baseline	0x4003707	3/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	baseline	0x5003707	3/1/2024	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	baseline	0x7002904	4/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	baseline	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	baseline	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	baseline	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	baseline	0xe000015	8/3/2023	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	baseline	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	baseline	0x3e	10/5/2023	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	baseline	0xd0003e7	4/1/2024	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series

Code Name	FMS	Plat ID	Servicing	MCU Rev	MCU Date	Brand Names
Ice Lake D	0x606c1 (06/6c/1)	0x10	baseline	0x10002b0	4/3/2024	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	baseline	0x4c000026	2/28/2024	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	baseline	0xb8	2/15/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	baseline	0x38	2/15/2024	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	baseline	0x52	2/15/2024	Intel Xeon W-11000E Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	baseline	0x2b0005c0	2/5/2024	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	baseline	0xf8	9/28/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	baseline	0xf6	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	baseline	0xf8	2/1/2024	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	baseline	0x100	2/5/2024	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	baseline	0x62	3/7/2024	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	baseline	0x125	4/16/2024	Intel Xeon E-2400 Series
Emerald Rapids SP	0xc06f2 (06/cf/2)	0x87	baseline	0x21000200	11/20/2023	Intel Xeon 8500 Series; Intel Xeon Gold 6500/5500 Series; Intel Xeon Silver 4500 Series; Intel Xeon Bronze 3500 Series

ESX 8.0 Update 3b contains the following security updates:

- OpenSSH is updated to version 9.8p1.

ESXi-8.0U3b-24280767

Name	ESXi
Version	ESXi-8.0U3b-24280767
Release Date	September 17, 2024
Category	Bugfix
Affected Components	• ESXi Component - core ESXi VIBs • ESXi Install/Upgrade Component • VMware BlueField RShim Driver • Networking Device Driver for NVIDIA BlueField RShim devices • NVIDIA BlueField boot control driver • VMware NVMe over TCP Driver
PRs Fixed	3426946, 3421084, 3421179, 3403683, 3388844, 3424051, 3422735, 3420488, 3421971, 3422005, 3417326, 3420850, 3420421, 3420586, 3410311, 3420907, 3419241, 3409528, 3404817, 3416221, 3421665, 3415365, 3415102, 3417329, 3421434, 3419074, 3417667, 3408802, 3417224, 3420702, 3385757, 3412010, 3418878, 3408477, 3406546, 3406968, 3406999, 3398549, 3402823, 3408281, 3409108, 3406627, 3407251, 3392225, 3412138, 3406037, 3406875, 3412536, 3407952, 3410114, 3386751, 3408145, 3403240, 3389766, 3283598, 3408302, 3403639, 3407532, 3394043, 3387721, 3408300, 3396479, 3401629, 3401392, 3409085, 3408640, 3405912, 3403680, 3405106, 3407204, 3407022, 3409176, 3408739, 3408740, 3390434, 3406053, 3400702, 3392173, 3403496, 3402955, 3433092, 3421179, 3414588
Related CVE numbers	N/A

This patch resolves the issues listed in [ESXi_8.0.3-0.35.24280767](#).

ESXi-8.0U3sb-24262298

Name	ESXi
Version	ESXi-8.0U3sb-24262298
Release Date	September 17, 2024
Category	Security
Affected Components	• ESXi Component - core ESXi VIBs • ESXi Install/Upgrade Component • ESXi Tools Component
PRs Fixed	3425039, 3423080, 3408352, 3415908, 3380359, 3404366, 3421357, 3421359, 3390663, 3404362, 3396479, 3404367, 3398130, 3398132, 3316536, 3410159, 3412279, 3395162, 3404369, 3404657, 3390662, 3409328
Related CVE numbers	N/A

This patch resolves the issues listed in [ESXi_8.0.3-0.30.24262298](#).

Known Issues

Installation and Upgrade Issues

If you use an offline bundle from a vendor, remediation of ESXi hosts with a Non-Critical Host Patches baseline might fail

If you import an offline bundle from a vendor depot, remediation of ESXi hosts with a Non-Critical Host Patches baseline might fail with a profile validation error.

Workaround: Instead of using a Non-Critical Host Patches baseline, create a new rollup bulletin that contains the necessary updates.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

You cannot update to ESXi 8.0 Update 2b by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see [Upgrade or Update a Host with Image Profiles](#).

The Cancel option in an interactive ESXi installation might not work as expected

Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings > <specific broadcom NIC> > Device Configuration Menu > Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

The `irdman` driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic

If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the `irdman` driver might fail. This issue is unlikely to occur, as the `irdman` driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

Workaround: Use ULPs with RC transport mode.

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1

The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

Capture of network packets by using the PacketCapture tool on ESXi does not work

Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in [Collecting network packets using the lightweight PacketCapture on ESXi](#).

Workaround: Use the pktcap-uw tool. For more information, see [Capture and Trace Network Packets by Using the pktcap-uw Utility](#).

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.

2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out

vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD

In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See [Reconciling Discrepancies in the Managed Virtual Disk Catalog](#).

vCenter Server and vSphere Client Issues

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

VMware ESXi 8.0 Update 3 Release Notes

This document contains the following sections

- [Introduction](#)
- [General Availability](#)
- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Product Support Notices](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 3 | 25 JUN 2024 | ISO Build 24022510
Check for additions and updates to these release notes.

General Availability

This ESXi 8.0 Update 3 release is a General Availability (GA) designation. For more information on the vSphere 8.0 IA/GA Release Model of vSphere Update releases, see [The vSphere 8 Release Model Evolves](#).

What's New

This release resolves CVE-2024-37085 and CVE-2024-37086. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0013](#).

For overview of the new features in vSphere 8.0 Update 3, see [What's New in vSphere 8 Update 3](#) and the following release notes:

- [vSphere IaaS Control Plane \(former VMware vSphere With Tanzu\)](#)
- [Tanzu Kubernetes releases](#)
- [VMware vSAN](#)
- [VMware Host Client](#)
- [VMware OVF Tool](#)

The list of new features and enhancements that follows adds some of the key highlights for vSphere 8.0 Update 3:

DPU/SmartNIC

- **High Availability with VMware vSphere Distributed Services Engine:** Starting with ESXi 8.0 Update 3, vSphere Distributed Services Engine adds support for 2 data processing units (DPUs) to provide high availability or increase offload capacity per ESXi host. Dual-DPU systems can use NVIDIA or Pensando devices. In ESXi 8.0 Update 3, dual-DPU systems are supported by Lenovo server designs. For more information, see [High Availability with VMware vSphere Distributed Services Engine](#).

vSphere IaaS control plane

- **Support for running vSphere IaaS control plane on vSAN stretched clusters:** vSphere 8.0 Update 3 adds support for vSphere IaaS control plane, formerly known as vSphere with Tanzu, on vSAN stretched clusters. This capability is available only for greenfield installations of the IaaS control plane on stretched vSAN clusters. If you already have IaaS control pane or vSphere with Tanzu plus Content Library running, you need to redeploy it before you use the vSAN Stretched Cluster Storage Policies.
Before configuring this capability, read [Running vSphere IaaS Control Plane on vSAN Stretched Cluster](#) and follow content on <https://core.vmware.com>.

Virtual Machine Management

- **New Virtual Machine Compute Policy for Best Effort Virtual Machine Evacuation:** vSphere 8.0 Update 3 adds a compute policy for best effort evacuation of virtual machines on ESXi hosts that are entering maintenance mode. When the host enters maintenance mode, all VMs are shut down. If shutdown fails, the VMs are powered off. If power off fails, you need to evacuate the VMs. With the new capability, when the VMs are in a powered-off state, vCenter attempts to power them on every few minutes on the best available ESXi host at the time. This policy overrides any DRS overrides set at the VM level, and the hosts on which the VMs are powered on might be different from the original host.

vSphere Cluster Service

- **Introducing Embedded vSphere Cluster Service (vCLS):** vSphere 8.0 Update 3 introduces a redesign of vCLS to Embedded vCLS, which utilizes vSphere Pod technology. Deployment and lifecycle of these VMs are managed within ESXi and are no longer managed by the vSphere ESX Agent Manager (EAM). Earlier versions of vCLS are termed External vCLS. Issues previously encountered with External vCLS are resolved in this release of Embedded vCLS. While existing API compatibility is preserved, minor modifications to customer scripts or automation tools might be necessary. Other products and solutions that have defined business logic around External vCLS might not work with Embedded vCLS. See individual product documentation to understand the interoperability support and impact. For more information, see [vSphere Cluster Services](#) and content available from vSphere Technical Marketing on the [Broadcom Support Portal](#).

vSAN

- **vSAN add-on licenses based on capacity per tebibyte (TiB):** With vSphere 8.0 Update 3, you can license your use of vSAN storage based on TiB capacity. The new capacity-based license replaces the core and CPU-based licenses. For more information on capacity reporting and licensing in vSAN, see [License Requirements](#) and [Counting Cores for VMware Cloud Foundation and vSphere Foundation and TiBs for vSAN](#).

Security

- **TLS 1.3 and 1.2 support by using TLS profiles:** Starting with vSphere 8.0 Update 3, you can use TLS profiles to simplify the configuration of TLS parameters and improve supportability in your vSphere system. With vSphere 8.0 Update 3, you get a default TLS profile on ESXi and vCenter Server hosts, COMPATIBLE, which supports TLS 1.3 and some TLS 1.2 connections. For more information, see [vSphere TLS Configuration](#).
- **PingFederate Identity Provider for vSphere:** With vSphere 8.0 Update 3, you can configure PingFederate as an external identity provider in your vSphere system. For more information, see [Configuring vCenter Server Identity Provider for PingFederate](#).

Storage/Memory

- **Memory Tiering:** vSphere 8.0 Update 3 launches in tech preview the Memory Tiering capability, which allows you to use NVMe devices that you can add locally to an ESXi host as tiered memory. Memory tiering over NVMe optimizes memory utilization by directing VM memory allocations to either NVMe devices or faster dynamic random access memory (DRAM) in the host. This allows you to increase your memory footprint and workload capacity, while reducing the total cost of ownership (TCO). For more details on the tech preview, see KB [95944](#).
- **Fabric Notification support for SAN clusters:** ESXi 8.0 Update 3 introduces support for Fabric Performance Impact Notifications Link Integrity (FPIN-LI). With FPIN-LI, the vSphere infrastructure layer can manage notifications from SAN switches or targets, identifying degraded SAN links and ensuring only healthy paths are used for storage devices. FPIN can also notify ESXi hosts for storage link congestion and errors.
- **Support for space reclamation requests from guest operating systems on NVMe-backed vSphere Virtual Volumes datastores and Config-vVol:** ESXi 8.0 Update 3 adds support for automatic space reclamation requests from guest operating systems on NVMe-backed vSphere Virtual Volumes datastores. ESXi 8.0 Update 3 also adds support for both command line-based and automatic unmap for vSphere Virtual Volumes objects of type Config-vVol, formatted with VMFS-6. For more information, see [Reclaim Space on the vSphere Virtual Volumes Datastores](#).
- **Manage the UNMAP load from ESXi hosts at a VMFS datastore level:** Starting with ESXi 8.0 Update 3, you can control the unmap load at datastore level to avoid time delays from space reclamation and reduce overall unmap load on the arrays in your environment. For more information, see [Space Reclamation on vSphere VMFS Datastores](#).
- **Windows Server Failover Clustering (WSFC) enhancements on vSphere Virtual Volumes:** vSphere 8.0 Update 3 adds support for a WSFC solution for NVMe-backed disks on vSphere Virtual Volumes. This capability allows NVMe reservations support in NVMe/TCP environments apart from Fibre Channel support for WSFC on vSphere Virtual Volumes. Virtual NVMe (vNVME) controllers are supported as the frontend for WSFC with NVMe-backed disks, not with SCSI-backed disks. For more information, see [VMware vSphere® Virtual Volumes Support for WSFC](#).
- **Support for active-active vSphere Metro Storage Cluster (vMSC) with vSphere Virtual Volumes:** vSphere 8.0 Update 3 introduces a new version of the VMware vSphere Storage APIs for Array Integration (VASA) to add support to active-active stretched storage clusters with vSphere Virtual Volumes, with active-active deployment topologies for SCSI block access between two sites. VASA version 6 includes new architecture and design for VASA Provider High Availability support for both stretched and non-stretched storage clusters. For more information, see [Using Stretched Storage Clustering with Virtual Volumes](#).
- **VMkernel port binding for NFS v4.1 datastores:** With vSphere 8.0 Update 3, you can bind an NFS 4.1 connection to a specific VM kernel adapter. If you use multipathing, you can provide multiple vmknics for each connection to ensure path isolation and security by directing NFS traffic across a specified subnet/VLAN, so that the NFS traffic does not use other vmknics. For more information, see [Configure VMkernel Binding for NFS 4.1 Datastores](#).
- **Support for nConnect for NFS v4.1 datastores:** ESXi 8.0 Update 3 adds support to multiple TCP connections for a NFS v4.1 volume, also referred to as nConnect. For NFS v4.1, multiple TCP connections are created for a single session that many datastores can share in parallel. It can be configured by using either the vSphere API or ESXCLI directly on an ESXi host. For more information, see [Configure Multiple TCP Connections for NFS](#).
- **Reduced time to inflate VMFS disks:** With vSphere 8.0 Update 3, a new VMFS API allows you to inflate a thin-provisioned disk to eagerzeroedthick (EZT) while the disk is in use and up to 10 times faster than previous alternative methods. During the inflation, all blocks are fully allocated and zeroed upfront to allow faster run-time performance. For more information, see [Virtual Disk Options](#).
- **Improved resiliency against memory corruption on RAM-heavy ESXi hosts:** vSphere 8.0 Update 3 adds proactive measures to prevent memory errors in systems with VMs of more than 1TB that might bring down an entire ESXi host and increase application downtime.

- **Advanced setting to block deletion and removal of disks for VMs with snapshots:** ESX 8.0 Update 3 adds a per-host advanced option `blockDiskRemoveIfSnapshot` to prevent the removal of disks from a VM that has snapshots, even if you choose to delete files, which might lead to orphaned disks. For more information, see VMware knowledge base article [94545](#).
- **Hardware accelerated move (clone operation) support on NVMe devices:** vSphere 8.0 Update 3 supports NVMe copy command for hardware accelerated move (also called clone blocks or copy offload) across or within NVMe namespaces that belong to the same NVMe subsystem.
- **Granular monitoring for VASA provider accessibility and certification authentication status on ESXi host level:** With vSphere 8.0 Update 3, you can monitor connection status and ESXi authentication with VASA storage providers from the vSphere Client and re-authenticate hosts as necessary. For more information, see [Reauthenticate VASA Client in ESXi](#).

GuestOS

- **Guest customization supports RHEL NetworkManager keyfile format:** In vSphere 8.0 Update 3, guest customization adds support for RHEL NetworkManager keyfile format and you can store network configuration in both keyfile and ifcfg format.

Drivers/Network

- **Support for Fibre Channel Extended Link Services (FC-ELS):** With vSphere 8.0 Update 3, you can use the command `esxcli storage fpin info set -e= <true/false>` to activate or deactivate the Fabric Performance Impact Notification (FPIN). The command saves the FPIN activation to both ConfigStore and the VMkernel System Interface Shell and persists across ESXi reboots. This is enabled by both Broadcom's `lpfc` and Marvell's `qlnativefc` drivers.
- **Unified Enhanced Networking Stack (UENS) driver for the Elastic Network Adapter (ENA):** vSphere 8.0 Update 3 adds a UENS driver for ENA, which provides connectivity to the AWS underlay Virtual Private Cloud (VPC).
- **Overlay Filters Supporting Tunnel End Point (TEP):** vSphere 8.0 Update 3 enhances the `i40en`, `qedentv`, and `sfvmk` NIC drivers by adding capabilities that expose overlay filters, supporting the TEP functionality.
- **Broadcom Driver Updates:**
 - **Broadcom `bnxtnet` driver:** Adds support for NetQ RSS to facilitate future Unified Enhanced Networking Stack (UENS) support.
- **Intel Driver Updates:**
 - **Intel `i40en` driver:** The number of queues per RSS engine is up from 4 to 8, in ENS mode it supports up to 16 queues.
 - **Intel `icen` driver:** Adds RoCE support on 1000GbE NIC.
 - **Intel `irdman` driver:** Adds RoCE support on 1000GbE NIC.
- **Marvell Driver Updates:**
 - **Marvell `qedentv` driver:** The max number of queues for the Default RSS engine is up from 4 to 16 in ENS mode.
- **Mellanox Driver Updates:**
 - **Mellanox `nmlx5` driver:** Adds support for hardware Large Receive Offload (LRO) in Enhanced Data Path mode to increase inbound throughput of high-bandwidth network connections by reducing CPU overhead.
 - Adds support for dual-DPU servers.
- **Pensando Driver Updates:**
 - **Pensando `ionic_en` driver:** Adds support for dual-DPU servers.
- **Routine Inbox Driver Updates and Bug Fixes**
 - Broadcom `bcm_mpi3`
 - Broadcom `lpfc`
 - Cisco `nfnc`
 - Marvell `qlnativefc`
 - Microchip `smartpqi`

CPU

- **PCIe hot plug is updated for server platforms utilizing newer generation AMD Genoa and Intel Sapphire Rapids CPUs:** Starting with vSphere 8.0 Update 3, kernel hot plug is supported for newer generation CPUs such as AMD Genoa and Intel Sapphire Rapids.
- **Support for Intel Xeon Max Series processors with integrated High Bandwidth Memory (HBM):** vSphere 8.0 Update 3 adds support for Intel Xeon Max Series processors (formerly with code name Sapphire Rapids HBM) with 64 GB of integrated HBM, aimed to enhance performance for workloads like high performance computing apps, artificial intelligence (AI), and machine learning (ML).
- **CPU C-State Power virtualization:** With vSphere 8.0 Update 3, for use cases such as Virtualized Radio Access Network (vRAN) workloads, you can configure and control the C-State power of the physical CPUs dedicated to vRAN VMs from the vSphere Client.
- **Cluster-wide option to retain virtual NUMA topology:** vSphere 8.0 Update 3 adds a setting to retain preconfigured vNUMA topology even if the VM moves, allowing for better NUMA topology tuning for VMs across all the hosts in the cluster. This is the advanced vCenter Server setting `VPXD_PersistVnuma` to keep the virtual NUMA topology on a cluster level under **Configure > Advanced Settings** in the vSphere Client.

Analytics and Metrics

- **vSphere green metrics with Running Average Power Limit (RAPL) technology:** Starting with vSphere 8.0 Update 3, in the **Advanced Performance Charts** in the vSphere Client you can see RAPL data on ESXi hosts that normally do not report individual subsystem power consumption, such as CPU and Memory, but only general host-level power consumption. With individual subsystem power consumption reports, you can plan on a more granular level to match your power and cooling budget.
- **Set VM log levels without powering off the VMs:** With vSphere 8.0 Update 3, you can set log levels between `VMW_LOG_TRIVIA` and `VMW_LOG_DEBUG_3` to avoid log spew in a healthy running VM without a power cycle by using the `SetLogLevel` service in the vAPI infrastructure.

GPU

- **Support for switching between Time Sliced and Multi-Instance GPU (MIG) modes for NVIDIA virtual GPUs:** Starting with vSphere 8.0 Update 3, you do not need to reboot an ESXi host to switch between time sliced and MIG modes for NVIDIA virtual GPUs. vGPU VMs can automatically set the correct device mode according to their vGPU type.
- **Zero-copy support for vGPUs to enhance vSphere vMotion and vSphere DRS tasks:** vSphere 8.0 Update 3 adds zero-copy support for vGPUs to enhance vSphere vMotion and vSphere DRS tasks by utilizing throughput of up to 100 Gbps.
- **Support for heterogeneous vGPU profiles on physical GPUs:** Starting with vSphere 8.0 Update 3, you can set vGPU profiles with different types or sizes on a single physical GPU to achieve greater flexibility with vGPU workloads and better utilization of GPU devices. For more information, see [Configuring vGPU Size](#).

vSphere Lifecycle Management

- **Support for parallel hardware and firmware upgrade with vSphere Lifecycle Manager:** With vCenter Server 8.0 Update 3, you can run parallel hardware and firmware remediation by using an integration between the vSphere Lifecycle Manager and the Hardware Support Manager.
- **VMware Photon™ 5.0 support for Update Manager Download Service (UMDS):** vSphere 8.0 Update 3 adds VMware Photon™ 5.0 to the supported Linux-based operating systems for installing UMDS. For more information, see [Installing UMDS](#).
- **Additional lifecycle management capabilities for standalone ESXi hosts:** Starting with vSphere 8.0 Update 3, you can add a standalone host to a data center or folder by importing an image from another ESXi host in the vCenter Server inventory or by using the current image on the host. For more information, see [Managing Standalone ESXi Hosts with vSphere Lifecycle Manager Images](#). When you move a host out of a cluster that you manage with a vSphere Lifecycle Manager image to a data center or a folder, the host becomes standalone and can retain the image from the cluster. For more information, see [Specifics of the Transitioning Workflow](#).

- **Lifecycle management of standalone ESXi hosts with VMware NSX:** Starting with vSphere 8.0 Update 3, NSX Manager and vSphere Lifecycle Manager work together to coordinate remediation of standalone ESXi hosts with VMware NSX. For more information, see [Using vSphere Lifecycle Manager Images for Standalone Hosts with NSX 4.2 and later](#).
- **Convert baseline-managed clusters to clusters managed by a single vSphere Lifecycle Manager image:** Starting with vSphere 8.0 Update 3, to start managing a cluster with a single vSphere Lifecycle Manager image, you can use the image installed on one of the ESXi hosts inside the cluster managed with baselines. For more information, see [Use an Image from a Host in the Cluster](#).
- **Patch VMX-related security vulnerabilities without any disruption to your workloads:** Starting with vSphere 8.0 Update 3, you can use the **Live Patch** capability to apply VMX-related security patches and bug fixes to ESXi hosts in a cluster managed with a vSphere Lifecycle Manager image. A pre-check prior to remediation lists all suitable hosts in a cluster. After you activate **LivePatch** in the remediation settings, qualified hosts in a cluster do not require maintenance mode, or reboot, or VM migration during the update procedure.
- **Customize vSphere Lifecycle Manager desired state images:** Starting with vSphere 8.0 Update 3, you can remove the Host Client and VMware Tools components from the base image, remove unnecessary drivers from vendor add-ons and components, and override existing drivers in a desired image. For more information, see [Edit a vSphere Lifecycle Manager Image](#).
- **Support for dual DPUs with vSphere Lifecycle Manager:** Starting with vSphere 8.0 Update 3, you can install dual DPUs on ESXi hosts and use the vSphere Lifecycle Manager workflow to upgrade the dual DPU system.
- **Extended support for vSphere Configuration Profiles (VCP):** With vSphere 8.0 Update 3, you can use VCP with the following new capabilities:
 - **Support to baseline-managed clusters (formerly referred to as VUM clusters):** Having an image-managed cluster is no longer a prerequisite for using VCP. You can use VCP to configure either baseline-managed clusters or image-managed clusters.
 - **Support for vSphere Distributed Switch (VDS):** VCP is fully integrated with VDS and supports drift detection and remediation of VDS configurations at a cluster level.
 - **Firewall ruleset management:** You can manage custom firewall rules at a cluster level by using VCP.
 - **ESXi Lockdown Mode:** vSphere admins can use the VCP desired configuration document to enforce Lockdown Mode on all hosts in a cluster.
 - **Support for SNMP and PCI device configurations:** You can manage SNMP and PCI devices at a cluster level by using VCP.

For more information, see [Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level](#).

vSphere Client and vCenter

- **Warning on the maximum number of remote https connections for vCenter:** Starting with vSphere 8.0 Update 3, to prevent a vCenter system to become unresponsive due to exceeding the limit of 2048 https connections, you will see HTTP error code **503 Service Unavailable** and **x-envoy-local-overloaded: true** in the response headers.
- **Merging the vSAN Management SDK with the Python SDK for the VMware vSphere API:** Starting with vSphere 8.0 Update 3, the [vSAN Management SDK for Python](#) is integrated into the Python SDK for the VMware vSphere API ([pyVmomi](#)). From the Python Package Index ([PyPI](#)), you can download a single package to manage vSAN, vCenter, and ESXi. This integration streamlines the discovery and installation process and enables automated pipelines instead of the series of manual steps previously.
- **vCenter Universally Unique Identifier (VC_UUID) field in the vSphere Client:** With vSphere 8.0 Update 3, the property table under **VMs > Virtual Machines** in the vSphere Client includes a new column that contains the VC_UUID. This identifier is automatically assigned to every virtual machine within a vCenter instance. The VC_UUID field helps clarify the correlation between the VM's UUID displayed on the switch's fabric and the VM name in vCenter.
- **Default HTTP response compression:** vSphere 8.0 Update 3 adds support for HTTP response compression by default on the edge proxies of vCenter and ESXi hosts, management traffic between vCenter and ESXi hosts, and for outgoing HTTP requests on sidecar proxies. HTTP response compression reduces the HTTP traffic in a vCenter Server system, shrinks page load time and speeds up API operations. You can deactivate HTTP response

compression if required, but the capability does not require any changes in your environment and is backward compatible.

- **Remove restrictions on virtual machine operations from the vSphere Client:** Starting with vCenter Server 8.0 Update 3, in the vSphere Client under **Virtual Machine > Configure > Disabled Methods** you can remove restrictions on virtual machine operations such as migration. For more information, see VMware knowledge base article [2044369](#).
- **Unified Management and Automation API Sessions:** Starting with vCenter 8.0 Update 3, you can combine vSphere Management API (VIM) and vSphere Automation API (vAPI) sessions, effectively unifying authentication across SOAP and REST vCenter endpoints. For more information see [Unified Management and Automation Session](#).
- **Migration of SPBM, SMS, EAM, and VLSM APIs to HTTP/JSON-based wire protocol:** Starting with vCenter Server 8.0 Update 3, the following 4 SOAP/XML service interfaces migrate to a HTTP/JSON-based wire protocol, providing OpenAPI 3.0 specifications for each, and integration with the [REST API](#) documentation:
 - The Storage Policy (SPBM) API, which simplifies the task of matching available storage to virtual machines.
 - VMware vCenter Storage Monitoring Service (SMS), which facilitates access to all vCenter storage information associated with VMware vCenter servers.
 - vSphere ESX Agent Manager API (EAM), which gives access to the objects that manage, monitor, and control lifecycle operations in vSphere.
 - Virtual Storage Lifecycle Management (VSLM) API that you use to manage first class disks (FCD).
 For more information, see [What's New in vSphere Automation: vSphere APIs, JSON, OpenAPI](#).

For vSphere Quick Boot support in ESXi 8.0 Update 3, see [Understanding ESXi Quick Boot Compatibility](#).

ESXi 8.0 Update 3 adds support for vSphere Quick Boot to Marvel OCTEON Fusion CNF105xx drivers.

For the full list of supported servers and drivers, see the [VMware Compatibility Guide](#).

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0 Update 2c Release Notes](#)
- [VMware ESXi 8.0 Update 1d Release Notes](#)
- [VMware ESXi 8.0 Update 2b Release Notes](#)
- [VMware ESXi 8.0 Update 2 Release Notes](#)
- [VMware ESXi 8.0 Update 1c Release Notes](#)
- [VMware ESXi 8.0 Update 1a Release Notes](#)
- [VMware ESXi 8.0 Update 1 Release Notes](#)
- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Product Support Notices

Deprecation of support on legacy NVIDIA GPUs for 3D hardware acceleration with Virtual Shared Graphics Acceleration (vSGA): Starting with vSphere 8.0 Update 3, support for 3D hardware acceleration by using vSGA on legacy NVIDIA GPUs is deprecated and will be removed in a future vSphere release. For more information, see KB [377624](#).

Product Support Notices from GA:

- **Deprecation of APIs for vCenter for Windows to Linux migration:**
Starting with vSphere 8.0 Update 3, APIs from vSphere 6.7.x for migrating vCenter for Windows to Linux during upgrade are deprecated and will be removed in a future vSphere release.

- **Deprecation of VMkernel API (vmkapi) version 2.5:**

Starting with vSphere 8.0 Update 3, vmkapi version 2.5 is deprecated and will be removed in a future major release. This requires an update of any third-party component released for vSphere 6.7.

- **Removal of Integrated Windows Authentication (IWA):**

vSphere 8.0 Update 3 is the final release to support Integrated Windows Authentication. IWA was deprecated in vSphere 7.0 and will be removed in the next major release. To ensure continued secure access, migrate from IWA to Active Directory over LDAPS or to Identity Federation with Multi-Factor Authentication. For more information, see [vSphere Authentication with vCenter Single Sign-On](#) and [Deprecation of Integrated Windows Authentication](#).

- **Removal of VMware Enhanced Authentication Plug-in (EAP):** Starting with vSphere 8.0 Update 3, you cannot use EAP to log in to a vCenter system by using the vSphere Client. For more information, see [Removing the deprecated VMware Enhanced Authentication Plug-in \(EAP\) to address CVE-2024-22245 and CVE-2024-22250](#).

Deprecation of the vCenter Service Lifecycle Management API: vSphere 8.0 Update 3 is deprecating the use of vCenter Service Lifecycle Management (vmonapi service) API and the service is not active by default, you must manually activate it. The service will be removed in a future release. For more information, see VMware knowledge base article [80775](#).

- **Removal of the internal runtime option execInstalledOnly:** Starting with ESXi 8.0 Update 3, the internal runtime option that deactivates the security option execInstalledOnly is deprecated and will be removed in the next major release. The boot option execInstalledOnly, which helps protect hosts against ransomware attacks, will be activated on ESXi hosts by default in the next major release.

- **Deprecation of Storage DRS Load Balancer and Storage I/O Control (SIOC):** The Storage DRS (SDRS) I/O Load Balancer, SDRS I/O Reservations-based load balancer, and vSphere Storage I/O Control Components will be deprecated in a future vSphere release. Existing 8.x and 7.x releases will continue to support this functionality. The deprecation affects I/O latency-based load balancing and I/O reservations-based load balancing among datastores within a Storage DRS datastore cluster. In addition, enabling of SIOC on a datastore and setting of Reservations and Shares by using SPBM Storage policies are also being deprecated. Storage DRS Initial placement and load balancing based on space constraints and SPBM Storage Policy settings for limits are not affected by the deprecation.

- **Deprecation of vSphere Trust Authority (vTA):**

Starting with vSphere 8.0 Update 3, vSphere Trust Authority is deprecated. vSphere continues to offer advanced workload attestation in its baseline functionality.

- **Removal of Patch Manager APIs:** Patch Manager APIs are supported in vSphere 8.x, but support will discontinue in a future release of vSphere. Instead of Patch Manager APIs, you can use the latest vSphere APIs, documented in the [vSphere API automation reference guide](#).

- **Deprecation of locales:**

Beginning with the next major release, we will be reducing the number of supported localization languages. The three supported languages will be:

- Japanese
- Spanish
- French
- The following languages will no longer be supported:

Impact:

- Users who have been using the deprecated languages will no longer receive updates or support in these languages.

All user interfaces, help documentation, and customer support will be available only in English or in the three supported languages mentioned above.

- Italian, German, Brazilian Portuguese, Traditional Chinese, Korean, Simplified Chinese

- **Removal of vSphere Lifecycle Manager baselines:**

Support for managing clusters with vSphere Lifecycle Manager baselines and baseline groups (legacy vSphere Update Manager workflows) will drop in a future vSphere release. Instead of baselines and baseline groups, you can use vSphere Lifecycle Manager images to perform tasks on a standalone host or on a cluster level such as install a desired

ESXi version on all hosts in a cluster, install and update third-party software, update, and upgrade ESXi or firmware, generate recommendations, and use a recommended image for your cluster.

Patches Contained in This Release

VMware ESXi 8.0 Update 3

Build Details

VMware vSphere Hypervisor (ESXi ISO) image

Download Filename:	VMware-VMvisor-Installer-8.0U3-24022510.x86_64.iso
Build:	24022510
Download Size:	605.63 MB
SHA256 checksum:	05ce214069a3e23265881fbd7a949fb93a99aa13059c6cac31920196e97ba4a1
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

VMware vSphere Hypervisor (ESXi) Offline Bundle

Download Filename:	VMware-ESXi-8.0U3-24022510-depot.zip
Build:	24022510
Download Size:	599 MB
SHA256 checksum:	82e963b196c9fca6ae2e25de2bfd353e221bf86f041f561b3ef42815e90eb604
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi80U3-24022510	Enhancement	Important

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U3-24022510-standard
ESXi-8.0U3-24022510-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi_8.0.3-0.0.24022510	06/25/2024	General	Bugfix image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

Virtual Machine Management Issues

PR 3237088: A vSphere system with EPYC processors of the series 7002, 7Fx2, 7Hx2, and 7001 might become unresponsive after more than 1000 days of continuous uptime

Due to AMD erratum 1474, if the core-C6 (CC6) sleep state is active on your EPYC processors of the series 7002, 7Fx2, 7Hx2, and 7001, a core might fail to exit CC6 in about 1044 days after the last system hardware reset, which excludes reboots by using VMware QuickBoot. As a result, your vSphere system might become unresponsive.

This issue is resolved in this release. The fix automatically deactivates CC6 until the next hardware reset.

PR 3035321: When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

This issue is resolved in this release.

Security Features Issues

PR 3373364: The CIMHttpServer firewall ruleset remains active even when you stop the CIM service

The `CIMHttpServer` firewall ruleset, which is for internal port 5988, might remain active even when you stop the CIM service.

This issue is resolved in this release. The `CIMHttpServer` firewall ruleset is permanently deactivated.

Miscellaneous Issues

PR 3359237: If a NVMe/TCP driver fails to process an I/O command, an ESXi host might disconnect from a storage array

If your ESXi hosts connect to storage arrays and perform read/write operations by using NVMe/TCP connections, when a NVMe/TCP driver fails to process an I/O command, ESXi might cancel the process with an abort command. As a result, an ESXi host might disconnect from a storage array and end previous I/O queues.

This issue is resolved in this release.

PR 3297607: An ESXi host might fail with a purple diagnostic screen while connecting to an NVMe over TCP controller

Due to a rare issue in the nvmetcp driver, an ESXi host might fail with a purple diagnostic screen while connecting to an NVMe over TCP controller. In the backtrace, you see a warning such as `WARNING: NVMFDEV:xxxx Failed to connect controller.`

This issue is resolved in this release.

PR 3355070: ESXi hosts might become unresponsive after VMFS heap exhaustion due to the pointer block cache size

The VMFS heap copy of the pointer block cache might fill up the heap memory and cause memory allocation failures for other operations that require memory from the VMFS heap. As a consequence, you might see some ESXi hosts becoming unresponsive. The issue is more likely to occur if the `maxAddressableSpaceTB` parameter is set to a higher value than the default of 32TB, as the time for cache refresh increases.

This issue is resolved in this release.

Guest OS Issues

PR 3091866: You see I/O errors in the logs of Linux virtual machines

In rare cases, unmap requests from a Linux Guest OS might compete with other workflows running on a different host such as block allocation on the same resource cluster. As a result, you see I/O errors in the logs of Linux VMs.

This issue is resolved in this release.

Known Issues

Storage Issues

vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out

vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD

In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See [Reconciling Discrepancies in the Managed Virtual Disk Catalog](#).

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

Miscellaneous Issues

The irdman driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic

If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the irdman driver might fail. This issue is unlikely to occur, as the irdman driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

Workaround: Use ULPs with RC transport mode.

Networking Issues

Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1

The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

Workaround: None.

Capture of network packets by using the PacketCapture tool on ESXi does not work

Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in [Collecting network packets using the lightweight PacketCapture on ESXi](#).

Workaround: Use the pktcap-uw tool. For more information, see [Capture and Trace Network Packets by Using the pktcap-uw Utility](#).

Installation, Upgrade, and Migration Issues

Fresh installation or creating VMFS partitions on a Micron 7500 or Intel D5-P5336 NVMe drives might fail with a purple diagnostic screen

UNMAP commands enable ESXi hosts to release storage space that is mapped to data deleted from the host. In NVMe, the equivalent of UNMAP commands is a deallocate DSM request. Micron 7500 and Intel D5-P5336 devices advertise a very large value in one of the deallocate limit attributes, DMSRL, which is the maximum number of logical blocks in a single range for a Dataset Management command. This leads to an integer overflow when the ESXi unmap split code converts number of blocks to number of bytes, which in turn might cause a failure of either installation or VMFS creation. You see a purple diagnostics screen with an error such as `Exception 14 or corruption in dlmalloc`. The issue affects ESXi 8.0 Update 2 and later.

Workaround: For a fresh ESXi installation, first install ESXi 8.0, deactivate UNMAP for the affected disk by using the command `esxcli storage core device vaai status set -D 0 -d <device-id>` and then upgrade to 8.0 Update 3. To create VMFS partitions, deactivate UNMAP for the affected disk by using the command `esxcli storage core device vaai status set -D 0 -d <device-id>` and then create VMFS as usual. You can reactivate UNMAP after you create a VMFS datastore, but if you delete or create a partition, UNMAP must remain deactivated.

The Cancel option in an interactive ESXi installation might not work as expected

Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

Standard image profiles for ESXi 8.0 Update 3 show last modified date as release date

The **Release Date** field of the standard image profile for ESXi 8.0 Update 3 shows the **Last Modified Date** value. The issue is only applicable to the image profiles used in Auto Deploy or ESXCLI. Base images used in vSphere Lifecycle Manager workflows display the release date correctly. This issue has no functional impact. The side effect is that if you search for profiles by release date, the profile does not show with the actual release date.

Workaround: Search by release version, such as 8.0.3.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

You cannot update to ESXi 8.0 Update 2b by using esxcli software vib commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

```
Please use a supported method to upgrade ESXi.
```

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see [Upgrade or Update a Host with Image Profiles](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the ntg3 driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on ntg3 drivers of versions earlier than 4.1.3 or the tg3 driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a ntg3 driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

VMware ESXi 8.0 Update 2e Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0 Update 2e | 15 JUL 2025 | Build 24789317

Check for additions and updates to these release notes.

Important:

ESXi 8.0 Update 2e delivers fixes for CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, and CVE-2025-2884. If you do not plan to update your environment to ESXi 8.0 Update 3f (build # 24784735), use 8.0 Update 2e to update your ESXi hosts of version 8.0 Update 2d and earlier with these security fixes. The supported update path from 8.0 Update 2e is to ESXi 8.0 Update 3f or later.

CAUTION:

Updates from ESXi 8.0 Update 2e to 8.0 Update 3 might expose your vSphere system to security vulnerabilities, because this is considered a back-in-time update. For more information, see the [Product Interoperability Matrix](#).

What's New

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

Patches Contained in This Release

Patch for VMware ESXi 8.0 Update 2e

This release of ESXi 8.0 Update 2e delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U2e-24789317-depot.zip
Build:	24789317
Download Size:	610.5 MB
sha256checksum:	c4080c8b63191f409ec13af5deafcbd6676032667835e2a1dc8487b628af42ec
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.2-0.48.24789317	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.2-0.48.24789317	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.2-0.48.24789317	Security	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U2e-24789317	Security	Critical	Security fixes only

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U2e-24789317-standard
ESXi-8.0U2e-24789317-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi 8.0 U2e - 24789317	July 15 2025	Security	Security fixes only

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.2-0.48.24789317

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-base_8.0.2-0.48.24789317 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.48.24789317 - VMware_bootbank_esxio-base_8.0.2-0.48.24789317 - VMware_bootbank_bmcad_8.0.2-0.48.24789317 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.48.24789317 - VMware_bootbank_vsan_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner_8.0.2-0.48.24789317 - VMware_bootbank_cpu-microcode_8.0.2-0.48.24789317 - VMware_bootbank_vsanhealth_8.0.2-0.48.24789317 - VMware_bootbank_gc_8.0.2-0.48.24789317 - VMware_bootbank_esx-xserver_8.0.2-0.48.24789317 - VMware_bootbank_bmcad-esxio_8.0.2-0.48.24789317 - VMware_bootbank_vdfs_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.48.24789317 - VMware_bootbank_crx_8.0.2-0.48.24789317 - VMware_bootbank_vds-vsip_8.0.2-0.48.24789317 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.48.24789317 - VMware_bootbank_infravisor_8.0.2-0.48.24789317 - VMware_bootbank_trx_8.0.2-0.48.24789317 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.48.24789317

	- VMware_bootbank_esxio_8.0.2-0.48.24789317 - VMware_bootbank_native-misc-drivers_8.0.2-0.48.24789317 - VMware_bootbank_gc-esxio_8.0.2-0.48.24789317 - VMware_bootbank_clusterstore_8.0.2-0.48.24789317
PRs Fixed	N/A
CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

The ESXi and `esx-update` bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

This patch updates the `esx-base` VIB. Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `esx-dvfilter-generic-fastpath`, `esxio-base`, `bmcad`, `native-misc-drivers-esxio`, `vsan`, `esxio-combiner`, `cpu-microcode`, `vsanhealth`, `gc`, `esx-xserver`, `bmcad-esxio`, `vdfs`, `esxio-combiner-esxio`, `crx`, `vds-vsip`, `drivervm-gpu-base`, `infravisor`, `trx`, `esxio-dvfilter-generic-fastpath`, `esxio`, `native-misc-drivers`, `gc-esxio`, and `clusterstore`.

This patch resolves the following issues:

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

esx-update_8.0.2-0.48.24789317

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-update_8.0.2-0.48.24789317 - VMware_bootbank_loadesx_8.0.2-0.48.24789317
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesx` and `esx-update`.

esxio-update_8.0.2-0.48.24789317

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	- VMware_bootbank_esxio-update_8.0.2-0.48.24789317 - VMware_bootbank_loadesxio_8.0.2-0.48.24789317
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesxio` and `esxio-update`.

ESXi-8.0U2e-24789317-standard

Profile Name	ESXi-8.0U2e-24789317-standard
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 15, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-base_8.0.2-0.48.24789317 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.48.24789317 - VMware_bootbank_esxio-base_8.0.2-0.48.24789317 - VMware_bootbank_bmcad_8.0.2-0.48.24789317 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.48.24789317 - VMware_bootbank_vsan_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner_8.0.2-0.48.24789317 - VMware_bootbank_cpu-microcode_8.0.2-0.48.24789317 - VMware_bootbank_vsanhealth_8.0.2-0.48.24789317 - VMware_bootbank_gc_8.0.2-0.48.24789317 - VMware_bootbank_esx-xserver_8.0.2-0.48.24789317 - VMware_bootbank_bmcad-esxio_8.0.2-0.48.24789317 - VMware_bootbank_vdfs_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.48.24789317 - VMware_bootbank_crx_8.0.2-0.48.24789317 - VMware_bootbank_vds-vsip_8.0.2-0.48.24789317 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.48.24789317 - VMware_bootbank_infravisor_8.0.2-0.48.24789317 - VMware_bootbank_trx_8.0.2-0.48.24789317 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.48.24789317 - VMware_bootbank_esxio_8.0.2-0.48.24789317 - VMware_bootbank_native-misc-drivers_8.0.2-0.48.24789317 - VMware_bootbank_gc-esxio_8.0.2-0.48.24789317 - VMware_bootbank_clusterstore_8.0.2-0.48.24789317 - VMware_bootbank_esx-update_8.0.2-0.48.24789317 - VMware_bootbank_loadesx_8.0.2-0.48.24789317 - VMware_bootbank_esxio-update_8.0.2-0.48.24789317

	VMware_bootbank_loadesxio_8.0.2-0.48.24789317
PRs Fixed	N/A
Related CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

This patch updates the following issues:

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

ESXi-8.0U2e-24789317-no-tools

Profile Name	ESXi-8.0U2e-24789317-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	Broadcom, Inc.
Release Date	July 15, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-base_8.0.2-0.48.24789317 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.48.24789317 - VMware_bootbank_esxio-base_8.0.2-0.48.24789317 - VMware_bootbank_bmcad_8.0.2-0.48.24789317 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.48.24789317 - VMware_bootbank_vsan_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner_8.0.2-0.48.24789317 - VMware_bootbank_cpu-microcode_8.0.2-0.48.24789317 - VMware_bootbank_vsanhealth_8.0.2-0.48.24789317 - VMware_bootbank_gc_8.0.2-0.48.24789317 - VMware_bootbank_esx-xserver_8.0.2-0.48.24789317 - VMware_bootbank_bmcad-esxio_8.0.2-0.48.24789317 - VMware_bootbank_vdfs_8.0.2-0.48.24789317 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.48.24789317 - VMware_bootbank_crx_8.0.2-0.48.24789317 - VMware_bootbank_vds-vsip_8.0.2-0.48.24789317 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.48.24789317 - VMware_bootbank_infravisor_8.0.2-0.48.24789317 - VMware_bootbank_trx_8.0.2-0.48.24789317 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.48.24789317 - VMware_bootbank_esxio_8.0.2-0.48.24789317 - VMware_bootbank_native-misc-drivers_8.0.2-0.48.24789317 - VMware_bootbank_gc-esxio_8.0.2-0.48.24789317 - VMware_bootbank_clusterstore_8.0.2-0.48.24789317

	- VMware_bootbank_esx-update_8.0.2-0.48.24789317 - VMware_bootbank_loadesx_8.0.2-0.48.24789317 - VMware_bootbank_esxio-update_8.0.2-0.48.24789317 - VMware_bootbank_loadesxio_8.0.2-0.48.24789317
PRs Fixed	N/A
Related CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

This patch updates the following issues:

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, and CVE-2025-41239. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0013](#).
- This release resolves [CVE-2025-2884](#). Broadcom has evaluated the severity of this issue to be in the moderate severity range.

ESXi 8.0 U2e - 24789317

Name	ESXi
Version	ESXi 8.0 U2e - 24789317
Release Date	July 15, 2025
Category	Security
Affected Components	- ESXi Component - ESXi Install/Upgrade Component
PRs Fixed	N/A
Related CVE numbers	CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, CVE-2025-2884

This patch resolves the issues listed in ESXi_8.0.2-0.48.24789317.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
```

```
Validation of vmknics failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed in the backtrace. The snapshot creation task completes upon retry.`

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') . The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.`

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 2d Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0 Update 2d | 04 MAR 2025 | Build 24585300

Check for additions and updates to these release notes.

Important:

ESXi 8.0 Update 2d delivers fixes for CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. If you do not plan to update your environment to ESXi 8.0 Update 3d (build # 24585383), use 8.0 Update 2d to update your ESXi hosts of version 8.0 Update 2c (build # 23825572) and earlier with these security fixes. The supported update path from 8.0 Update 2d is to ESXi 8.0 Update 3d or later.

CAUTION:

Updates from ESXi 8.0 Update 2d to 8.0 Update 3 might expose your vSphere system to security vulnerabilities, because this is considered a back-in-time update. For more information, see the [Product Interoperability Matrix](#).

What's New

This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

Patches Contained in This Release

Patch for VMware ESXi 8.0 Update 2d

This release of ESXi 8.0 Update 2d delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U2d-24585300-depot.zip
Build:	24585300
Download Size:	610.3 MB
sha256checksum:	225843151c6858279da8956107917db187948d5ff6af1964314c8d ff2c3b6d5f

Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.2-0.45.24585300	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.2-0.45.24585300	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.2-0.45.24585300	Security	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U2d-24585300	Security	Critical	Security fixes only

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U2d-24585300-standard
ESXi-8.0U2d-24585300-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi 8.0 U2d - 24585300	March 04 2025	Security	Security fixes only

For information about the individual components and bulletins, see the [Resolved Issues](#) section.

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues**ESXi_8.0.2-0.45.24585300**

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-base_8.0.2-0.45.24585300 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.45.24585300 - VMware_bootbank_esxio-base_8.0.2-0.45.24585300 - VMware_bootbank_bmcad_8.0.2-0.45.24585300 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.45.24585300 - VMware_bootbank_vsan_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner_8.0.2-0.45.24585300 - VMware_bootbank_cpu-microcode_8.0.2-0.45.24585300 - VMware_bootbank_vsanhealth_8.0.2-0.45.24585300 - VMware_bootbank_gc_8.0.2-0.45.24585300 - VMware_bootbank_esx-xserver_8.0.2-0.45.24585300 - VMware_bootbank_bmcad-esxio_8.0.2-0.45.24585300 - VMware_bootbank_vdfs_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.45.24585300 - VMware_bootbank_crx_8.0.2-0.45.24585300 - VMware_bootbank_vds-vsip_8.0.2-0.45.24585300 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.45.24585300 - VMware_bootbank_infravisor_8.0.2-0.45.24585300 - VMware_bootbank_trx_8.0.2-0.45.24585300 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.45.24585300 - VMware_bootbank_esxio_8.0.2-0.45.24585300 - VMware_bootbank_native-misc-drivers_8.0.2-0.45.24585300 - VMware_bootbank_gc-esxio_8.0.2-0.45.24585300 - VMware_bootbank_clusterstore_8.0.2-0.45.24585300
PRs Fixed	N/A
CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch updates the `esx-base` VIB. Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `esx-dvfilter-generic-fastpath`, `esxio-base`, `bmcad`, `native-misc-drivers-esxio`, `vsan`, `esxio-combiner`, `cpu-microcode`, `vsanhealth`, `gc`, `esx-xserver`, `bmcad-esxio`, `vdfs`, `esxio-combiner-esxio`, `crx`, `vds-vsip`, `drivervm-gpu-base`, `infravisor`, `trx`, `esxio-dvfilter-generic-fastpath`, `esxio`, `native-misc-drivers`, `gc-esxio`, and `clusterstore`. This patch resolves the following issue:

This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

esx-update_8.0.2-0.45.24585300

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-update_8.0.2-0.45.24585300 - VMware_bootbank_loadesx_8.0.2-0.45.24585300
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesx` and `esx-update`.

esxio-update_8.0.2-0.45.24585300

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-update_8.0.2-0.45.24585300 - VMware_bootbank_loadesxio_8.0.2-0.45.24585300
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesxio` and `esxio-update`.

ESXi-8.0U2d-24585300-standard

Profile Name	ESXi-8.0U2d-24585300-standard
Build	For build information, see Patches Contained in This Release .
Vendor	Broadcom, Inc.
Release Date	March 04, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	VMware_bootbank_esx-base_8.0.2-0.45.24585300

	- VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.45.24585300 - VMware_bootbank_esxio-base_8.0.2-0.45.24585300 - VMware_bootbank_bmcal_8.0.2-0.45.24585300 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.45.24585300 - VMware_bootbank_vsan_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner_8.0.2-0.45.24585300 - VMware_bootbank_cpu-microcode_8.0.2-0.45.24585300 - VMware_bootbank_vsanhealth_8.0.2-0.45.24585300 - VMware_bootbank_gc_8.0.2-0.45.24585300 - VMware_bootbank_esx-xserver_8.0.2-0.45.24585300 - VMware_bootbank_bmcal-esxio_8.0.2-0.45.24585300 - VMware_bootbank_vdfs_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.45.24585300 - VMware_bootbank_crx_8.0.2-0.45.24585300 - VMware_bootbank_vds-vsip_8.0.2-0.45.24585300 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.45.24585300 - VMware_bootbank_infravisor_8.0.2-0.45.24585300 - VMware_bootbank_trx_8.0.2-0.45.24585300 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.45.24585300 - VMware_bootbank_esxio_8.0.2-0.45.24585300 - VMware_bootbank_native-misc-drivers_8.0.2-0.45.24585300 - VMware_bootbank_gc-esxio_8.0.2-0.45.24585300 - VMware_bootbank_clusterstore_8.0.2-0.45.24585300 - VMware_bootbank_esx-update_8.0.2-0.45.24585300 - VMware_bootbank_loadesx_8.0.2-0.45.24585300 - VMware_bootbank_esxio-update_8.0.2-0.45.24585300 - VMware_bootbank_loadesxio_8.0.2-0.45.24585300
PRs Fixed	N/A
Related CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch updates the following issue:

This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

ESXi-8.0U2d-24585300-no-tools

Profile Name	ESXi-8.0U2d-24585300-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	Broadcom, Inc.
Release Date	March 04, 2025
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	- VMware_bootbank_esx-base_8.0.2-0.45.24585300 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.45.24585300 - VMware_bootbank_esxio-base_8.0.2-0.45.24585300 - VMware_bootbank_bmcad_8.0.2-0.45.24585300 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.45.24585300 - VMware_bootbank_vsan_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner_8.0.2-0.45.24585300 - VMware_bootbank_cpu-microcode_8.0.2-0.45.24585300 - VMware_bootbank_vsanhealth_8.0.2-0.45.24585300 - VMware_bootbank_gc_8.0.2-0.45.24585300 - VMware_bootbank_esx-xserver_8.0.2-0.45.24585300 - VMware_bootbank_bmcad-esxio_8.0.2-0.45.24585300 - VMware_bootbank_vdfs_8.0.2-0.45.24585300 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.45.24585300 - VMware_bootbank_crx_8.0.2-0.45.24585300 - VMware_bootbank_vds-vsip_8.0.2-0.45.24585300 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.45.24585300 - VMware_bootbank_infravisor_8.0.2-0.45.24585300 - VMware_bootbank_trx_8.0.2-0.45.24585300 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.45.24585300 - VMware_bootbank_esxio_8.0.2-0.45.24585300 - VMware_bootbank_native-misc-drivers_8.0.2-0.45.24585300 - VMware_bootbank_gc-esxio_8.0.2-0.45.24585300 - VMware_bootbank_clusterstore_8.0.2-0.45.24585300 - VMware_bootbank_esx-update_8.0.2-0.45.24585300 - VMware_bootbank_loadesx_8.0.2-0.45.24585300 - VMware_bootbank_esxio-update_8.0.2-0.45.24585300 - VMware_bootbank_loadesxio_8.0.2-0.45.24585300
PRs Fixed	N/A
Related CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch updates the following issue:

This release resolves CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0004](#).

ESXi 8.0 U2d - 24585300

Name	ESXi
Version	ESXi 8.0 U2d - 24585300
Release Date	March 04, 2025
Category	Bugfix
Affected Components	- ESXi Component - ESXi Install/Upgrade Component

PRs Fixed	N/A
Related CVE numbers	CVE-2025-22224, CVE-2025-22225, CVE-2025-22226

This patch resolves the issues listed in [ESXi_8.0.2-0.45.24585300](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS

certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
```

```
Validation of vmknics failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might

see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade.`

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console.` The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 2c Release Notes

This document contains the following sections

- [Introduction](#)
- [Earlier Releases of ESXi 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0 Update 2c | 21 MAY 2024 | Build 23825572
Check for additions and updates to these release notes.

This [release](#) is applicable only for ESXi hosts with vSAN. For more information, see the [Resolved Issues](#) section.

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

[VMware ESXi 8.0 Update 1d Release Notes](#)

[VMware ESXi 8.0 Update 2b Release Notes](#)

[VMware ESXi 8.0 Update 2 Release Notes](#)

[VMware ESXi 8.0 Update 1c Release Notes](#)

[VMware ESXi 8.0 Update 1a Release Notes](#)

[VMware ESXi 8.0 Update 1 Release Notes](#)

[VMware ESXi 8.0c Release Notes](#)

[VMware ESXi 8.0b Release Notes](#)

[VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Patches Contained in This Release

Patch for VMware ESXi 8.0 Update 2c

This release of ESXi 8.0 Update 2c delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U2c-23825572-depot.zip
Build:	23825572
Download Size:	611.1 MB
sha256checksum:	e1b4988a49165f354c668bc4e39334e248f6987f22a09fec7cc2c2d000d9a407
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.2-0.40.23825572	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.2-0.40.23825572	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.2-0.40.23825572	Bugfix	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U2c-23825572	Bugfix	Critical	Bug fixes only

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U2c-23825572-standard
ESXi-8.0U2c-23825572-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi 8.0 U2c - 23825572	05/21/2024	Bugfix	Bug fixes only

For information about the individual components and bulletins, see the [Resolved Issues](#) section.

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.2-0.40.23825572

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.40.23825572 - VMware_bootbank_esxio_8.0.2-0.40.23825572 - VMware_bootbank_native-misc-drivers_8.0.2-0.40.23825572 - VMware_bootbank_vds-vsip_8.0.2-0.40.23825572 - VMware_bootbank_trx_8.0.2-0.40.23825572 - VMware_bootbank_bmcad-esxio_8.0.2-0.40.23825572 - VMware_bootbank_bmcad_8.0.2-0.40.23825572 - VMware_bootbank_vsanhealth_8.0.2-0.40.23825572 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.40.23825572 - VMware_bootbank_gc_8.0.2-0.40.23825572 - VMware_bootbank_cpu-microcode_8.0.2-0.40.23825572 - VMware_bootbank_vdfs_8.0.2-0.40.23825572 - VMware_bootbank_esxio-base_8.0.2-0.40.23825572 - VMware_bootbank_esxio-combiner_8.0.2-0.40.23825572 - VMware_bootbank_gc-esxio_8.0.2-0.40.23825572 - VMware_bootbank_crx_8.0.2-0.40.23825572 - VMware_bootbank_vsan_8.0.2-0.40.23825572 - VMware_bootbank_esx-xserver_8.0.2-0.40.23825572

	- VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.40.23825572 - VMware_bootbank_esx-base_8.0.2-0.40.23825572 - VMware_bootbank_clusterstore_8.0.2-0.40.23825572 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.40.23825572 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.40.23825572 - VMware_bootbank_infravisor_8.0.2-0.40.23825572
PRs Fixed	3361939
CVE numbers	N/A

This patch updates the `esx-base` VIB. Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `native-misc-drivers-esxio`, `esxio`, `native-misc-drivers`, `vds-vsip`, `trx`, `bmcad-esxio`, `bmcad`, `vsanhealth`, `esxio-combiner-esxio`, `gc`, `cpu-microcode`, `vdfs`, `esxio-base`, `esxio-combiner`, `gc-esxio`, `crx`, `vsan`, `esx-xserver`, `esxio-dvfilter-generic-fastpath`, `clusterstore`, `gpu`, `infravisor`, and `esx-dvfilter-generic-fastpath`. This patch resolves the following issue:

PR 3381464: VMs might fail to boot up after a vSphere Storage vMotion migration from vSAN ESA as the source datastore

Under a rare condition, a virtual machine migration by using Storage vMotion from a vSAN ESA datatore as the source might result in data inconsistency of VMDKs at the destination datastore. As a result, the VMs might fail to boot up after the migration.

This issue is resolved in this release.

PR 3361939: Applications using AVX-512 instructions in a vSAN environment might report data consistency errors

If an application in a vSAN environment uses AVX-512 instructions, it might produce incorrect results or report data consistency errors, related to the data calculated by using such results.

This issue is resolved in this release.

esx-update_8.0.2-0.40.23825572

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_loadesx_8.0.2-0.40.23825572 - VMware_bootbank_esx-update_8.0.2-0.40.23825572
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesx` and `esx-update`.

esxio-update_8.0.2-0.40.23825572

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_loadesxio_8.0.2-0.40.23825572 - VMware_bootbank_esxio-update_8.0.2-0.40.23825572
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesxio` and `esxio-update`.

ESXi-8.0U2c-23825572-standard

Profile Name	ESXi-8.0U2c-23825572-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	May 21, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.40.23825572 - VMware_bootbank_esxio_8.0.2-0.40.23825572 - VMware_bootbank_native-misc-drivers_8.0.2-0.40.23825572 - VMware_bootbank_vds-vsip_8.0.2-0.40.23825572 - VMware_bootbank_trx_8.0.2-0.40.23825572 - VMware_bootbank_bmc-al-esxio_8.0.2-0.40.23825572 - VMware_bootbank_bmc-al_8.0.2-0.40.23825572 - VMware_bootbank_vsanhealth_8.0.2-0.40.23825572 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.40.23825572 - VMware_bootbank_gc_8.0.2-0.40.23825572 - VMware_bootbank_cpu-microcode_8.0.2-0.40.23825572 - VMware_bootbank_vdfs_8.0.2-0.40.23825572 - VMware_bootbank_esxio-base_8.0.2-0.40.23825572 - VMware_bootbank_esxio-combiner_8.0.2-0.40.23825572 - VMware_bootbank_gc-esxio_8.0.2-0.40.23825572 - VMware_bootbank_crx_8.0.2-0.40.23825572 - VMware_bootbank_vsan_8.0.2-0.40.23825572 - VMware_bootbank_esx-xserver_8.0.2-0.40.23825572 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.40.23825572 - VMware_bootbank_esx-base_8.0.2-0.40.23825572

	- VMware_bootbank_clusterstore_8.0.2-0.40.23825572 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.40.23825572 - VMware_bootbank_loadesx_8.0.2-0.40.23825572 - VMware_bootbank_esx-update_8.0.2-0.40.23825572 - VMware_bootbank_loadesxio_8.0.2-0.40.23825572 - VMware_bootbank_esxio-update_8.0.2-0.40.23825572 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.40.23825572 - VMware_bootbank_infravisor_8.0.2-0.40.23825572
PRs Fixed	3361939
Related CVE numbers	N/A

This patch updates the following issue:

PR 3381464: VMs might fail to boot up after a vSphere Storage vMotion migration from vSAN ESA as the source datastore

Under a rare condition, a virtual machine migration by using Storage vMotion from a vSAN ESA datastore as the source might result in data inconsistency of VMDKs at the destination datastore. As a result, the VMs might fail to boot up after the migration.

This issue is resolved in this release.

PR 3361939: Applications using AVX-512 instructions in a vSAN environment might report data consistency errors

If an application in a vSAN environment uses AVX-512 instructions, it might produce incorrect results or report data consistency errors, related to the data calculated by using such results.

This issue is resolved in this release.

ESXi-8.0U2c-23825572-no-tools

Profile Name	ESXi-8.0U2c-23825572-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	May 21, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.40.23825572 - VMware_bootbank_esxio_8.0.2-0.40.23825572 - VMware_bootbank_native-misc-drivers_8.0.2-0.40.23825572 - VMware_bootbank_vds-vsip_8.0.2-0.40.23825572 - VMware_bootbank_trx_8.0.2-0.40.23825572 - VMware_bootbank_bmcal-esxio_8.0.2-0.40.23825572 - VMware_bootbank_bmcal_8.0.2-0.40.23825572 - VMware_bootbank_vsanhealth_8.0.2-0.40.23825572 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.40.23825572 - VMware_bootbank_gc_8.0.2-0.40.23825572 - VMware_bootbank_cpu-microcode_8.0.2-0.40.23825572

	• VMware_bootbank_vdfs_8.0.2-0.40.23825572 • VMware_bootbank_esxio-base_8.0.2-0.40.23825572 • VMware_bootbank_esxio-combiner_8.0.2-0.40.23825572 • VMware_bootbank_gc-esxio_8.0.2-0.40.23825572 • VMware_bootbank_crx_8.0.2-0.40.23825572 • VMware_bootbank_vsan_8.0.2-0.40.23825572 • VMware_bootbank_esx-xserver_8.0.2-0.40.23825572 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.40.23825572 • VMware_bootbank_esx-base_8.0.2-0.40.23825572 • VMware_bootbank_clusterstore_8.0.2-0.40.23825572 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.40.23825572 • VMware_bootbank_loadesx_8.0.2-0.40.23825572 • VMware_bootbank_esx-update_8.0.2-0.40.23825572 • VMware_bootbank_loadesxio_8.0.2-0.40.23825572 • VMware_bootbank_esxio-update_8.0.2-0.40.23825572 • VMware_bootbank_drivervm-gpu-base_8.0.2-0.40.23825572 • VMware_bootbank_infravisor_8.0.2-0.40.23825572
PRs Fixed	3361939
Related CVE numbers	N/A

This patch updates the following issue:

PR 3381464: VMs might fail to boot up after a vSphere Storage vMotion migration from vSAN ESA as the source datastore

Under a rare condition, a virtual machine migration by using Storage vMotion from a vSAN ESA datatore as the source might result in data inconsistency of VMDKs at the destination datastore. As a result, the VMs might fail to boot up after the migration.

This issue is resolved in this release.

PR 3361939: Applications using AVX-512 instructions in a vSAN environment might report data consistency errors

If an application in a vSAN environment uses AVX-512 instructions, it might produce incorrect results or report data consistency errors, related to the data calculated by using such results.

This issue is resolved in this release.

ESXi 8.0 U2c - 23825572

Name	ESXi
Version	ESXi80U2c-23825572
Release Date	May 21, 2024
Category	Bugfix
Affected Components	• ESXi Component • ESXi Install/Upgrade Component
PRs Fixed	3361939
Related CVE numbers	N/A

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some nmlx5_core driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliance` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
```

```
Validation of vmknics failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the `ntg3` driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 2b Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Product Support Notices](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 2b | 29 FEB 2024 | Build 23305546
Check for additions and updates to these release notes.

What's New

Solution License for VMware Cloud Foundation (VCF) and VMware vSphere Foundation (VVF):

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VVF and VCF.

For more information on applying the Solution License to VVF components, see [VMware vSphere Foundation \(VVF\) Licensing](#).

For more information on VCF Solution License, see [VMware Cloud Foundation 5.1.1 Release Notes](#).

100 GiB of included storage capacity per licensed core:

Starting with vSphere 8.0 Update 2b, as part of the VMware vSphere Foundation Solution License, you can use up to 100 gibibytes (GiB) of included vSAN storage per host licensed core. For capacity larger than 100 GiB per core, you must purchase vSAN capacity per terabyte (TiB) and apply a vSAN license key that reflects the total raw storage capacity of the vSAN cluster. For more information on capacity reporting and licensing in vSAN, see [Demystifying Capacity Reporting in vSAN](#) and [Counting Cores for VMware Cloud Foundation and vSphere Foundation and TiBs for vSAN](#).

Dynamic firewall rules when configuring syslog collectors with non-standard ports:

With ESXi 8.0 Update 2b, when you configure syslog remote hosts, or loghosts, with non-standard ports, the vmsyslogd service automatically creates persistent dynamic firewall rules. You no longer need to manually open the firewall at ports that are different from the default 514 for TCP/UDP and 1514 for SSL protocols respectively. When configuring remote hosts with the standard ports, you still need to enable the syslog firewall ruleset.

ESXi 8.0 Update 2b adds support to vSphere Quick Boot for multiple servers, including:

- Dell
 - MX760c vSAN Ready Node
 - PowerEdge C6615
 - PowerEdge XR8610t
 - R7625 vSAN Ready Node
 - VxRail VP-760xa
 - VxRail VS-760
- HPE
 - Alletra Storage Server 4110
 - Alletra Storage Server 4140
 - HPE Cray XD670
 - ProLiant DL110 Gen11
 - ProLiant DL20 Gen11
 - ProLiant ML30 Gen11
- Lenovo
 - DN8848 V2

For the full list of supported servers, see the [VMware Compatibility Guide](#).

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0 Update 2 Release Notes](#)
- [VMware ESXi 8.0 Update 1c Release Notes](#)
- [VMware ESXi 8.0 Update 1a Release Notes](#)
- [VMware ESXi 8.0 Update 1 Release Notes](#)
- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Product Support Notices

- **Reduction in the memory usage of the ps command:** In ESXi 8.0 Update 2b, an optimization of memory usage for the `ps` command reduces memory required to run the command by up to 96%.

Patches Contained in This Release

VMware ESXi 8.0 Update 2b

Build Details

Download Filename:	VMware-ESXi-8.0U2b-23305546-depot.zip
Build:	23305546
Download Size:	987.6 MB
sha256checksum:	264ebd69044a85fc4ade78ff088f5d4fa9af3cbd23eca241f0855cb7bdeb8258
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U2b-23305546	Bugfix	Critical	Security fixes and Bug fixes
ESXi80U2sb-23305545	Security	Critical	Security only fixes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.2-0.30.23305546	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.2-0.30.23305546	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.2-0.30.23305546	Bugfix	Critical
Host Based Replication Server for ESX	VMware-HBR-UW_8.0.2-0.30.23305546	Bugfix	Critical
ESXi Component - core ESXi VIBs	ESXi_8.0.2-0.25.23305545	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.2-0.25.23305545	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.2-0.25.23305545	Security	Critical
ESXi Tools Component	VMware-VM-Tools_12.3.5.2254099-23305545	Security	Critical

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U2b-23305546-standard
ESXi-8.0U2b-23305546-no-tools
ESXi-8.0U2sb-23305545-standard
ESXi-8.0U2sb-23305545-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi-8.0U2b-23305546	02/29/2024	General	Security and Bugfix image
ESXi-8.0U2sb-23305545	02/29/2024	Security	Security only image

For information about the individual components and bulletins, see the [Resolved Issues](#) section.

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.2-0.30.23305546

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_vds-vsip_8.0.2-0.30.23305546 - VMware_bootbank_crx_8.0.2-0.30.23305546 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.30.23305546 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.30.23305546 - VMware_bootbank_vsanhealth_8.0.2-0.30.23305546 - VMware_bootbank_gc-esxio_8.0.2-0.30.23305546 - VMware_bootbank_clusterstore_8.0.2-0.30.23305546 - VMware_bootbank_trx_8.0.2-0.30.23305546 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.30.23305546

	- VMware_bootbank_cpu-microcode_8.0.2-0.30.23305546 - VMware_bootbank_infravisor_8.0.2-0.30.23305546 - VMware_bootbank_esx-base_8.0.2-0.30.23305546 - VMware_bootbank_vsan_8.0.2-0.30.23305546 - VMware_bootbank_bmcad_8.0.2-0.30.23305546 - VMware_bootbank_esxio_8.0.2-0.30.23305546 - VMW_bootbank_pensandoatlas_1.46.0.E.28.1.314-2vmw.802.0.0.22939414 - VMware_bootbank_gc_8.0.2-0.30.23305546 - VMware_bootbank_vdfs_8.0.2-0.30.23305546 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.30.23305546 - VMware_bootbank_esxio-combiner_8.0.2-0.30.23305546 - VMware_bootbank_native-misc-drivers_8.0.2-0.30.23305546 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.30.23305546 - VMware_bootbank_esx-xserver_8.0.2-0.30.23305546 - VMware_bootbank_bmcad-esxio_8.0.2-0.30.23305546 - VMware_bootbank_esxio-base_8.0.2-0.30.23305546
PRs Fixed	3326852, 3326624, 3319724, 3316967, 3307362, 3318781, 3331997, 3317441, 3317938, 3336882, 3317707, 3316767, 3306552, 3210610, 3262965, 3312008, 3304392, 3308819, 3313089, 3313912, 3313572, 3308844, 3298815, 3312006, 3311820, 3308054, 3308132, 3302896, 3308133, 3312713, 3301092, 3296088, 3311524, 3311123, 3307579, 3308347, 3309887, 3308458, 3309849, 3303232, 3311989, 3308023, 3308025, 3301665, 3302759, 3305165, 3304472, 3303651, 3293470, 3303807, 3307435, 3282049, 3293466, 3293464, 3296047, 3300173, 3294828, 3300031, 3295243, 3268113, 3284909, 3256705, 3292627, 3287767, 3289239, 3297899, 3295250, 3287747, 3222955, 3300326, 3292550, 3300705, 3287772, 3294624, 3291105, 3296694, 3291219, 3280469, 3288307, 3287769, 3287762, 3289397, 3293199, 3285514, 3275379, 3288872
CVE numbers	N/A

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

Updates the `bootbank_vds-vsip`, `bootbank_crx`, `bootbank_esxio-combiner-esxio`, `bootbank_esxio-dvfilter-generic-fastpath`, `bootbank_vsanhealth`, `bootbank_gc-esxio`, `bootbank_clusterstore`, `bootbank_trx`, `bootbank_native-misc-drivers-esxio`, `bootbank_cpu-microcode`, `bootbank_infravisor`, `bootbank_esx-base`, `bootbank_vsan`, `bootbank_bmcad`, `bootbank_esxio`, `VMW_bootbank_pensandoatlas`, `bootbank_gc`, `bootbank_vdfs`, `bootbank_drivervm-gpu-base`, `bootbank_esxio-combiner`, `bootbank_native-misc-drivers`, `bootbank_esx-dvfilter-generic-fastpath`, `bootbank_esx-xserver`, `bootbank_bmcad-esxio`, and `bootbank_esxio-base` VIBs to resolve the following issues:

PR 3303807: ESXi hosts with active Trusted Platform Module (TPM) encryption might fail with a purple diagnostic screen when using vSphere Quick Boot for upgrade to ESXi 8.0 Update 1 or later

If you use vSphere Quick Boot to upgrade ESXi hosts with active TPM encryption to ESXi 8.0 Update 1 or later, such hosts might fail with a purple diagnostic screen and a message similar to `security violation was detected`.

This issue is resolved in this release. Upgrade of ESXi hosts with TPM encryption and vSphere Quick Boot is deactivated for ESXi 8.0 Update 2b.

PR 3316767: VM I/Os might prematurely fail or ESXi hosts crash while HPP claims a local multipath SCSI device

Due to a rare issue with the failover handling in the High Performance Plug-in (HPP) module, if HPP claims a local multipath SCSI device in active-active mode, you might see the following issues:

- in case of connection loss or a NO CONNECT error, VM I/Os can fail prematurely instead of retrying on another path.

- in case of all paths down (APD) state, if a complete storage rescan is triggered from vCenter or the ESXi host, the host might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3302759: Pending I/Os might cause the ESXi host running a Primary VM to fail with a purple diagnostic screen during failover

During a failover of a Primary VM protected by vSphere Fault Tolerance, if some I/Os between the VM and the ESXi host are still unresolved, the host might intermittently fail with a purple diagnostic screen.

This issue is resolved in this release. The fix makes sure that no pending I/Os exist at the time an FT session info cleanup runs.

PR 3291105: While detaching an Enhanced Networking Stack (ENS) port from an ESXi host, the host might fail with a purple diagnostic screen

During deactivation of ENS in scenarios such as an upgrade or a VNIC configuration change, while detaching an ENS port from an ESXi host with active ENS, the port might be removed from the port list and not be available for further reference, which might cause an invalid memory access error. As a result, the ESXi host fails with a purple diagnostic screen.

This issue is resolved in this release.

PR 3297899: vSAN file server continues restarting

In case all AD servers are inaccessible, vSAN Skyline Health might show a file server that is always in restarting state, with repeated failovers from one host to another. If this occurs, all file shares on the file server, including NFS shares with `securityType=AUTH_SYS`, become inaccessible.

This issue is resolved in this release.

PR 3288307: Cannot upgrade vSAN File Service or enter maintenance mode

In rare cases, if the `hostd` service reboots, upgrade of the vSAN File Service or entering maintenance mode might be blocked.

This issue is resolved in this release.

PR 3294624: Skyline Health fails to load for file shares in an inextensible state

When a file share is in an inextensible state, the current storage policy cannot be fulfilled. The file share might also have other issues, such as exceeding the soft quota with its used capacity. Skyline Health fails to load for objects in this state.

This issue is resolved in this release.

PR 3293199: The file `/var/run/log/clusterAgent.stderr` on an ESXi host becomes extremely large, filling up the OSDATA partition

In case of certain rare error conditions, a component within the `clusterAgent` ESXi service might start periodically writing to `/var/run/log/clusterAgent.stderr` instead of the regular `syslog` destination on an ESXi host. Since this is not expected behavior, the `stderr` file is not rotated or otherwise monitored. If the error conditions persist over time, for example several months, the file can expand to several GB in size, filling up the OSDATA partition. When disk space is exhausted, ESXi services and VMs might lock up.

This issue is resolved in this release. The fix makes sure that if your environment had the issue, after upgrading to ESXi 8.0 Update 2b or later, the oversized log file is compressed. The file is not automatically deleted, but you can safely manually remove the compressed file.

PR 3319724: vSAN ESA VMs with large-sized VMDKs become inaccessible

In vSAN Express Storage Architecture (ESA) environments of version 8.0 and later, VMDK objects with a large amount of used capacity might overload the object scrubber, causing such objects to become inaccessible.

This issue is resolved in this release.

PR 3294828: The hostd service might intermittently fail due to an invalid sensor data record

In rare cases, the Baseboard Management Controller (BMC) might provide an invalid sensor data record to the hostd service and cause it to fail. As a result, ESXi hosts become unresponsive or disconnect from vCenter.

This issue is resolved in this release. The fix implements checks to validate the sensor data record before passing it to the hostd service.

PR 3287769: Unexpected failover of vSAN File Service containers

The EMM monitor can trigger unexpectedly and block the File Service EPiC heartbeat update to the arbitrator host. This problem can cause an unexpected container failover, which interrupts the I/O of SMB file shares.

This issue is resolved in this release.

PR 3293464: The VMX service might fail during migration of virtual machines with vSphere Virtual Volumes if the volume closure on the source takes long

In some scenarios, such as virtual machines with vSphere Virtual Volumes and Changed Block Tracking (CBT) or Content Based Read Cache (CBRC) enabled, during storage migration, flushing pending I/Os on volume closure at the source might take long, like 10 sec. If clearing the pending I/Os is not complete by the time the VMX service tries to reach the volumes on the destinations host, the service fails.

This issue is resolved in this release. The fix makes sure that all vSphere Virtual Volumes and disks are closed on the source host before the destination host tries to open them.

PR 3308819: You cannot add allowed IP addresses for an ESXi host

With ESXi 8.0 Update 2, some ESXi firewall rulesets, such as `dhcp`, are system-owned by default and prevent manual adding of allowed IP addresses to avoid possible break of service. With ESXi 8.0 Update 2b, you can manually add allowed IP addresses to all rulesets, except for `nfsClient`, `nfs41Client`, `trusted-infrastructure-kmxd`, `trusted-infrastructure-kmxa`, and `vsanEncryption`.

This issue is resolved in this release.

PR 3308461: When MAC learning is not active, newly added ports on a vSphere Distributed Switch (VDS) might go into blocked state

If MAC learning on a VDS is not active, when from the guest OS you change the MAC address of a virtual machine to a new proxy switch port, the newly added port might be in a blocked state due to a L2 violation error, even when the **Mac address changes option** is set to **Accept**. The issue occurs only when MAC learning on a VDS is not active.

This issue is resolved in this release.

PR 3332617: vGPU VMs with NVIDIA RTX 5000 Ada or NVIDIA RTX 6000 Ada graphics card fail to start after update to ESXi 8.0 Update 2

After an update to ESXi 8.0 Update 2, booting of a vGPU VM with NVIDIA RTX 5000 Ada or NVIDIA RTX 6000 Ada graphics card fails with an error such as `vmiop-display unable to reserve vgpu` in the logs. In the vSphere Client, you see an error such as `Module DevicePowerOnEarly power on failed..`

This issue is resolved in this release.

PR 3326624: An ESXi host might fail during upgrade due to an unmap I/O request larger than 2040 MB on a vSAN-backed object

During an upgrade to ESXi 8.0 Update 2, in rare cases, a guest virtual machine might issue an unmap I/O request larger than 2040 MB on a vSAN-backed object and the request might split in two parts. As a result, the ESXi host receives duplicate metadata entries and might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3304472: vSphere vMotion operations might fail due to a rare issue with no namespaces present in the namespacemgr.db file

When a virtual machine is configured with namespaces and for some reason the namespaces table in the namespacemgr.db file is empty in the ESXi host database, migration of such a VM by using vSphere vMotion might fail. In the vmware.log file, you see errors such as:

```
vmx - NamespaceMgrCheckpoint: No valid queue found while restoring the namespace events.
The migrate type is 1.
```

This issue is resolved in this release.

PR 3293470: In the vSphere Client, you see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore

In the vSphere Client, when you navigate to **Monitor > Advanced > Datastore** to see the advanced performance metrics of a virtual machine, you might see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore. The issue occurs due to incorrect calculation of the counters for the read/write IO latency as a cumulative average.

This issue is resolved in this release.

PR 3300031: The Network Time Protocol (NTP) peer and loop statistics are not automatically captured in the vm-support bundle

Prior to vSphere 8.0 Update 2b, unless you manually add the NTP peer and loop statistics to the vm-support bundle, NTP logs are not available for troubleshooting.

This issue is resolved in this release. The fix adds support for automatically capturing NTP peer and loop statistics from the /var/log/ntp/ directory in the vm-support bundle.

PR 3293466: vSphere vMotion operations between ESXi 7.x hosts fail with the error "Destination failed to preopen disks"

If on an ESXi 7.x host you use an older Virtual Volumes storage provider, also called a VASA provider, it might not allow binding multiple vSphere Virtual Volumes in batches to optimize the performance of migration and power on. As a result, in the vSphere Client you might see the error `Destination failed to preopen disks` when running vSphere vMotion operations between 7.x hosts.

This issue is resolved in this release.

PR 3316967: Changed Block Tracking (CBT) might not work as expected on a hot extended virtual disk

In vSphere 8.0 Update 2, to optimize the open and close process of virtual disks during hot extension, the disk remains open during hot extend operations. Due to this change, incremental backup of virtual disks with CBT enabled might be incomplete, because the CBT in-memory bitmap does not resize, and CBT cannot record the changes to the extended disk block. As a result, when you try to restore a VM from an incremental backup of virtual disks with CBT, the VM might fail to start.

This issue is resolved in this release.

PR 3317441: Enhanced Network Datapath (EDP) might reserve too much memory on a large vSphere system and cause resource issues

EDP reservation is linear to the total memory by reserving up to 1.6% of the total. For a VM with high memory reservation in a large vSphere system, such a memory buffer might be too big. For example, if the system has a 100TB of memory, the EDP reservation might be 1.6TB. As a result, you might see resource issues such as not being able to hot-add VNICs to VMs with high memory reservation.

This issue is resolved in this release. The fix sets the EDP reservation limit at 1% of the total memory and puts a 5GB cap on a single memory reservation. You can adjust the memory reservation limit by using the VMkernel boot option `ensMbufPoolMaxMinMB` or by using the following command: `esxcli system settings kernel set -s ensMbufPoolMaxMBPerGB -v <value-in-MB>`. In any case, you must reboot the ESXi host to enable the change.

PR 3275379: VBS-enabled Windows VMs might intermittently fail with a blue diagnostic screen on ESXi hosts running on AMD processors.

Windows virtual machines with virtualization-based security (VBS) enabled might intermittently fail with a blue diagnostic screen on ESXi hosts running on AMD processors. The BSOD has the following signature:

```
SECURE_KERNEL_ERROR (18b)The secure kernel has encountered a fatal error.
Arguments:
Arg1: 000000000000018c
Arg2: 000000000000100b
Arg3: 0000000000000000
Arg4: 000000000000409b
```

This issue is resolved in this release.

PR 3311820: Creation of VMFS volumes on local software emulated 4Kn devices fails

In some environments, after an update to ESXi 8.0 Update 2, you might not be able to create VMFS volumes on local software emulated 4Kn devices. The issue occurs when a SCSI driver fails to discover and register a 4Kn NVMe device as a 4Kn SWE device. In the vSphere On-disk Metadata Analyzer (VOMA), you see an error such as `ON-DISK ERROR: Invalid disk block size 512, should be 4096.`

This issue is resolved in this release.

PR 3296047: The vvolld service fails with Unable to Allocate Memory error

If a vSphere API for Storage Awareness (VASA) provider is down for a long time, frequent reconnection attempts from the vvolld service might exhaust its memory, because of repetitively creating and destroying session objects. The issue can also occur when an event such as a change in a VASA provider triggers an update to the vvolld service. The vvolld service restarts within seconds, but if the restart interrupts a running operation of the VASA provider, the operation might fail and affect related virtual machine operations as well. In the backtrace, you see an error such as `Unable To Allocate Memory.`

This issue is resolved in this release.

PR 3309887: You cannot use host profiles to configure a path selection policy for storage disks

In some cases, when you try to change or configure the path selection policy of storage disks by using a host profile, the operation fails. In the `syslog.log` file, you see lines such as

```
Errors: Unable to set SATP configuration: Invalid options provided for SATP VMW_SATP_ALUA
```

This issue is resolved in this release.

PR 3308844: The hostd service might become unresponsive after a restart due to a rare issue with Inter Process Communication (IPC)

If an IPC process has incorrect payload in the IPC packet, for example: an incorrect version in the header, a deadlock in the iSCSI daemon might cause the hostd service to become unresponsive after a hostd restart. As a result, you cannot perform any operations on the affected ESXi host and VMs by using the vSphere Client.

This issue is resolved in this release.

PR 3256705: ESXi hosts might fail to boot due to a rare deadlock between open and rename operations on objects with parent-child relationship

A rename operation locks the source and destination directories based on the object pointer address with the assumption that the parent object has the lower address. In rare cases, the assumption might not be correct and the rename lock can happen for the child first and then the parent, while the open operation locks the parent first and then the child. As a result, a deadlock between the open and rename operations occurs and ESXi hosts fail to boot.

This issue is resolved in this release.

PR 3311989: Linux-based virtual machines randomly stop responding to input from the keyboard and the mouse

A race condition on Linux-based virtual machines between the interrupt requests from the keyboard and the mouse, such as moving the mouse and typing at the same time, might cause the VM to stop responding to input from the keyboard and the mouse.

This issue is resolved in this release.

PR 3307579: You see virtual machines in environments with Enhanced Data Path enabled randomly losing connectivity

A rare memory leak might cause VMkernel NICs connected to ESXi host switches, working on the networking stack mode known as Enhanced Data Path or Enhanced Networking Stack, randomly lose connectivity.

This issue is resolved in this release.

PR 3311123: ESXi hosts randomly fail with a purple diagnostic screen and errors such as Spinlock spinout NMI or Panic requested by another PCPU

A rare race condition between an I/O thread and a Storage I/O Control thread on an ESXi host, waiting for each other to release read/write locks, might cause the host to fail with a purple diagnostic screen. In the panic backtrace you see a call to either `PsaStorSchedQUpdateMaxQDepth` or `PsaStorDirectSubmitOrQueueCommand` and the error on the screen is either `Spin count exceeded` or `Panic requested by another PCPU` respectively.

This issue is resolved in this release.

PR 3312713: An ESXi host might fail with purple diagnostic screen due to rare issue with the offline consolidation of SEsparse delta disks

Due to slow storage, the offline consolidation of a large SEsparse delta disk, such as 3TB or more, might take longer than the default 10-hour timeout. In such cases, a clearing process stops the consolidation task before it completes. As a result, threads trying to access the consolidation task cause a PCPU lock up and lead to a failure of the ESXi host with a purple diagnostic screen with an error such as `Spin count exceeded - possible deadlock`.

This issue is resolved in this release.

PR 3307435: The hostd service might fail after an update in environments with Virtual Shared Graphics Acceleration (vSGA) and Enhanced vMotion Compatibility (EVC)

In EVC-enabled clusters with active vSGA to share GPU resources across multiple virtual machines, an update operation, such as VIB installation on an ESXi host, might cause the hostd service to fail and start generating frequent coredumps. The issue occurs while adding the GPU-related features to the ESXi host after exiting maintenance mode.

This issue is resolved in this release.

PR 3303651: In a High Performance Plug-in (HPP) environment, a network outage might trigger host failure

In HPP environments, which aim to improve the performance of NVMe devices on ESXi hosts, in case of a network outage, you might see two issues:

1. In case of multi-pathed local SCSI devices in active-active mode claimed by HPP, VM I/Os might fail prematurely instead of being retried on the other path.
2. In case of local SCSI devices claimed by HPP, in an all paths down (APD) scenario, if a complete storage rescan is triggered from vCenter, the ESXi host might fail with a purple diagnostic screen.

Both issues occur only with SCSI local devices and when I/Os fail due to lack of network or APD.

This issue is resolved in this release. The fix makes sure that I/Os fail only in case of `CHECK CONDITION` error indicating a change in a SCSI local device.

esx-update_8.0.2-0.30.23305546

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esx-update_8.0.2-0.30.23305546 - VMware_bootbank_loadesx_8.0.2-0.30.23305546
PRs Fixed	3317203, 3313226, 3309606, 3261182, 3286057
CVE numbers	N/A

Updates the `esx-update` and `loadesx` VIBs to resolve the following issue:

PR 3309606: Using the `poll()` system call on character devices in some drivers might lead to an ESXi host failure with a purple diagnostic screen

Depending on the implementation, for example in the Integrated Lights Out (ILO) driver of HPE, using the `poll()` system call on character devices might lead to an ESXi host failure with a purple diagnostic screen. The issue is due to an internal memory leak in some cases where the driver returns an error to the `poll()` system call.

This issue is resolved in this release.

esxio-update_8.0.2-0.30.23305546

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esxio-update_8.0.2-0.30.23305546 - VMware_bootbank_loadesxio_8.0.2-0.30.23305546
PRs Fixed	N/A
CVE numbers	N/A

Updates the `esxio-update` and `loadesxio` VIBs.

VMware-HBR-UW_8.0.2-0.30.23305546

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	No
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs Included	VMware_bootbank_vmware-hbrsrv_8.0.2-0.30.23305546
PRs Fixed	3312725, 3312593, 3311987, 3311956, 3311691, 3295096, 3295060, 3306346, 3306344, 3306345
CVE numbers	N/A

Updates the hbrsrv VIB.

ESXi_8.0.2-0.25.23305545

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_vsan_8.0.2-0.25.23305545 - VMware_bootbank_esxio-combiner_8.0.2-0.25.23305545 - VMware_bootbank_infravisor_8.0.2-0.25.23305545 - VMware_bootbank_clusterstore_8.0.2-0.25.23305545 - VMware_bootbank_native-misc-drivers_8.0.2-0.25.23305545 - VMW_bootbank_pensandoatlas_1.46.0.E.28.1.314-2vmw.802.0.0.22939414 - VMware_bootbank_crx_8.0.2-0.25.23305545 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.25.23305545 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.25.23305545 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.25.23305545 - VMware_bootbank_esx-xserver_8.0.2-0.25.23305545 - VMware_bootbank_bmcad-esxio_8.0.2-0.25.23305545 - VMware_bootbank_trx_8.0.2-0.25.23305545 - VMware_bootbank_esxio_8.0.2-0.25.23305545 - VMware_bootbank_esxio-base_8.0.2-0.25.23305545 - VMware_bootbank_gc-esxio_8.0.2-0.25.23305545 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.25.23305545 - VMware_bootbank_bmcad_8.0.2-0.25.23305545 - VMware_bootbank_gc_8.0.2-0.25.23305545 - VMware_bootbank_vsanhealth_8.0.2-0.25.23305545 - VMware_bootbank_vdfs_8.0.2-0.25.23305545 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.25.23305545 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.25.23305545 - VMware_bootbank_vds-vsip_8.0.2-0.25.23305545 - VMware_bootbank_esx-base_8.0.2-0.25.23305545
PRs Fixed	3333007, 3328623, 3317046, 3315810, 3309040, 3341832, 3309039, 3293189, 3309041, 3291298, 3304476, 3260854, 3291282, 3291430, 3302555, 3301475, 3291280, 3291297, 3291275, 3291295
CVE numbers	N/A

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

Updates the `bootbank_vsan`, `bootbank_esxio-combiner`, `bootbank_infravisor`, `bootbank_clusterstore`, `bootbank_native-misc-drivers`, `VMW_bootbank_pensandoatlas`, `bootbank_crx`, `bootbank_esxio-combiner-esxio`, `bootbank_esxio-dvfilter-generic-fastpath`, `bootbank_esx-xserver`, `bootbank_bmc-al-esxio`, `bootbank_trx`, `bootbank_esxio`, `bootbank_esxio-base`, `bootbank_gc-esxio`, `bootbank_esx-dvfilter-generic-fastpath`, `bootbank_bmc-al`, `bootbank_gc`, `bootbank_vsanhealth`, `bootbank_vdfs`, `bootbank_native-misc-drivers-esxio`, `bootbank_drivervm-gpu-base`, `bootbank_vds-vsip`, and `bootbank_esx-base` VIBs to resolve the following issues:

ESXi 8.0 Update 2b provides the following security updates:

- OpenSSL is updated to version 3.0.11.
- The Python library is updated to version 3.8.18.
- The Envoy proxy is updated to version 1.25.11.
- The cURL library is updated to version 8.3.0.
- The c-ares C library is updated to version 1.19.1.
- The Go-lang library is updated to version 1.19.11.
- The pyOpenSSL library is updated to version 2.0.
- The Libxslt library is updated to version 1.1.28.
- The libarchive library is updated to version 3.5.1.

ESXi 8.0 Update 2b includes the following Intel microcode:

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Broadwell H	0x40671 (06/47/1)	0x22	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	0x4003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	0x5003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	0x7002703	3/21/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	0xe000014	9/18/2021	Intel Xeon D-1600 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Skylake H/S	0x506e3 (06/5e/3)	0x36	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	0x38	12/2/2021	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	0xd0003b9	9/1/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	0x1000268	9/8/2023	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	0xb4	9/7/2023	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	0x34	9/7/2023	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	0x4e	9/7/2023	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	0x2c000321	8/21/2023	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	0x2b000541	8/21/2023	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	0xf4	2/23/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	0xfa	2/27/2023	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	0x5d	9/3/2023	Intel Xeon E-2300 Series

Code Name	FMS	PIt ID	MCU Rev	MCU Date	Brand Names
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	0x11e	8/31/2023	Intel Xeon E-2400 Series

esx-update_8.0.2-0.25.23305545

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_loadesx_8.0.2-0.25.23305545 - VMware_bootbank_esx-update_8.0.2-0.25.23305545
PRs Fixed	N/A
CVE numbers	N/A

Updates the `esx-update` and `loadesx` VIBs.

esxio-update_8.0.2-0.25.23305545

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	- VMware_bootbank_esxio-update_8.0.2-0.25.23305545 - VMware_bootbank_loadesxio_8.0.2-0.25.23305545
PRs Fixed	N/A
CVE numbers	N/A

Updates the `esxio-update` and `loadesxio` VIBs.

VMware-VM-Tools_12.3.5.22544099-23305545

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	No
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs Included	VMware_locker_tools-light_12.3.5.22544099-23305545
PRs Fixed	3304809

CVE numbers	N/A
-------------	-----

Updates the `locker_tools-light` VIB to resolve the following issue:

VMware Tools Bundling Changes in ESXi 8.0 Update 2b

- The following VMware Tools ISO images are bundled with ESXi 8.0 Update 2b:
 - **windows.iso**: VMware Tools 12.3.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.
- The following VMware Tools ISO images are available for download:
- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
 - VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
 - **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
 - **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.3.5 Release Notes](#)
- [Earlier versions of VMware Tools](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Updating VMware Tools](#)

ESXi-8.0U2b-23305546-standard

Profile Name	ESXi-8.0U2b-23305546-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	February 29, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	• VMware_bootbank_vds-vsip_8.0.2-0.30.23305546 • VMware_bootbank_crx_8.0.2-0.30.23305546 • VMware_bootbank_esxio-combiner-esxio_8.0.2-0.30.23305546 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.30.23305546 • VMware_bootbank_vsanhealth_8.0.2-0.30.23305546 • VMware_bootbank_gc-esxio_8.0.2-0.30.23305546 • VMware_bootbank_clusterstore_8.0.2-0.30.23305546 • VMware_bootbank_trx_8.0.2-0.30.23305546 • VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.30.23305546 • VMware_bootbank_cpu-microcode_8.0.2-0.30.23305546 • VMware_bootbank_infravisor_8.0.2-0.30.23305546 • VMware_bootbank_esx-base_8.0.2-0.30.23305546 • VMware_bootbank_vsan_8.0.2-0.30.23305546

	- VMware_bootbank_bmcad_8.0.2-0.30.23305546 - VMware_bootbank_esxio_8.0.2-0.30.23305546 - VMW_bootbank_pensandoatlas_1.46.0.E.28.1.314-2 - VMware_bootbank_gc_8.0.2-0.30.23305546 - VMware_bootbank_vdfs_8.0.2-0.30.23305546 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.30.23305546 - VMware_bootbank_esxio-combiner_8.0.2-0.30.23305546 - VMware_bootbank_native-misc-drivers_8.0.2-0.30.23305546 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.30.23305546 - VMware_bootbank_esx-xserver_8.0.2-0.30.23305546 - VMware_bootbank_bmcad-esxio_8.0.2-0.30.23305546 - VMware_bootbank_esxio-base_8.0.2-0.30.23305546 - VMware_bootbank_esx-update_8.0.2-0.30.23305546 - VMware_bootbank_loadesx_8.0.2-0.30.23305546 - VMware_bootbank_esxio-update_8.0.2-0.30.23305546 - VMware_bootbank_loadesxio_8.0.2-0.30.23305546 - VMware_bootbank_vmware-hbrsrv_8.0.2-0.30.23305546 - VMware_locker_tools-light_12.3.5.22544099-23305545
PRs Fixed	3326852, 3326624, 3319724, 3316967, 3307362, 3318781, 3331997, 3317441, 3317938, 3336882, 3317707, 3316767, 3306552, 3210610, 3262965, 3312008, 3304392, 3308819, 3313089, 3313912, 3313572, 3308844, 3298815, 3312006, 3311820, 3308054, 3308132, 3302896, 3308133, 3312713, 3301092, 3296088, 3311524, 3311123, 3307579, 3308347, 3309887, 3308458, 3309849, 3303232, 3311989, 3308023, 3308025, 3301665, 3302759, 3305165, 3304472, 3303651, 3293470, 3303807, 3307435, 3282049, 3293466, 3293464, 3296047, 3300173, 3294828, 3300031, 3295243, 3268113, 3284909, 3256705, 3292627, 3287767, 3289239, 3297899, 3295250, 3287747, 3222955, 3300326, 3292550, 3300705, 3287772, 3294624, 3291105, 3296694, 3291219, 3280469, 3288307, 3287769, 3287762, 3289397, 3293199, 3285514, 3275379, 3288872, 3317203, 3313226, 3309606, 3261182, 3286057, 3312725, 3312593, 3311987, 3311956, 3311691, 3295096, 3295060, 3306346, 3306344, 3306345
Related CVE numbers	N/A

This patch updates the following issues:

PR 3316767: VM I/Os might prematurely fail or ESXi hosts crash while HPP claims a local multipath SCSI device

Due to a rare issue with the failover handling in the High Performance Plug-in (HPP) module, if HPP claims a local multipath SCSI device in active-active mode, you might see the following issues:

- in case of connection loss or a NO CONNECT error, VM I/Os can fail prematurely instead of retrying on another path.
- in case of all paths down (APD) state, if a complete storage rescan is triggered from vCenter or the ESXi host, the host might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3302759: Pending I/Os might cause the ESXi host running a Primary VM to fail with a purple diagnostic screen during failover

During a failover of a Primary VM protected by vSphere Fault Tolerance, if some I/Os between the VM and the ESXi host are still unresolved, the host might intermittently fail with a purple diagnostic screen.

This issue is resolved in this release. The fix makes sure that no pending I/Os exist at the time an FT session info cleanup runs.

PR 3309606: Using the poll() system call on character devices in some drivers might lead to an ESXi host failure with a purple diagnostic screen

Depending on the implementation, for example in the Integrated Lights Out (ILO) driver of HPE, using the `poll()` system call on character devices might lead to an ESXi host failure with a purple diagnostic screen. The issue is due to an internal memory leak in some cases where the driver returns an error to the `poll()` system call.

This issue is resolved in this release.

PR 3291105: While detaching an Enhanced Networking Stack (ENS) port from an ESXi host, the host might fail with a purple diagnostic screen

During deactivation of ENS in scenarios such as an upgrade or a VNIC configuration change, while detaching an ENS port from an ESXi host with active ENS, the port might be removed from the port list and not be available for further reference, which might cause an invalid memory access error. As a result, the ESXi host fails with a purple diagnostic screen.

This issue is resolved in this release.

PR 3297899: vSAN file server continues restarting

In case all AD servers are inaccessible, vSAN Skyline Health might show a file server that is always in restarting state, with repeated failovers from one host to another. If this occurs, all file shares on the file server, including NFS shares with `securityType=AUTH_SYS`, become inaccessible.

This issue is resolved in this release.

PR 3288307: Cannot upgrade vSAN File Service or enter maintenance mode

In rare cases, if the `hostd` service reboots, upgrade of the vSAN File Service or entering maintenance mode might be blocked.

This issue is resolved in this release.

PR 3294624: Skyline Health fails to load for file shares in an inextensible state

When a file share is in an inextensible state, the current storage policy cannot be fulfilled. The file share might also have other issues, such as exceeding the soft quota with its used capacity. Skyline Health fails to load for objects in this state.

This issue is resolved in this release.

PR 3293199: The file `/var/run/log/clusterAgent.stderr` on an ESXi host becomes extremely large, filling up the OSDATA partition

In case of certain rare error conditions, a component within the `clusterAgent` ESXi service might start periodically writing to `/var/run/log/clusterAgent.stderr` instead of the regular `syslog` destination on an ESXi host. Since this is not expected behavior, the `stderr` file is not rotated or otherwise monitored. If the error conditions persist over time, for example several months, the file can expand to several GB in size, filling up the OSDATA partition. When disk space is exhausted, ESXi services and VMs might lock up.

This issue is resolved in this release. The fix makes sure that if your environment had the issue, after upgrading to ESXi 8.0 Update 2b or later, the oversized log file is compressed. The file is not automatically deleted, but you can safely manually remove the compressed file.

PR 3319724: vSAN ESA VMs with large-sized VMDKs become inaccessible

In vSAN Express Storage Architecture (ESA) environments of version 8.0 and later, VMDK objects with a large amount of used capacity might overload the object scrubber, causing such objects to become inaccessible.

This issue is resolved in this release.

PR 3294828: The `hostd` service might intermittently fail due to an invalid sensor data record

In rare cases, the Baseboard Management Controller (BMC) might provide an invalid sensor data record to the `hostd` service and cause it to fail. As a result, ESXi hosts become unresponsive or disconnect from vCenter.

This issue is resolved in this release. The fix implements checks to validate the sensor data record before passing it to the `hostd` service.

PR 3287769: Unexpected failover of vSAN File Service containers

The EMM monitor can trigger unexpectedly and block the File Service EPiC heartbeat update to the arbitrator host. This problem can cause an unexpected container failover, which interrupts the I/O of SMB file shares.

This issue is resolved in this release.

PR 3293464: The VMX service might fail during migration of virtual machines with vSphere Virtual Volumes if the volume closure on the source takes long

In some scenarios, such as virtual machines with vSphere Virtual Volumes and Changed Block Tracking (CBT) or Content Based Read Cache (CBRC) enabled, during storage migration, flushing pending I/Os on volume closure at the source might take long, like 10 sec. If clearing the pending I/Os is not complete by the time the VMX service tries to reach the volumes on the destinations host, the service fails.

This issue is resolved in this release. The fix makes sure that all vSphere Virtual Volumes and disks are closed on the source host before the destination host tries to open them.

PR 3308819: You cannot add allowed IP addresses for an ESXi host

With ESXi 8.0 Update 2, some ESXi firewall rulesets, such as `dhcp`, are system-owned by default and prevent manual adding of allowed IP addresses to avoid possible break of service. With ESXi 8.0 Update 2b, you can manually add allowed IP addresses to all rulesets, except for `nfsClient`, `nfs41Client`, `trusted-infrastructure-kmxd`, `trusted-infrastructure-kmxa`, and `vsanEncryption`.

This issue is resolved in this release.

PR 3308461: When MAC learning is not active, newly added ports on a vSphere Distributed Switch (VDS) might go into blocked state

If MAC learning on a VDS is not active, when from the guest OS you change the MAC address of a virtual machine to a new proxy switch port, the newly added port might be in a blocked state due to a L2 violation error, even when the **Mac address changes option** is set to **Accept**. The issue occurs only when MAC learning on a VDS is not active.

This issue is resolved in this release.

PR 3332617: vGPU VMs with NVIDIA RTX 5000 Ada or NVIDIA RTX 6000 Ada graphics card fail to start after update to ESXi 8.0 Update 2

After an update to ESXi 8.0 Update 2, booting of a vGPU VM with NVIDIA RTX 5000 Ada or NVIDIA RTX 6000 Ada graphics card fails with an error such as `vmiop-display unable to reserve vgpu` in the logs. In the vSphere Client, you see an error such as `Module DevicePowerOnEarly power on failed..`

This issue is resolved in this release.

PR 3326624: An ESXi host might fail during upgrade due to an unmap I/O request larger than 2040 MB on a vSAN-backed object

During an upgrade to ESXi 8.0 Update 2, in rare cases, a guest virtual machine might issue an unmap I/O request larger than 2040 MB on a vSAN-backed object and the request might split in two parts. As a result, the ESXi host receives duplicate metadata entries and might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3304472: vSphere vMotion operations might fail due to a rare issue with no namespaces present in the namespacesmgr.db file

When a virtual machine is configured with namespaces and for some reason the namespaces table in the `namespacemgr.db` file is empty in the ESXi host database, migration of such a VM by using vSphere vMotion might fail. In the `vmware.log` file, you see errors such as:

vmx - NamespaceMgrCheckpoint: No valid queue found while restoring the namespace events. The migrate type is 1.

This issue is resolved in this release.

PR 3293470: In the vSphere Client, you see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore

In the vSphere Client, when you navigate to **Monitor >Advanced >Datastore** to see the advanced performance metrics of a virtual machine, you might see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore. The issue occurs due to incorrect calculation of the counters for the read/write IO latency as a cumulative average.

This issue is resolved in this release.

PR 3300031: The Network Time Protocol (NTP) peer and loop statistics are not automatically captured in the vm-support bundle

Prior to vSphere 8.0 Update 2b, unless you manually add the NTP peer and loop statistics to the vm-support bundle, NTP logs are not available for troubleshooting.

This issue is resolved in this release. The fix adds support for automatically capturing NTP peer and loop statistics from the `/var/log/ntp/` directory in the vm-support bundle.

PR 3293466: vSphere vMotion operations between ESXi 7.x hosts fail with the error "Destination failed to preopen disks"

If on an ESXi 7.x host you use an older Virtual Volumes storage provider, also called a VASA provider, it might not allow binding multiple vSphere Virtual Volumes in batches to optimize the performance of migration and power on. As a result, in the vSphere Client you might see the error `Destination failed to preopen disks` when running vSphere vMotion operations between 7.x hosts.

This issue is resolved in this release.

PR 3316967: Changed Block Tracking (CBT) might not work as expected on a hot extended virtual disk

In vSphere 8.0 Update 2, to optimize the open and close process of virtual disks during hot extension, the disk remains open during hot extend operations. Due to this change, incremental backup of virtual disks with CBT enabled might be incomplete, because the CBT in-memory bitmap does not resize, and CBT cannot record the changes to the extended disk block. As a result, when you try to restore a VM from an incremental backup of virtual disks with CBT, the VM might fail to start.

This issue is resolved in this release.

PR 3317441: Enhanced Network Datapath (EDP) might reserve too much memory on a large vSphere system and cause resource issues

EDP reservation is linear to the total memory by reserving up to 1.6% of the total. For a VM with high memory reservation in a large vSphere system, such a memory buffer might be too big. For example, if the system has a 100TB of memory, the EDP reservation might be 1.6TB. As a result, you might see resource issues such as not being able to hot-add VNICS to VMs with high memory reservation.

This issue is resolved in this release. The fix sets the EDP reservation limit at 1% of the total memory and puts a 5GB cap on a single memory reservation. You can adjust the memory reservation limit by using the VMkernel boot option `ensMbufPoolMaxMinMB` or by using the following command: `esxcli system settings kernel set -s ensMbufPoolMaxMBPerGB -v <value-in-MB>`. In any case, you must reboot the ESXi host to enable the change.

PR 3275379: VBS-enabled Windows VMs might intermittently fail with a blue diagnostic screen on ESXi hosts running on AMD processors.

Windows virtual machines with virtualization-based security (VBS) enabled might intermittently fail with a blue diagnostic screen on ESXi hosts running on AMD processors. The BSOD has the following signature:

SECURE_KERNEL_ERROR (18b)The secure kernel has encountered a fatal error.

Arguments:

Arg1: 000000000000018c

Arg2: 000000000000100b

Arg3: 0000000000000000

Arg4: 000000000000409b

This issue is resolved in this release.

PR 3311820: Creation of VMFS volumes on local software emulated 4Kn devices fails

In some environments, after an update to ESXi 8.0 Update 2, you might not be able to create VMFS volumes on local software emulated 4Kn devices. The issue occurs when a SCSI driver fails to discover and register a 4Kn NVMe device as a 4Kn SWE device. In the vSphere On-disk Metadata Analyzer (VOMA), you see an error such as `ON-DISK ERROR: Invalid disk block size 512, should be 4096`.

This issue is resolved in this release.

PR 3296047: The vvoid service fails with Unable to Allocate Memory error

If a vSphere API for Storage Awareness (VASA) provider is down for a long time, frequent reconnection attempts from the vvoid service might exhaust its memory, because of repetitively creating and destroying session objects. The issue can also occur when an event such as a change in a VASA provider triggers an update to the vvoid service. The vvoid service restarts within seconds, but if the restart interrupts a running operation of the VASA provider, the operation might fail and affect related virtual machine operations as well. In the backtrace, you see an error such as `Unable To Allocate Memory`.

This issue is resolved in this release.

PR 3309887: You cannot use host profiles to configure a path selection policy for storage disks

In some cases, when you try to change or configure the path selection policy of storage disks by using a host profile, the operation fails. In the `syslog.log` file, you see lines such as

Errors: Unable to set SATP configuration: Invalid options provided for SATP VMW_SATP_ALUA

This issue is resolved in this release.

PR 3308844: The hostd service might become unresponsive after a restart due to a rare issue with Inter Process Communication (IPC)

If an IPC process has incorrect payload in the IPC packet, for example: an incorrect version in the header, a deadlock in the iSCSI daemon might cause the hostd service to become unresponsive after a hostd restart. As a result, you cannot perform any operations on the affected ESXi host and VMs by using the vSphere Client.

This issue is resolved in this release.

PR 3256705: ESXi hosts might fail to boot due to a rare deadlock between open and rename operations on objects with parent-child relationship

A rename operation locks the source and destination directories based on the object pointer address with the assumption that the parent object has the lower address. In rare cases, the assumption might not be correct and the rename lock can happen for the child first and then the parent, while the open operation locks the parent first and then the child. As a result, a deadlock between the open and rename operations occurs and ESXi hosts fail to boot.

This issue is resolved in this release.

PR 3311989: Linux-based virtual machines randomly stop responding to input from the keyboard and the mouse

A race condition on Linux-based virtual machines between the interrupt requests from the keyboard and the mouse, such as moving the mouse and typing at the same time, might cause the VM to stop responding to input from the keyboard and the mouse.

This issue is resolved in this release.

PR 3307579: You see virtual machines in environments with Enhanced Data Path enabled randomly losing connectivity

A rare memory leak might cause VMkernel NICs connected to ESXi host switches, working on the networking stack mode known as Enhanced Data Path or Enhanced Networking Stack, randomly lose connectivity.

This issue is resolved in this release.

PR 3311123: ESXi hosts randomly fail with a purple diagnostic screen and errors such as Spinlock spinout NMI or Panic requested by another PCPU

A rare race condition between an I/O thread and a Storage I/O Control thread on an ESXi host, waiting for each other to release read/write locks, might cause the host to fail with a purple diagnostic screen. In the panic backtrace you see a call to either `PsaStorSchedQUpdateMaxQDepth` or `PsaStorDirectSubmitOrQueueCommand` and the error on the screen is either `Spin count exceeded` or `Panic requested by another PCPU` respectively.

This issue is resolved in this release.

PR 3312713: An ESXi host might fail with purple diagnostic screen due to rare issue with the offline consolidation of SEsparse delta disks

Due to slow storage, the offline consolidation of a large SEsparse delta disk, such as 3TB or more, might take longer than the default 10-hour timeout. In such cases, a clearing process stops the consolidation task before it completes. As a result, threads trying to access the consolidation task cause a PCPU lock up and lead to a failure of the ESXi host with a purple diagnostic screen with an error such as `Spin count exceeded - possible deadlock`.

This issue is resolved in this release.

PR 3307435: The hostd service might fail after an update in environments with Virtual Shared Graphics Acceleration (vSGA) and Enhanced vMotion Compatibility (EVC)

In EVC-enabled clusters with active vSGA to share GPU resources across multiple virtual machines, an update operation, such as VIB installation on an ESXi host, might cause the hostd service to fail and start generating frequent coredumps. The issue occurs while adding the GPU-related features to the ESXi host after exiting maintenance mode.

This issue is resolved in this release.

PR 3303651: In a High Performance Plug-in (HPP) environment, a network outage might trigger host failure

In HPP environments, which aim to improve the performance of NVMe devices on ESXi hosts, in case of a network outage, you might see two issues:

1. In case of multi-pathed local SCSI devices in active-active mode claimed by HPP, VM I/Os might fail prematurely instead of being retried on the other path.
2. In case of local SCSI devices claimed by HPP, in an all paths down (APD) scenario, if a complete storage rescan is triggered from vCenter, the ESXi host might fail with a purple diagnostic screen.

Both issues occur only with SCSI local devices and when I/Os fail due to lack of network or APD.

This issue is resolved in this release. The fix makes sure that I/Os fail only in case of `CHECK CONDITION` error indicating a change in a SCSI local device.

ESXi-8.0U2b-23305546-no-tools

Profile Name	ESXi-8.0U2b-23305546-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	February 29, 2024

Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_vds-vsip_8.0.2-0.30.23305546 - VMware_bootbank_crx_8.0.2-0.30.23305546 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.30.23305546 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.30.23305546 - VMware_bootbank_vsanhealth_8.0.2-0.30.23305546 - VMware_bootbank_gc-esxio_8.0.2-0.30.23305546 - VMware_bootbank_clusterstore_8.0.2-0.30.23305546 - VMware_bootbank_trx_8.0.2-0.30.23305546 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.30.23305546 - VMware_bootbank_cpu-microcode_8.0.2-0.30.23305546 - VMware_bootbank_infravisor_8.0.2-0.30.23305546 - VMware_bootbank_esx-base_8.0.2-0.30.23305546 - VMware_bootbank_vsan_8.0.2-0.30.23305546 - VMware_bootbank_bmcad_8.0.2-0.30.23305546 - VMware_bootbank_esxio_8.0.2-0.30.23305546 - VMware_bootbank_pensandoatlas_1.46.0.E.28.1.314-2 - VMware_bootbank_gc_8.0.2-0.30.23305546 - VMware_bootbank_vdfs_8.0.2-0.30.23305546 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.30.23305546 - VMware_bootbank_esxio-combiner_8.0.2-0.30.23305546 - VMware_bootbank_native-misc-drivers_8.0.2-0.30.23305546 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.30.23305546 - VMware_bootbank_esx-xserver_8.0.2-0.30.23305546 - VMware_bootbank_bmcad-esxio_8.0.2-0.30.23305546 - VMware_bootbank_esxio-base_8.0.2-0.30.23305546 - VMware_bootbank_esx-update_8.0.2-0.30.23305546 - VMware_bootbank_loadesx_8.0.2-0.30.23305546 - VMware_bootbank_esxio-update_8.0.2-0.30.23305546 - VMware_bootbank_loadesxio_8.0.2-0.30.23305546 - VMware_bootbank_vmware-hbrsrv_8.0.2-0.30.23305546
PRs Fixed	3326852, 3326624, 3319724, 3316967, 3307362, 3318781, 3331997, 3317441, 3317938, 3336882, 3317707, 3316767, 3306552, 3210610, 3262965, 3312008, 3304392, 3308819, 3313089, 3313912, 3313572, 3308844, 3298815, 3312006, 3311820, 3308054, 3308132, 3302896, 3308133, 3312713, 3301092, 3296088, 3311524, 3311123, 3307579, 3308347, 3309887, 3308458, 3309849, 3303232, 3311989, 3308023, 3308025, 3301665, 3302759, 3305165, 3304472, 3303651, 3293470, 3303807, 3307435, 3282049, 3293466, 3293464, 3296047, 3300173, 3294828, 3300031, 3295243, 3268113, 3284909, 3256705, 3292627, 3287767, 3289239, 3297899, 3295250, 3287747, 3222955, 3300326, 3292550, 3300705, 3287772, 3294624, 3291105, 3296694, 3291219, 3280469, 3288307, 3287769, 3287762, 3289397, 3293199, 3285514, 3275379, 3288872, 3317203, 3313226, 3309606, 3261182, 3286057, 3312725, 3312593, 3311987, 3311956, 3311691, 3295096, 3295060, 3306346, 3306344, 3306345
Related CVE numbers	N/A

This patch updates the following issues:

PR 3302759: Pending I/Os might cause the ESXi host running a Primary VM to fail with a purple diagnostic screen during failover

During a failover of a Primary VM protected by vSphere Fault Tolerance, if some I/Os between the VM and the ESXi host are still unresolved, the host might intermittently fail with a purple diagnostic screen.

This issue is resolved in this release. The fix makes sure that no pending I/Os exist at the time an FT session info cleanup runs.

PR 3309606: Using the poll() system call on character devices in some drivers might lead to an ESXi host failure with a purple diagnostic screen

Depending on the implementation, for example in the Integrated Lights Out (ILO) driver of HPE, using the `poll()` system call on character devices might lead to an ESXi host failure with a purple diagnostic screen. The issue is due to an internal memory leak in some cases where the driver returns an error to the `poll()` system call.

This issue is resolved in this release.

PR 3291105: While detaching an Enhanced Networking Stack (ENS) port from an ESXi host, the host might fail with a purple diagnostic screen

During deactivation of ENS in scenarios such as an upgrade or a VNIC configuration change, while detaching an ENS port from an ESXi host with active ENS, the port might be removed from the port list and not be available for further reference, which might cause an invalid memory access error. As a result, the ESXi host fails with a purple diagnostic screen.

This issue is resolved in this release.

PR 3297899: vSAN file server continues restarting

In case all AD servers are inaccessible, vSAN Skyline Health might show a file server that is always in restarting state, with repeated failovers from one host to another. If this occurs, all file shares on the file server, including NFS shares with `securityType=AUTH_SYS`, become inaccessible.

This issue is resolved in this release.

PR 3288307: Cannot upgrade vSAN File Service or enter maintenance mode

In rare cases, if the `hostd` service reboots, upgrade of the vSAN File Service or entering maintenance mode might be blocked.

This issue is resolved in this release.

PR 3294624: Skyline Health fails to load for file shares in an inextensible state

When a file share is in an inextensible state, the current storage policy cannot be fulfilled. The file share might also have other issues, such as exceeding the soft quota with its used capacity. Skyline Health fails to load for objects in this state.

This issue is resolved in this release.

PR 3293199: The file `/var/run/log/clusterAgent.stderr` on an ESXi host becomes extremely large, filling up the OSDATA partition

In case of certain rare error conditions, a component within the `clusterAgent` ESXi service might start periodically writing to `/var/run/log/clusterAgent.stderr` instead of the regular syslog destination on an ESXi host. Since this is not expected behavior, the `stderr` file is not rotated or otherwise monitored. If the error conditions persist over time, for example several months, the file can expand to several GB in size, filling up the OSDATA partition. When disk space is exhausted, ESXi services and VMs might lock up.

This issue is resolved in this release. The fix makes sure that if your environment had the issue, after upgrading to ESXi 8.0 Update 2b or later, the oversized log file is compressed. The file is not automatically deleted, but you can safely manually remove the compressed file.

PR 3319724: vSAN ESA VMs with large-sized VMDKs become inaccessible

In vSAN Express Storage Architecture (ESA) environments of version 8.0 and later, VMDK objects with a large amount of used capacity might overload the object scrubber, causing such objects to become inaccessible.

This issue is resolved in this release.

PR 3294828: The hostd service might intermittently fail due to an invalid sensor data record

In rare cases, the Baseboard Management Controller (BMC) might provide an invalid sensor data record to the hostd service and cause it to fail. As a result, ESXi hosts become unresponsive or disconnect from vCenter.

This issue is resolved in this release. The fix implements checks to validate the sensor data record before passing it to the hostd service.

PR 3287769: Unexpected failover of vSAN File Service containers

The EMM monitor can trigger unexpectedly and block the File Service EPiC heartbeat update to the arbitrator host. This problem can cause an unexpected container failover, which interrupts the I/O of SMB file shares.

This issue is resolved in this release.

PR 3293464: The VMX service might fail during migration of virtual machines with vSphere Virtual Volumes if the volume closure on the source takes long

In some scenarios, such as virtual machines with vSphere Virtual Volumes and Changed Block Tracking (CBT) or Content Based Read Cache (CBRC) enabled, during storage migration, flushing pending I/Os on volume closure at the source might take long, like 10 sec. If clearing the pending I/Os is not complete by the time the VMX service tries to reach the volumes on the destinations host, the service fails.

This issue is resolved in this release. The fix makes sure that all vSphere Virtual Volumes and disks are closed on the source host before the destination host tries to open them.

PR 3308819: You cannot add allowed IP addresses for an ESXi host

With ESXi 8.0 Update 2, some ESXi firewall rulesets, such as `dhcp`, are system-owned by default and prevent manual adding of allowed IP addresses to avoid possible break of service. With ESXi 8.0 Update 2b, you can manually add allowed IP addresses to all rulesets, except for `nfsClient`, `nfs41Client`, `trusted-infrastructure-kmxd`, `trusted-infrastructure-kmxa`, and `vsanEncryption`.

This issue is resolved in this release.

PR 3308461: When MAC learning is not active, newly added ports on a vSphere Distributed Switch (VDS) might go into blocked state

If MAC learning on a VDS is not active, when from the guest OS you change the MAC address of a virtual machine to a new proxy switch port, the newly added port might be in a blocked state due to a L2 violation error, even when the **Mac address changes option** is set to **Accept**. The issue occurs only when MAC learning on a VDS is not active.

This issue is resolved in this release.

PR 3332617: vGPU VMs with NVIDIA RTX 5000 Ada or NVIDIA RTX 6000 Ada graphics card fail to start after update to ESXi 8.0 Update 2

After an update to ESXi 8.0 Update 2, booting of a vGPU VM with NVIDIA RTX 5000 Ada or NVIDIA RTX 6000 Ada graphics card fails with an error such as `vmiop-display unable to reserve vgpu` in the logs. In the vSphere Client, you see an error such as `Module DevicePowerOnEarly power on failed..`

This issue is resolved in this release.

PR 3326624: An ESXi host might fail during upgrade due to an unmap I/O request larger than 2040 MB on a vSAN-backed object

During an upgrade to ESXi 8.0 Update 2, in rare cases, a guest virtual machine might issue an unmap I/O request larger than 2040 MB on a vSAN-backed object and the request might split in two parts. As a result, the ESXi host receives duplicate metadata entries and might fail with a purple diagnostic screen.

This issue is resolved in this release.

PR 3304472: vSphere vMotion operations might fail due to a rare issue with no namespaces present in the namespacemgr.db file

When a virtual machine is configured with namespaces and for some reason the namespaces table in the namespacemgr.db file is empty in the ESXi host database, migration of such a VM by using vSphere vMotion might fail. In the vmware.log file, you see errors such as:

```
vmx - NamespaceMgrCheckpoint: No valid queue found while restoring the namespace events.
The migrate type is 1.
```

This issue is resolved in this release.

PR 3293470: In the vSphere Client, you see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore

In the vSphere Client, when you navigate to **Monitor > Advanced > Datastore** to see the advanced performance metrics of a virtual machine, you might see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore. The issue occurs due to incorrect calculation of the counters for the read/write IO latency as a cumulative average.

This issue is resolved in this release.

PR 3300031: The Network Time Protocol (NTP) peer and loop statistics are not automatically captured in the vm-support bundle

Prior to vSphere 8.0 Update 2b, unless you manually add the NTP peer and loop statistics to the vm-support bundle, NTP logs are not available for troubleshooting.

This issue is resolved in this release. The fix adds support for automatically capturing NTP peer and loop statistics from the /var/log/ntp/ directory in the vm-support bundle.

PR 3293466: vSphere vMotion operations between ESXi 7.x hosts fail with the error "Destination failed to preopen disks"

If on an ESXi 7.x host you use an older Virtual Volumes storage provider, also called a VASA provider, it might not allow binding multiple vSphere Virtual Volumes in batches to optimize the performance of migration and power on. As a result, in the vSphere Client you might see the error `Destination failed to preopen disks` when running vSphere vMotion operations between 7.x hosts.

This issue is resolved in this release.

PR 3316967: Changed Block Tracking (CBT) might not work as expected on a hot extended virtual disk

In vSphere 8.0 Update 2, to optimize the open and close process of virtual disks during hot extension, the disk remains open during hot extend operations. Due to this change, incremental backup of virtual disks with CBT enabled might be incomplete, because the CBT in-memory bitmap does not resize, and CBT cannot record the changes to the extended disk block. As a result, when you try to restore a VM from an incremental backup of virtual disks with CBT, the VM might fail to start.

This issue is resolved in this release.

PR 3317441: Enhanced Network Datapath (EDP) might reserve too much memory on a large vSphere system and cause resource issues

EDP reservation is linear to the total memory by reserving up to 1.6% of the total. For a VM with high memory reservation in a large vSphere system, such a memory buffer might be too big. For example, if the system has a 100TB of memory,

the EDP reservation might be 1.6TB. As a result, you might see resource issues such as not being able to hot-add VNICs to VMs with high memory reservation.

This issue is resolved in this release. The fix sets the EDP reservation limit at 1% of the total memory and puts a 5GB cap on a single memory reservation. You can adjust the memory reservation limit by using the VMkernel boot option `ensMbufPoolMaxMinMB` or by using the following command: `esxcli system settings kernel set -s ensMbufPoolMaxMBPerGB -v <value-in-MB>`. In any case, you must reboot the ESXi host to enable the change.

PR 3275379: VBS-enabled Windows VMs might intermittently fail with a blue diagnostic screen on ESXi hosts running on AMD processors.

Windows virtual machines with virtualization-based security (VBS) enabled might intermittently fail with a blue diagnostic screen on ESXi hosts running on AMD processors. The BSOD has the following signature:

```
SECURE_KERNEL_ERROR (18b)The secure kernel has encountered a fatal error.
Arguments:
Arg1: 0000000000000018c
Arg2: 0000000000000100b
Arg3: 0000000000000000
Arg4: 000000000000409b
```

This issue is resolved in this release.

PR 3311820: Creation of VMFS volumes on local software emulated 4Kn devices fails

In some environments, after an update to ESXi 8.0 Update 2, you might not be able to create VMFS volumes on local software emulated 4Kn devices. The issue occurs when a SCSI driver fails to discover and register a 4Kn NVMe device as a 4Kn SWE device. In the vSphere On-disk Metadata Analyzer (VOMA), you see an error such as `ON-DISK ERROR: Invalid disk block size 512, should be 4096`.

This issue is resolved in this release.

PR 3296047: The vvolld service fails with Unable to Allocate Memory error

If a vSphere API for Storage Awareness (VASA) provider is down for a long time, frequent reconnection attempts from the vvolld service might exhaust its memory, because of repetitively creating and destroying session objects. The issue can also occur when an event such as a change in a VASA provider triggers an update to the vvolld service. The vvolld service restarts within seconds, but if the restart interrupts a running operation of the VASA provider, the operation might fail and affect related virtual machine operations as well. In the backtrace, you see an error such as `Unable To Allocate Memory`.

This issue is resolved in this release.

PR 3309887: You cannot use host profiles to configure a path selection policy for storage disks

In some cases, when you try to change or configure the path selection policy of storage disks by using a host profile, the operation fails. In the `syslog.log` file, you see lines such as

```
Errors: Unable to set SATP configuration: Invalid options provided for SATP VMW_SATP_ALUA
```

This issue is resolved in this release.

PR 3308844: The hostd service might become unresponsive after a restart due to a rare issue with Inter Process Communication (IPC)

If an IPC process has incorrect payload in the IPC packet, for example: an incorrect version in the header, a deadlock in the iSCSI daemon might cause the hostd service to become unresponsive after a hostd restart. As a result, you cannot perform any operations on the affected ESXi host and VMs by using the vSphere Client.

This issue is resolved in this release.

PR 3256705: ESXi hosts might fail to boot due to a rare deadlock between open and rename operations on objects with parent-child relationship

A rename operation locks the source and destination directories based on the object pointer address with the assumption that the parent object has the lower address. In rare cases, the assumption might not be correct and the rename lock can happen for the child first and then the parent, while the open operation locks the parent first and then the child. As a result, a deadlock between the open and rename operations occurs and ESXi hosts fail to boot.

This issue is resolved in this release.

PR 3311989: Linux-based virtual machines randomly stop responding to input from the keyboard and the mouse

A race condition on Linux-based virtual machines between the interrupt requests from the keyboard and the mouse, such as moving the mouse and typing at the same time, might cause the VM to stop responding to input from the keyboard and the mouse.

This issue is resolved in this release.

PR 3307579: You see virtual machines in environments with Enhanced Data Path enabled randomly losing connectivity

A rare memory leak might cause VMkernel NICs connected to ESXi host switches, working on the networking stack mode known as Enhanced Data Path or Enhanced Networking Stack, randomly lose connectivity.

This issue is resolved in this release.

PR 3311123: ESXi hosts randomly fail with a purple diagnostic screen and errors such as Spinlock spinout NMI or Panic requested by another PCPU

A rare race condition between an I/O thread and a Storage I/O Control thread on an ESXi host, waiting for each other to release read/write locks, might cause the host to fail with a purple diagnostic screen. In the panic backtrace you see a call to either `PsaStorSchedQUpdateMaxQDepth` or `PsaStorDirectSubmitOrQueueCommand` and the error on the screen is either `Spin count exceeded` or `Panic requested by another PCPU` respectively.

This issue is resolved in this release.

PR 3312713: An ESXi host might fail with purple diagnostic screen due to rare issue with the offline consolidation of SEsparse delta disks

Due to slow storage, the offline consolidation of a large SEsparse delta disk, such as 3TB or more, might take longer than the default 10-hour timeout. In such cases, a clearing process stops the consolidation task before it completes. As a result, threads trying to access the consolidation task cause a PCPU lock up and lead to a failure of the ESXi host with a purple diagnostic screen with an error such as `Spin count exceeded - possible deadlock`.

This issue is resolved in this release.

PR 3307435: The hostd service might fail after an update in environments with Virtual Shared Graphics Acceleration (vSGA) and Enhanced vMotion Compatibility (EVC)

In EVC-enabled clusters with active vSGA to share GPU resources across multiple virtual machines, an update operation, such as VIB installation on an ESXi host, might cause the hostd service to fail and start generating frequent coredumps. The issue occurs while adding the GPU-related features to the ESXi host after exiting maintenance mode.

This issue is resolved in this release.

PR 3303651: In a High Performance Plug-in (HPP) environment, a network outage might trigger host failure

In HPP environments, which aim to improve the performance of NVMe devices on ESXi hosts, in case of a network outage, you might see two issues:

1. In case of multi-pathed local SCSI devices in active-active mode claimed by HPP, VM I/Os might fail prematurely instead of being retried on the other path.
2. In case of local SCSI devices claimed by HPP, in an all paths down (APD) scenario, if a complete storage rescan is triggered from vCenter, the ESXi host might fail with a purple diagnostic screen.

Both issues occur only with SCSI local devices and when I/Os fail due to lack of network or APD.

This issue is resolved in this release. The fix makes sure that I/Os fail only in case of `CHECK CONDITION` error indicating a change in a SCSI local device.

ESXi-8.0U2sb-23305545-standard

Profile Name	ESXi-8.0U2sb-23305545-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	February 29, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_vsan_8.0.2-0.25.23305545 - VMware_bootbank_esxio-combiner_8.0.2-0.25.23305545 - VMware_bootbank_infravisor_8.0.2-0.25.23305545 - Mware_bootbank_clusterstore_8.0.2-0.25.23305545 - VMware_bootbank_native-misc-drivers_8.0.2-0.25.23305545 - VMW_bootbank_pensandoatlas_1.46.0.E.28.1.314 - VMware_bootbank_crx_8.0.2-0.25.23305545 - VMware_bootbank_esxio-combiner-esxio_8.0.2-0.25.23305545 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.25.23305545 - VMware_bootbank_esx-xserver_8.0.2-0.25.23305545 - VMware_bootbank_bmc-al-esxio_8.0.2-0.25.23305545 - VMware_bootbank_trx_8.0.2-0.25.23305545 - VMware_bootbank_esxio_8.0.2-0.25.23305545 - VMware_bootbank_esxio-base_8.0.2-0.25.23305545 - VMware_bootbank_gc-esxio_8.0.2-0.25.23305545 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.25.23305545 - VMware_bootbank_bmc-al_8.0.2-0.25.23305545 - VMware_bootbank_gc_8.0.2-0.25.23305545 - VMware_bootbank_vsanhealth_8.0.2-0.25.23305545 - VMware_bootbank_vdfs_8.0.2-0.25.23305545 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.25.23305545 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.25.23305545 - VMware_bootbank_vds-vsip_8.0.2-0.25.23305545 - VMware_bootbank_esx-base_8.0.2-0.25.23305545 - VMware_bootbank_loadesx_8.0.2-0.25.23305545 - VMware_bootbank_esx-update_8.0.2-0.25.23305545 - VMware_bootbank_esxio-update_8.0.2-0.25.23305545 - VMware_bootbank_loadesxio_8.0.2-0.25.23305545 - VMware_locker_tools-light_12.3.5.22544099-23305545
PRs Fixed	3333007, 3328623, 3317046, 3315810, 3309040, 3341832, 3309039, 3293189, 3309041, 3291298, 3304476, 3260854, 3291282, 3291430, 3302555, 3301475, 3291280, 3291297, 3291275, 3291295, 3304809
Related CVE numbers	N/A

This patch updates the following issues:

PR 3316767: VM I/Os might prematurely fail or ESXi hosts crash while HPP claims a local multipath SCSI device

Due to a rare issue with the failover handling in the High Performance Plug-in (HPP) module, if HPP claims a local multipath SCSI device in active-active mode, you might see the following issues:

- in case of connection loss or a NO CONNECT error, VM I/Os can fail prematurely instead of retrying on another path.
- in case of all paths down (APD) state, if a complete storage rescan is triggered from vCenter or the ESXi host, the host might fail with a purple diagnostic screen.

This issue is resolved in this release.

VMware Tools Bundling Changes in ESXi 8.0 Update 2b

- The following VMware Tools ISO images are bundled with ESXi 8.0 Update 2b:
 - **windows.iso**: VMware Tools 12.3.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - **linux.iso**: VMware Tools 10.3.26 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.3.5 Release Notes](#)
- [Earlier versions of VMware Tools](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Updating VMware Tools](#)

ESXi 8.0 Update 2b provides the following security updates:

- OpenSSL is updated to version 3.0.11.
- The Python library is updated to version 3.8.18.
- The Envoy proxy is updated to version 1.25.11.
- The cURL library is updated to version 8.3.0.
- The c-ares C library is updated to version 1.19.1.
- The Go-lang library is updated to version 1.19.11.
- The pyOpenSSL library is updated to version 2.0.
- The Libxslt library is updated to version 1.1.28.
- The libarchive library is updated to version 3.5.1.

ESXi 8.0 Update 2b includes the following Intel microcode:

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Clarkdale	0x20652 (06/25/2)	0x12	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Haswell DT	0x306c3 (06/3c/3)	0x32	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
					D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	0x4003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	0x5003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	0x7002703	3/21/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	0xe000014	9/18/2021	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	0x38	12/2/2021	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	0xd0003b9	9/1/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	0x1000268	9/8/2023	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	0xb4	9/7/2023	Intel Core i3/i5/i7-1100 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	0x34	9/7/2023	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	0x4e	9/7/2023	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	0x2c000321	8/21/2023	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	0x2b000541	8/21/2023	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	0xf4	2/23/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	0xfa	2/27/2023	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	0x5d	9/3/2023	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	0x11e	8/31/2023	Intel Xeon E-2400 Series

ESXi-8.0U2sb-23305545-no-tools

Profile Name	ESXi-8.0U2sb-23305545-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	February 29, 2024
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_vsan_8.0.2-0.25.23305545 - VMware_bootbank_esxio-combiner_8.0.2-0.25.23305545 - VMware_bootbank_infravisor_8.0.2-0.25.23305545 - Mware_bootbank_clusterstore_8.0.2-0.25.23305545 - VMware_bootbank_native-misc-drivers_8.0.2-0.25.23305545 - VMW_bootbank_pensandoatlas_1.46.0.E.28.1.314 - VMware_bootbank_crx_8.0.2-0.25.23305545

	- VMware_bootbank_esxio-combiner-esxio_8.0.2-0.25.23305545 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.2-0.25.23305545 - VMware_bootbank_esx-xserver_8.0.2-0.25.23305545 - VMware_bootbank_bmcad-esxio_8.0.2-0.25.23305545 - VMware_bootbank_trx_8.0.2-0.25.23305545 - VMware_bootbank_esxio_8.0.2-0.25.23305545 - VMware_bootbank_esxio-base_8.0.2-0.25.23305545 - VMware_bootbank_gc-esxio_8.0.2-0.25.23305545 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.2-0.25.23305545 - VMware_bootbank_bmcad_8.0.2-0.25.23305545 - VMware_bootbank_gc_8.0.2-0.25.23305545 - VMware_bootbank_vsanhealth_8.0.2-0.25.23305545 - VMware_bootbank_vdfs_8.0.2-0.25.23305545 - VMware_bootbank_native-misc-drivers-esxio_8.0.2-0.25.23305545 - VMware_bootbank_drivervm-gpu-base_8.0.2-0.25.23305545 - VMware_bootbank_vds-vsiop_8.0.2-0.25.23305545 - VMware_bootbank_esx-base_8.0.2-0.25.23305545 - VMware_bootbank_loadesx_8.0.2-0.25.23305545 - VMware_bootbank_esx-update_8.0.2-0.25.23305545 - VMware_bootbank_esxio-update_8.0.2-0.25.23305545 - VMware_bootbank_loadesxio_8.0.2-0.25.23305545
PRs Fixed	3333007, 3328623, 3317046, 3315810, 3309040, 3341832, 3309039, 3293189, 3309041, 3291298, 3304476, 3260854, 3291282, 3291430, 3302555, 3301475, 3291280, 3291297, 3291275, 3291295
Related CVE numbers	N/A

This patch updates the following issues:

ESXi 8.0 Update 2b provides the following security updates:

- OpenSSL is updated to version 3.0.11.
- The Python library is updated to version 3.8.18.
- The Envoy proxy is updated to version 1.25.11.
- The cURL library is updated to version 8.3.0.
- The c-ares C library is updated to version 1.19.1.
- The Go-lang library is updated to version 1.19.11.
- The pyOpenSSL library is updated to version 2.0.
- The Libxslt library is updated to version 1.1.28.
- The libarchive library is updated to version 3.5.1.

ESXi 8.0 Update 2b includes the following Intel microcode:

Code Name	FMS	Plt ID	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	0x7	4/23/2018	Intel Core i7-620LE Processor

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	0x4003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	0x5003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	0x7002703	3/21/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	0xe000014	9/18/2021	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	0x38	12/2/2021	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	0xd0003b9	9/1/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	0x1000268	9/8/2023	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	0xb4	9/7/2023	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	0x34	9/7/2023	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	0x4e	9/7/2023	Intel Xeon W-11000E Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	0x2c000321	8/21/2023	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	0x2b000541	8/21/2023	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	0xf4	2/23/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	0xfa	2/27/2023	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	0x5d	9/3/2023	Intel Xeon E-2300 Series
Raptor Lake E/HX/S	0xb0671 (06/b7/1)	0x32	0x11e	8/31/2023	Intel Xeon E-2400 Series

ESXi-8.0U2b-23305546

Name	ESXi
Version	ESXi-8.0U2b-23305546
Release Date	February 29, 2024
Category	Bugfix
Affected Components	- ESXi Component - ESXi Install/Upgrade Component - Host Based Replication Server for ESX
PRs Fixed	3326852, 3326624, 3319724, 3316967, 3307362, 3318781, 3331997, 3317441, 3317938, 3336882, 3317707, 3316767, 3306552, 3210610, 3262965, 3312008, 3304392, 3308819, 3313089, 3313912, 3313572, 3308844, 3298815, 3312006, 3311820, 3308054, 3308132, 3302896, 3308133, 3312713, 3301092, 3296088, 3311524, 3311123, 3307579, 3308347, 3309887, 3308458, 3309849, 3303232, 3311989, 3308023, 3308025, 3301665, 3302759, 3305165, 3304472, 3303651, 3293470, 3303807, 3307435, 3282049, 3293466, 3293464, 3296047, 3300173, 3294828, 3300031, 3295243, 3268113, 3284909, 3256705, 3292627, 3287767, 3289239, 3297899, 3295250, 3287747, 3222955, 3300326, 3292550, 3300705, 3287772, 3294624, 3291105, 3296694, 3291219, 3280469, 3288307, 3287769, 3287762, 3289397, 3293199, 3285514, 3275379, 3288872, 3317203, 3313226, 3309606, 3261182, 3286057, 3312725, 3312593, 3311987, 3311956, 3311691, 3295096, 3295060, 3306346, 3306344, 3306345

Related CVE numbers	N/A
---------------------	-----

ESXi-8.0U2sb-23305545

Name	ESXi
Version	ESXi-8.0U2sb-23305545
Release Date	February 29, 2024
Category	Security
Affected Components	• ESXi Component • ESXi Install/Upgrade Component • ESXi Tools Component
PRs Fixed	3333007, 3328623, 3317046, 3315810, 3309040, 3341832, 3309039, 3293189, 3309041, 3291298, 3304476, 3260854, 3291282, 3291430, 3302555, 3301475, 3291280, 3291297, 3291275, 3291295, 3304809
Related CVE numbers	N/A

Known Issues

Miscellaneous Issues

You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot

If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

Workaround: Ignore the compliance errors and proceed with the upgrade.

Installation and Upgrade Issues

You cannot update to ESXi 8.0 Update 2b by using `esxcli software vib` commands

Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.
```

Please use a supported method to upgrade ESXi.

```
vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see [Upgrade or Update a Host with Image Profiles](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a

VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.

2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors

might be Transition to Non-critical from OK, while in the XCC and IPMI, the description is Transition to OK.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

VMware ESXi 8.0 Update 2 Release Notes

This document contains the following sections

- [Introduction](#)
- [General Availability](#)
- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Patches Contained in This Release](#)
- [Product Support Notices](#)
- [Installation and Upgrade Notes for This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 2 | 21 SEP 2023 | GA ISO Build 22380479
Check for additions and updates to these release notes.

General Availability

This [ESXi 8.0 Update 2](#) release is a General Availability (GA) designation. For more information on the vSphere 8.0 IA/GA Release Model of vSphere Update releases, see [The vSphere 8 Release Model Evolves](#).

What's New

For overview of the new features in vSphere 8.0 Update 2, see [What's New in vSphere 8 Update 2](#) and the following release notes:

- [VMware vSphere With Tanzu](#)

- [Tanzu Kubernetes releases](#)
- [VMware vSAN](#)
- [VMware Host Client](#)
- [VMware vSphere+ Release Notes](#)

The list of new features and enhancements that follows adds some of the key highlights for ESXi 8.0 Update 2:

- **With vSphere 8.0 Update 2, vSphere Distributed Services Engine adds support for:**
 - NVIDIA BlueField-2 DPUs to server designs from Fujitsu (Intel Sapphire Rapids).
- **ESXi 8.0 Update 2 adds support to vSphere Quick Boot for multiple servers, including:**
 - Dell
 - VxRail VD-4510c
 - VxRail VD-4520c
 - Fujitsu
 - PRIMERGY RX2540 M7
 - PRIMERGY RX2530 M7
 - Lenovo
 - ThinkSystem SR635 V3
 - ThinkSystem SR645 V3
 - ThinkSystem SR655 V3
 - ThinkSystem ST650 V3

For the full list of supported servers, see the [VMware Compatibility Guide](#).
- **In-Band Error-Correcting Code (IB ECC) support:** With vSphere 8.0 Update 2, you can use IB ECC on hardware platforms that support this option to perform data integrity checks without the need for actual ECC type DDR memory.
- **Support for Graphics and AI/ML workloads on Intel ATS-M:** vSphere 8.0 Update 2 adds support for graphics and AI/ML workloads on Intel ATS-M.
- **Enhanced ESXi CPU Scheduler:** vSphere 8.0 Update 2 adds performance enhancements in the ESXi CPU scheduler for newer generation systems that use high core count CPUs, such as Intel Sapphire Rapids with up to 60 cores per socket.
- **Broadcom lpfc driver update:** With vSphere 8.0 Update 2, the lpfc driver can generate and provide information about Fibre Channel Port Identifier (FPIN) reports.
- **Mellanox (nmmlx5) driver update:** With vSphere 8.0 Update 2, the nmmlx5 driver supports NVIDIA ConnectX-7 SmartNIC and improves performance, such as support for up to 8 Mellanox uplinks per ESXi host, offload support for up to 200G speeds, hardware offloads for inner IPv4/IPv6 checksum offloads (CSO), TCP Segment Offloads (TSO) with outer IPv6 for both GENEVE and VXLAN, and enabling NetQ Receive-Side Scaling (RSS).
- **Marvell (qedentv) driver update:** With vSphere 8.0 Update 2, the qedentv driver supports NetQ Receive-Side Scaling (RSS) and Hardware Large Receive Offload (HW LRO) to enhance performance, scalability, and efficiency.
- **Other driver updates:**
 - Broadcom bcm_mpi3 - bug fixes
 - Broadcom bnxtnet - accumulated bnxtnet driver updates from async driver, including queueGetStats callback of vmk_UplinkQueueOps , enhanced debugging
 - Broadcom lsi_mr3 - routine update
 - Intel icen - enables the RSS feature in native mode for the icen NIC device driver
 - Microchip smartpqi - bug fixes and addition of OEM branded PCI IDs
 - Pensando ionic_cloud: support for cloud vendors
- **IPv6 driver enhancements:** With vSphere 8.0 Update 2, ESXi drivers add offload capabilities for performance improvements with IPv6 when used as an overlay.
- **Uniform Passthrough (UPT) mode support in the nvmxnet3 driver:** vSphere 8.0 Update 2 adds support for UPT to allow faster vSphere vMotion operations in nested ESXi environments.

- **Support 8 ports of 100G on a single host with both Broadcom and Mellanox NICs:** vSphere 8.0 Update 2 increases support for 100GB NIC ports from 4 to 8 in ESXi for Broadcom and Mellanox.
- **CIM Services Tickets for REST Authentication:** In addition to the JWT-based authentication, vSphere 8.0 Update 2 adds the option to authenticate with ESXi host by using CIM services tickets with the `acquireCimServicesTicket()` API for SREST plug-ins.
- **glibc library update:** The glibc library is updated to version 2.28 to align with NIAP requirements.

Guest platform for workloads:

- **Virtual hardware version 21:** vSphere 8.0 Update 2 introduces virtual hardware version 21 to enhance latest guest operating system support and increase maximums for vGPU and vNVME as follows:
 - 16 vGPU devices per VM (see [Configure Virtual Graphics on vSphere](#))
 - 256 vNVMe disks per VM (64 x 4 vNVMe adapters)
 - NVMe 1.3 support for Windows 11 and Windows Server 2022
 - NVMe support for Windows Server Failover Clustering (WSFC)
- **Hot-extend a shared vSphere Virtual Volumes Disk:** vSphere 8.0 Update 2 supports hot extension of shared vSphere Virtual Volumes disks, which allows you to increase the size of a shared disk without deactivating the cluster, and effectively no downtime, and is helpful for VM clustering solutions such as Windows Server Failover Cluster (WSFC). For more information, see [VMware vSphere Virtual Volumes Support for WSFC](#) and [You Can Hot Extend a Shared vVol Disk](#).
- **Use vNVME controller for WSFC:** With vSphere 8.0 Update 2, you can use NVMe controller in addition to existing Paravirtual controller for WSFC with Clustered VMDK for Windows Server 2022 (OS Build 20348.1547) and later. To use NVMe Controller virtual machine hardware version must be 21 or later. For more information, see [Setup for Windows Server Failover Clustering](#).
- **Use vNVME controller for disks in multi-writer mode:** Starting with vSphere 8.0 Update 2, you can use virtual NVMe controllers for virtual disks with multi-writer mode, used by third party cluster applications such as Oracle RAC.
- **USB 3.2 Support in virtual eXtensible Host Controller Interface (xHCI):** With vSphere 8.0 Update 2, the virtual xHCI controller is 20 Gbps compatible.
- **Read-only mode for attached virtual disks:** With vSphere 8.0 Update 2, you can attach a virtual disk as read-only to a virtual machine to avoid temporary redo logs and improve performance for use cases such as VMware App Volumes.
- **Support VM clone when a First Class Disk (FCD) is attached:** With vSphere 8.0 Update 2, you can use the `cloneVM()` API to clone a VM with a FCD attached.

GPU

- **GPU Driver VM for any passthrough GPU card:** With vSphere 8.0 Update 2, a GPU Driver VM facilitates support of new GPU vendors for the virtual SVGA device (vSGA).

Storage

- **Support for multiple TCP connections on a single NFS v3 volume:** With vSphere 8.0 Update 2, the capability that allows support of multiple TCP connections for a single NFS v3 volume by using ESXCLI, nConnect, becomes fully supported for on-prem environments. For more information, see VMware knowledge base articles [91497](#) and [91479](#).
- **ESXCLI support for SCSI UNMAP operations for vSphere Virtual Volumes:** Starting with vSphere 8.0 Update 2, you can use command line, ESXCLI, for SCSI UNMAP operations for vSphere Virtual Volumes.

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0 Update 1c Release Notes](#)
- [VMware ESXi 8.0 Update 1a Release Notes](#)
- [VMware ESXi 8.0 Update 1 Release Notes](#)

- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Patches Contained in This Release

VMware ESXi 8.0 Update 2

Build Details

Download Filename:	VMware-ESXi-8.0U2-22380479-depot.zip
Build:	22380479
Download Size:	615.7 MB
sha256checksum:	76dbd1f0077b20b765c8084e74a5fd012c0c485297f60530e1f9703bdf051b5c
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi80U2-22380479	Enhancement	Important

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U2-22380479-standard
ESXi-8.0U2-22380479-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi_8.0.2-0.0.22380479	09/21/2023	General	Bugfix image

For information about the individual components and bulletins, see the [Resolved Issues](#) section.

Patch Download and Installation

You can download ESXi 8.0 Update 2 from the [Broadcom Support Portal](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of

vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Product Support Notices

- **Restricting the use of the ptrace system call by default:** Starting with vSphere 8.0 Update 2, VMware restricts the use of the `ptrace` system call by default due to security concerns.
- **Deprecation of the Internet Wide Area RDMA Protocol (iWARP) on Intel X722 NIC hardware:** With vSphere 8.0 Update 2, iWARP on Intel X722 NIC hardware is no longer supported.
- **Deprecation of legacy SOAP API contracts:** With vSphere 8.0 Update 2, WSDL-based SOAP API contracts of vSphere 5.5 and earlier releases are deprecated and will be removed in a future vSphere release. All vSphere automation based on VMware or third-party tooling that use vSphere 5.5 or earlier contracts must be upgraded to more recent tooling.
- **Default NVMe version 1.3 for Windows virtual machines:** Starting with vSphere 8.0 Update 2, the NVMe version of Windows Server 2022 or Windows 11 and later virtual machines with hardware version 21 and later, and a virtual NVMe controller, is set to 1.3 by default. For proper NVMe hot add, remove, and extend operations, you must upgrade Windows at least to KB5029250 (OS Build 20348.1906) for Windows Server 2022 and KB5028254 (OS Build 22621.2070) for Windows 11.
- **TLS 1.3 support:** vSphere 8.0 Update 2 introduces initial support for TLS 1.3 in ESXi. vCenter stays on TLS 1.2.
- **Support of NVMe/TCP in Enhanced Networking Stack (ENS) mode for the qedentv driver:** Starting with vSphere 8.0 Update 2, the VMware qedentv NIC driver from Marvell supports NVMe/TCP in ENS mode.

Installation and Upgrade Notes for This Release

- When updating vSphere Distributed Service Engine hosts configured with NVIDIA BlueField-2 (BF2) DPUs to ESXi 8.0 Update 2, you must follow a specific order for ESXi and BF2 DPU firmware upgrade in the same maintenance window:
 - a. Enter the vSphere Distributed Service Engine host into maintenance mode.
 - b. Make sure that the NIC and ARM firmware of the BF2 DPU are not upgraded and they still run firmware applicable for ESXi versions earlier than 8.0 Update 2, such as NIC firmware version 24.33.1246 and ARM firmware version 18.2.0.12580.
 - c. Update ESXi to 8.0 Update 2.
 - d. Upgrade the BF2 DPU ARM (UEFI & ATF) firmware to 4.0.2.12722, which is the firmware version required for ESXi 8.0 Update 2.
 - e. Upgrade the BF2 DPU NIC firmware to 24.36.7506, which is the firmware version required for ESXi 8.0 Update 2 drivers.
 - f. Exit the vSphere Distributed Service Engine host from maintenance mode.

If you do not follow the steps in this order, ESXi rollback fails and you cannot recover the BF2 DPU based vSphere Distributed Service Engine host configuration if upgrade fails.

• VMware Tools Bundling Changes in ESXi 8.0 Update 2

– The following VMware Tools ISO images are bundled with ESXi 8.0 Update 2:

- **windows.iso**: VMware Tools 12.3.0 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
- **linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- **darwin.iso**: Supports Mac OS X versions 10.11 and later.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [VMware Tools 12.3.0 Release Notes](#)
- [Earlier versions of VMware Tools](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [Updating VMware Tools](#)

Resolved Issues

Security Issues

ESXi 8.0 Update 2 includes the following Intel microcode:

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	0x1d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	0x11	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	0x7	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	0x2f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	0x1f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	0x621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	0x71a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	0xd	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	0x3b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	0x21	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	0x28	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	0x42e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	0x715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Haswell EP	0x306f2 (06/3f/2)	0x6f	0x49	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	0x1a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	0x22	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	0x12d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	0xb000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	0x2007006	3/6/2023	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	0x4003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	0x5003604	3/17/2023	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Cooper Lake	0x5065b (06/55/b)	0xbf	0x7002703	3/21/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	0x1c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	0x700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	0xf00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	0xe000014	9/18/2021	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	0xf0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	0x38	12/2/2021	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	0xd0003a5	3/30/2023	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Silver 4300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	0x1000230	1/27/2023	Intel Xeon D-2700 Series; Intel Xeon D-1700 Series
Snow Ridge	0x80665 (06/86/5)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	0x4c000023	2/22/2023	Intel Atom P5000 Series
Tiger Lake U	0x806c1 (06/8c/1)	0x80	0xac	2/27/2023	Intel Core i3/i5/i7-1100 Series
Tiger Lake U Refresh	0x806c2 (06/8c/2)	0xc2	0x2c	2/27/2023	Intel Core i3/i5/i7-1100 Series
Tiger Lake H	0x806d1 (06/8d/1)	0xc2	0x46	2/27/2023	Intel Xeon W-11000E Series
Sapphire Rapids SP HBM	0x806f8 (06/8f/8)	0x10	0x2c0001d1	2/14/2023	Intel Xeon Max 9400 Series
Sapphire Rapids SP	0x806f8 (06/8f/8)	0x87	0x2b000461	3/13/2023	Intel Xeon Platinum 8400 Series; Intel Xeon Gold 6400/5400 Series; Intel Xeon Silver 4400 Series; Intel Xeon Bronze 3400 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	0xf4	2/23/2023	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Coffee Lake	0x906ea (06/9e/a)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	0xf4	2/23/2023	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	0xfa	2/27/2023	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	0x59	2/26/2023	Intel Xeon E-2300 Series

ESXi 8.0 Update 2 includes the following AMD CPU microcode:

Code Name	FMS	MCU Rev	MCU Date	Brand Names
Bulldozer	0x600f12 (15/01/2)	0x0600063e	2/7/2018	Opteron 6200/4200/3200 Series
Piledriver	0x600f20 (15/02/0)	0x06000852	2/6/2018	Opteron 6300/4300/3300 Series
Zen-Naples	0x800f12 (17/01/2)	0x0800126e	11/11/2021	EPYC 7001 Series
Zen2-Rome	0x830f10 (17/31/0)	0x0830107a	5/17/2023	EPYC 7002/7Fx2/7Hx2 Series
Zen3-Milan-B1	0xa00f11 (19/01/1)	0x0a0011d1	7/10/2023	EPYC 7003/7003X Series
Zen3-Milan-B2	0xa00f12 (19/01/2)	0x0a001234	7/10/2023	EPYC 7003/7003X Series

Installation, Upgrade and Migration Issues

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS.

This issue is resolved in this release.

Miscellaneous Issues

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
```

```
Validation of vmknfc failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

This issue is resolved in this release.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

This issue is resolved in this release.

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

This issue is resolved in this release.

Networking Issues

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

This issue is resolved in this release.

Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations

vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

This issue is resolved in this release.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

This issue is resolved in this release.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

This issue is resolved in this release.

Known Issues

Miscellaneous Issues

If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

In a High Performance Plug-in (HPP) environment, a network outage might trigger host failure

In HPP environments, which aim to improve the performance of NVMe devices on ESXi hosts, in case of a network outage, you might see two issues:

1. In case of multi-pathed local SCSI devices (active-active presentation) claimed by HPP, VM I/Os might fail prematurely instead of being retried on the other path.
2. In case of local SCSI devices claimed by HPP, in an all paths down (APD) scenario, if a complete storage rescan is triggered from vCenter, the ESXi host might fail with a purple diagnostic screen.

Both issues occur only with SCSI local devices and when I/Os fail due to lack of network or APD.

Workaround: Add user claimrules to claim devices by the ESXi native multipathing plug-in (NMP) instead of HPP. For more information, see VMware knowledge base article [94377](#).

TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack

If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

Security Issues

Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot

Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see [Specify Lockdown Mode Exception Users](#) and [Specify Lockdown Mode Exception Users in the VMware Host Client](#).

Networking Issues

Transfer speed in IPv6 environments with active TCP segmentation offload is slow

In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop

When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option **Config.HostAgent.ssl.keyStore.allowSelfSigned** is set to **false** by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is not active, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some nmlx5_core driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss` : Use the DRSS parameter.
- `drss` : Use the DRSS parameter.
- `rss` : Use the RSS parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 1d Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0 Update 1d | 05 MAR 2024 | Build 23299997

Check for additions and updates to these release notes.

Important:

ESXi 8.0 Update 1d delivers fixes for CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, and CVE-2024-22255. If you do not plan to update your environment to ESXi 8.0 Update 2b (build # 23305546), use 8.0 Update 1d to update your ESXi hosts of version 8.0 Update 1c (build # 22088125) and earlier with these security fixes. The supported update path from 8.0 Update 1d is to ESXi 8.0 Update 2b or later.

CAUTION:

Updates from ESXi 8.0 Update 1d to 8.0 Update 2 might expose your vSphere system to security vulnerabilities, because this is considered a back-in-time update. For more information, see the [Product Interoperability Matrix](#).

What's New

This [release](#) resolves CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, and CVE-2024-22255. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0006](#).

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0 Update 1c Release Notes](#)
- [VMware ESXi 8.0 Update 1a Release Notes](#)
- [VMware ESXi 8.0 Update 1 Release Notes](#)
- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Patches Contained in This Release

Patch for VMware ESXi 8.0 Update 1d

This release of ESXi 8.0 Update 1d delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U1d-23299997-depot.zip
Build:	23299997
Download Size:	585.2 MB
sha256checksum:	eb65a987843db3e7da326b3a03f4770d5d5f5b165d0f160e04c43d ccbc9e64bf
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.1-0.35.23299997	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.1-0.35.23299997	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.1-0.35.23299997	Security	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U1d-23299997	Security	Critical	Security only fixes

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U1d-23299997-standard
ESXi-8.0U1d-23299997-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi 8.0 U1d - 23299997	03/05/2024	Security	Security only image

For information about the individual components and bulletins, see the [Resolved Issues](#) section.

Patch Download and Installation

You can download ESXi 8.0 Update 1d from the [Broadcom Support Portal](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of

vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.1-0.35.23299997

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.35.23299997 - VMware_bootbank_esxio_8.0.1-0.35.23299997 - VMware_bootbank_native-misc-drivers_8.0.1-0.35.23299997 - VMware_bootbank_vds-vsip_8.0.1-0.35.23299997 - VMware_bootbank_trx_8.0.1-0.35.23299997 - VMware_bootbank_bmcad-esxio_8.0.1-0.35.23299997 - VMware_bootbank_bmcad_8.0.1-0.35.23299997 - VMware_bootbank_vsanhealth_8.0.1-0.35.23299997 - VMware_bootbank_esxio-combiner-esxio_8.0.1-0.35.23299997 - VMware_bootbank_gc_8.0.1-0.35.23299997 - VMware_bootbank_cpu-microcode_8.0.1-0.35.23299997 - VMware_bootbank_vdfs_8.0.1-0.35.23299997 - VMware_bootbank_esxio-base_8.0.1-0.35.23299997 - VMware_bootbank_esxio-combiner_8.0.1-0.35.23299997 - VMware_bootbank_gc-esxio_8.0.1-0.35.23299997 - VMware_bootbank_crx_8.0.1-0.35.23299997 - VMware_bootbank_vsan_8.0.1-0.35.23299997 - VMware_bootbank_esx-xserver_8.0.1-0.35.23299997 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.35.23299997 - VMware_bootbank_esx-base_8.0.1-0.35.23299997 - VMware_bootbank_clusterstore_8.0.1-0.35.23299997 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.35.23299997
PRs Fixed	N/A
CVE numbers	CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, CVE-2024-22255

This patch updates the `esx-base` VIB. Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `native-misc-drivers-esxio`, `esxio`, `native-misc-drivers`, `vds-vsip`, `trx`, `bmcad-esxio`, `bmcad`, `vsanhealth`, `esxio-combiner-esxio`, `gc`, `cpu-microcode`, `vdfs`, `esxio-base`, `esxio-combiner`, `gc-esxio`, `crx`, `vsan`, `esx-xserver`,

`esxio-dvfilter-generic-fastpath`, `clusterstore`, and `esx-dvfilter-generic-fastpath`. This patch resolves the following issue:

This release resolves CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, and CVE-2024-22255. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0006](#).

esx-update_8.0.1-0.35.23299997

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_loadesx_8.0.1-0.35.23299997 - VMware_bootbank_esx-update_8.0.1-0.35.23299997
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesx` and `esx-update`.

esxio-update_8.0.1-0.35.23299997

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_loadesxio_8.0.1-0.35.23299997 - VMware_bootbank_esxio-update_8.0.1-0.35.23299997
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency on the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `loadesxio` and `esxio-update`.

ESXi-8.0U1d-23299997-standard

Profile Name	ESXi-8.0U1d-23299997-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	March 05, 2024
Acceptance Level	Partner Supported

Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.35.23299997 - VMware_bootbank_esxio_8.0.1-0.35.23299997 - VMware_bootbank_native-misc-drivers_8.0.1-0.35.23299997 - VMware_bootbank_vds-vsip_8.0.1-0.35.23299997 - VMware_bootbank_trx_8.0.1-0.35.23299997 - VMware_bootbank_bmcad-esxio_8.0.1-0.35.23299997 - VMware_bootbank_bmcad_8.0.1-0.35.23299997 - VMware_bootbank_vsanhealth_8.0.1-0.35.23299997 - VMware_bootbank_esxio-combiner-esxio_8.0.1-0.35.23299997 - VMware_bootbank_gc_8.0.1-0.35.23299997 - VMware_bootbank_cpu-microcode_8.0.1-0.35.23299997 - VMware_bootbank_vdfs_8.0.1-0.35.23299997 - VMware_bootbank_esxio-base_8.0.1-0.35.23299997 - VMware_bootbank_esxio-combiner_8.0.1-0.35.23299997 - VMware_bootbank_gc-esxio_8.0.1-0.35.23299997 - VMware_bootbank_crx_8.0.1-0.35.23299997 - VMware_bootbank_vsan_8.0.1-0.35.23299997 - VMware_bootbank_esx-xserver_8.0.1-0.35.23299997 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.35.23299997 - VMware_bootbank_esx-base_8.0.1-0.35.23299997 - VMware_bootbank_clusterstore_8.0.1-0.35.23299997 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.35.23299997 - VMware_bootbank_loadesx_8.0.1-0.35.23299997 - VMware_bootbank_esx-update_8.0.1-0.35.23299997 - VMware_bootbank_loadesxio_8.0.1-0.35.23299997 - VMware_bootbank_esxio-update_8.0.1-0.35.23299997
PRs Fixed	N/A
Related CVE numbers	CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, CVE-2024-22255

This patch updates the following issue:

This release resolves CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, and CVE-2024-22255. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0006](#).

ESXi-8.0U1d-23299997-no-tools

Profile Name	ESXi-8.0U1d-23299997-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware by Broadcom, Inc.
Release Date	March 05, 2024
Acceptance Level	Partner Supported

Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.35.23299997 - VMware_bootbank_esxio_8.0.1-0.35.23299997 - VMware_bootbank_native-misc-drivers_8.0.1-0.35.23299997 - VMware_bootbank_vds-vsip_8.0.1-0.35.23299997 - VMware_bootbank_trx_8.0.1-0.35.23299997 - VMware_bootbank_bmcal-esxio_8.0.1-0.35.23299997 - VMware_bootbank_bmcal_8.0.1-0.35.23299997 - VMware_bootbank_vsanhealth_8.0.1-0.35.23299997 - VMware_bootbank_esxio-combiner-esxio_8.0.1-0.35.23299997 - VMware_bootbank_gc_8.0.1-0.35.23299997 - VMware_bootbank_cpu-microcode_8.0.1-0.35.23299997 - VMware_bootbank_vdfs_8.0.1-0.35.23299997 - VMware_bootbank_esxio-base_8.0.1-0.35.23299997 - VMware_bootbank_esxio-combiner_8.0.1-0.35.23299997 - VMware_bootbank_gc-esxio_8.0.1-0.35.23299997 - VMware_bootbank_crx_8.0.1-0.35.23299997 - VMware_bootbank_vsan_8.0.1-0.35.23299997 - VMware_bootbank_esx-xserver_8.0.1-0.35.23299997 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.35.23299997 - VMware_bootbank_esx-base_8.0.1-0.35.23299997 - VMware_bootbank_clusterstore_8.0.1-0.35.23299997 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.35.23299997 - VMware_bootbank_loadesx_8.0.1-0.35.23299997 - VMware_bootbank_esx-update_8.0.1-0.35.23299997 - VMware_bootbank_loadesxio_8.0.1-0.35.23299997 - VMware_bootbank_esxio-update_8.0.1-0.35.23299997
PRs Fixed	N/A
Related CVE numbers	CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, CVE-2024-22255

This patch updates the following issue:

This release resolves CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, and CVE-2024-22255. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0006](#).

ESXi 8.0 U1d - 23299997

Name	ESXi
Version	ESXi80U1d-23299997
Release Date	March 05, 2024
Category	Security
Affected Components	- ESXi Component - ESXi Install/Upgrade Component

PRs Fixed	N/A
Related CVE numbers	CVE-2024-22252, CVE-2024-22253, CVE-2024-22254, CVE-2024-22255

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail

as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some nmlx5_core driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00::xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
```

```
Validation of vmknics failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the `ntg3` driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 1c Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0 Update 1c | 27 JUL 2023 | Build 22088125
Check for additions and updates to these release notes.

What's New

[ESXi 8.0 Update 1c](#) supports vSphere Quick Boot on the following servers:

- Cisco Systems Inc
 - UCSC-C225-M6S

- Dell Inc.
 - R660 vSAN Ready Node
 - R760 vSAN Ready Node
 - PowerEdge R660xs
 - PowerEdge R760xd2
 - PowerEdge R760xa
 - PowerEdge R760xs
 - PowerEdge R860
 - PowerEdge R960
 - PowerEdge T560
- HPE
 - Alletra 4120
 - HPE Cray XD220v
 - ProLiant DL320 Gen11
 - ProLiant DL360 Gen11
 - ProLiant DL380 Gen11
 - ProLiant DL380a Gen11
 - ProLiant DL560 Gen11
 - ProLiant ML110 Gen11
 - ProLiant ML350 Gen11
- Lenovo
 - ThinkSystem SR630 V3
 - ThinkSystem SR650 V3

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0 Update 1a Release Notes](#)
- [VMware ESXi 8.0 Update 1 Release Notes](#)
- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Patches Contained in This Release

Patch for VMware ESXi 8.0 Update 1c

This release of ESXi 8.0 Update 1c delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U1c-22088125-depot.zip
Build:	22088125
Download Size:	949.1 MB
sha256checksum:	ab25ede2b6c40d6bb551f41ae43187fe124cc83853e9df25cd79ccf6836546a8

Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.1-0.25.22088125	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.1-0.25.22088125	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.1-0.25.22088125	Bugfix	Critical
Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers for VMware ESXi	Mellanox-nmlx5_4.23.0.36-15vmw.801.0.25.22088125	Bugfix	Critical
Broadcom NetXtreme I ESX VMKAPI ethernet driver	Broadcom-ntg3_4.1.10.0-5vmw.801.0.25.22088125	Bugfix	Critical
VMware NVMe over TCP Driver	VMware-NVMeoF-TCP_1.0.1.7-1vmw.801.0.25.22088125	Bugfix	Critical
ESXi Component - core ESXi VIBs	ESXi_8.0.1-0.20.22082334	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.1-0.20.22082334	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.1-0.20.22082334	Security	Critical
ESXi Tools Component	VMware-VM-Tools_12.2.5.21855600-22082334	Security	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity	Detail
ESXi80U1c-22088125	Bugfix	Critical	Security fixes and Bug fixes
ESXi80U1sc-22082334	Security	Critical	Security only fixes

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U1c-22088125-standard
ESXi-8.0U1c-22088125-no-tools
ESXi-8.0U1sc-22082334-standard
ESXi-8.0U1sc-22082334-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi 8.0 U1c - 22088125	07/27/2023	Bugfix	Security and Bugfix image
ESXi 8.0 U1sc - 22082334	07/27/2023	Security	Security only image

For information about the individual components and bulletins, see the [Resolved Issues](#) section.

Patch Download and Installation

You can download ESXi 8.0 Update 1c from the [Broadcom Support Portal](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

ESXi_8.0.1-0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-combiner-esxio_8.0.1-0.25.22088125 - VMware_bootbank_esxio-base_8.0.1-0.25.22088125 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.25.22088125 - VMware_bootbank_vdfs_8.0.1-0.25.22088125 - VMware_bootbank_bmcad_8.0.1-0.25.22088125 - VMware_bootbank_clusterstore_8.0.1-0.25.22088125 - VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.25.22088125 - VMware_bootbank_gc-esxio_8.0.1-0.25.22088125 - VMware_bootbank_esxio_8.0.1-0.25.22088125 - VMware_bootbank_bmcad-esxio_8.0.1-0.25.22088125 - VMware_bootbank_esxio-combiner_8.0.1-0.25.22088125 - VMware_bootbank_gc_8.0.1-0.25.22088125 - VMware_bootbank_esx-xserver_8.0.1-0.25.22088125 - VMware_bootbank_vsan_8.0.1-0.25.22088125 - VMware_bootbank_esx-base_8.0.1-0.25.22088125 - VMware_bootbank_trx_8.0.1-0.25.22088125 - VMware_bootbank_native-misc-drivers_8.0.1-0.25.22088125 - VMware_bootbank_cpu-microcode_8.0.1-0.25.22088125 - VMware_bootbank_vsanhealth_8.0.1-0.25.22088125 - VMware_bootbank_crx_8.0.1-0.25.22088125

PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236, 3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3038908, 3179111
CVE numbers	N/A

Updates the `esxio-combiner-esxio`, `esxio-base`, `esxio-dvfilter-generic-fastpath`, `vdfs`, `bmcal`, `clusterstore`, `native-misc-drivers-esxio`, `gc-esxio`, `esxio`, `bmcal-esxio`, `esxio-combiner`, `gc`, `esx-xserver`, `vsan`, `esx-base`, `trx`, `native-misc-drivers`, `cpu-microcode`, `vsanhealth`, and `crx` VIBs to resolve the following issues:

Transient vSAN health check warning: Network configuration is out of sync

vSAN Skyline health might randomly report that network configuration is out of sync. This transient issue occurs when the vSAN health service uses an outdated vCenter configuration to perform unicast check.

This issue is resolved in this release.

A race condition between two operations could potentially lead to inconsistent metadata

In extremely rare cases, a race condition between an unmap operation on a VMDK and a concurrent snapshot creation operation on the same object could cause inconsistent vSAN metadata for that object in a vSAN host. This inconsistency could potentially lead to VM failure or host failure.

This issue is resolved in this release.

ESX hosts might fail with a purple diagnostic screen and an error NMI IPI: Panic requested by another PCPU

The resource pool cache is a VMFS specific volume level cache that stores the resource clusters corresponding to the VMFS volume. While searching for priority clusters, the cache flusher workflow iterates through a large list of cached resource clusters, which can cause lockup of the physical CPUs. As a result, ESX hosts might fail with a purple diagnostic screen. In the `logDump` file, you see an error such as:

```
^[[7m2022-10-22T07:56:47.322Z cpu13:2101160)WARNING: Heartbeat: 827: PCPU 0 didn't have a
heartbeat for 7 seconds, timeout is 14, 1 IPIs sent; *may* be locked up.^[[0m
^[[31;1m2022-10-22T07:56:47.322Z cpu0:2110633)ALERT: NMI: 710: NMI IPI:
RIPOFF(base):RBP:CS
```

This issue is resolved in this release.

ESXi ConfigStore database fills up and writes fail

Stale data related to block devices might not be deleted in time from the ESXi ConfigStore database and cause an out of space condition. As a result, write operations to ConfigStore start to fail. In the backtrace, you see logs such as:

2022-12-19T03:51:42.733Z cpu53:26745174)WARNING: VisorFSRam: 203: Cannot extend visorfs file /etc/vmware/configstore/current-store-1-journal because its ramdisk (configstore) is full.

This issue is resolved in this release.

Delay before evacuating vSAN storage device in UNHEALTHY state

After detecting devices as UNHEALTHY, the Local Log-Structured Object Manager (LSOM) might wait for 10 minutes (LSOM_DEVICE_MONITORING_INTERVAL) before initiating evacuation on these devices.

This issue is resolved in this release.

vSAN precheck for maintenance mode or disk decommission doesn't list objects that might lose accessibility

This issue affects objects with resyncing components, and some components reside on a device to be removed or placed into maintenance mode. When you run a precheck with the **No-Action** option, the precheck does not evaluate the object correctly to report it in the `inaccessibleObjects` list.

This issue is resolved in this release. Precheck includes all affected objects in the `inaccessibleObjects` list.

The vSAN iSCSI Target service might fail due to a rare race condition

When you run the ESXCLI command `esxcli network firewall load`, the operation does not reload existing dynamic firewall rules and they are lost. When you run the ESXCLI command `esxcli network firewall refresh`, the operation reloads existing dynamic firewall rules, but in some cases a race condition might cause some rules to be lost. The issue occurs only when multiple firewall refresh commands run at the same time, which leads to the race condition. As a result, the vSAN iSCSI Target service might fail.

This issue is resolved in this release.

Setting the screen resolution of a powered off, encrypted virtual machine does not always work when manually editing a VMX file

If you manually specify the screen resolution of a powered off, encrypted virtual machine by editing the VMX file, the change might not take effect.

This issue is resolved in this release.

If the Internet Control Message Protocol (ICMPA) is not active, ESXi host reboot might take long after upgrading to vSphere 8.0 and later

If ICMPA is not active on the NFS servers in your environment, after upgrading your system to vSphere 8.0 and later, ESXi hosts reboot might take an hour to complete, because restore operations for NFS datastores fail. NFS uses the `vmkping` utility to identify reachable IPs of the NFS servers before executing a mount operation and when ICMP is not active, mount operations fail.

This issue is resolved in this release. To remove dependency on the ICMP protocol to find reachable IPs, the fix adds socket APIs to ensure that IPs on a given NFS server are available.

When you migrate a VM with recently hot-added memory, an ESXi host might repeatedly fail with a purple diagnostic screen

Due to a race condition while the memory hotplug module recomputes the NUMA memory layout of a VM on a destination host after migration, an ESXi host might repeatedly fail with a purple diagnostic screen. In the backtrace, you see errors such as:

```
0x452900262cf0:[0x4200138fee8b]PanicvPanicInt@vmkernel#nover+0x327 stack: 0x452900262dc8,
0x4302f6c06508, 0x4200138fee8b, 0x420013df1300, 0x452900262cf0 0x452900262dc0:
[0x4200138ff43d]Panic_WithBacktrace@vmkernel#nover+0x56 stack: 0x452900262e30,
0x452900262de0, 0x452900262e40, 0x452900262df0, 0x3e7514 0x452900262e30:
[0x4200138fbb90]NMI_Interrupt@vmkernel#nover+0x561 stack: 0x0, 0xf48, 0x0,
```

```

0x0, 0x0 0x452900262f00:[0x420013953392]IDTNMIWork@vmkernel#nover+0x7f stack:
0x420049800000, 0x4200139546dd, 0x0, 0x452900262fd0, 0x0 0x452900262f20:
[0x4200139546dc]Int2_NMI@vmkernel#nover+0x19 stack: 0x0, 0x42001394e068,
0xf50, 0xf50, 0x0 0x452900262f40:[0x42001394e067]gate_entry@vmkernel#nover
+0x68 stack: 0x0, 0x43207bc02088, 0xd, 0x0, 0x43207bc02088 0x45397b61bd30:
[0x420013be7514]NUMASched_PageNum2PhysicalDomain@vmkernel#nover+0x58 stack:
0x1, 0x420013be34c3, 0x45396f79f000, 0x1, 0x100005cf757 0x45397b61bd50:
[0x420013be34c2]NUMASched_UpdateAllocStats@vmkernel#nover+0x4b stack: 0x100005cf757, 0x0,
0x0, 0x4200139b36d9, 0x0 0x45397b61bd80:[0x4200139b36d8]VmMem_NodeStatsSub@vmkernel#nover
+0x59 stack: 0x39, 0x45396f79f000, 0xbce0dbf, 0x100005cf757, 0x0 0x45397b61bdc0:
[0x4200139b4372]VmMem_FreePageNoBackmap@vmkernel#nover+0x8b stack: 0x465ec0001be0, 0xa,
0x465ec18748b0, 0x420014e7685f, 0x465ec14437d0

```

This issue is resolved in this release.

vSAN cluster shutdown fails on a cluster with IPV6 disabled

This issue occurs on vSAN hosts running ESXi 8.0 Update 1 with IPv6 disabled. When you use the vSAN cluster shutdown wizard, the workflow fails with the following error message: 'NoneType' object is not iterable.

This issue is resolved in this release.

You see trap files in a SNMP directory under /var/spool even though SNMP is not enabled

After the hostd service starts, for example after an ESXi host reboot, it might create a SNMP directory under /var/spool and you see many .trap files to pile up in this directory.

This issue is resolved in this release. The fix makes sure that the directory /var/spool/snmp exists only when SNMP is enabled.

The hostd service repeatedly fails and the ESXi host disconnects from the vCenter system

If for any reason an ESXi host is temporarily in a state of insufficient memory, the hostd service might repeatedly fail due to a vSphere Replication filter that prevents the allocation of bitmaps. As a result, the ESXi host disconnects from the vCenter system and cannot connect back.

This issue is resolved in this release.

During service insertion for NSX-managed workload VMs, some VMS might become intermittently unresponsive and the virtual device might reset

During service insertion for NSX-managed workload VMs, a packet list might be reinjected from the input chain of one switchport to another switchport. In such cases, the source switchport does not correspond to the actual portID of the input chain and the virtual device does not get completion status for the transmitted frame. As a result, when you run a service insertion task, some VMs might become intermittently unresponsive due to network connectivity issues.

This issue is resolved in this release.

Performance of certain nested virtual machines on AMD CPUs might degrade

Nested virtual machines on AMD CPUs with operational systems such as Windows with virtualization-based security (VBS) might experience performance degradation, timeouts, or unresponsiveness due to an issue with the virtualization of AMD's Rapid Virtualization Indexing (RVI), also known as Nested Page Tables (NPT).

This issue is resolved in this release.

Changing the mode of the virtual disk on a running virtual machine might cause the VM to fail

If you use the VMware Host Client to edit the disk mode of a running virtual machine, for example from Independent - Nonpersistent to Dependent or Independent - Persistent, the operation fails and might cause the VM to fail. In the vmware.log, you see errors such as:

`msg.disk.notConfigured2]` Failed to configure disk 'scsi0:4'. The virtual machine cannot be powered on with an unconfigured disk.

`[msg.checkpoint.continuesync.error]` An operation required the virtual machine to quiesce and the virtual machine was unable to continue running.

This issue is resolved in this release. The fix blocks changing the mode of an Independent - Nonpersistent disk on a running virtual machine by using the VMware Host Client. The vSphere Client already blocks such operations.

ESXi NVMe/TCP initiator fails to recover paths after target failure recovery

When an NVMe/TCP target recovers from a failure, ESXi cannot recover the path.

This issue is resolved in this release.

ESXi host becomes unresponsive and you cannot put the host in Maintenance Mode or migrate VMs from that host

Asynchronous reads of metadata on a VMFS volume attached to an ESXi host might cause a race condition with other threads on the host and make the host unresponsive. As a result, you cannot put the host in Maintenance Mode or migrate VMs from that host.

This issue is resolved in this release.

A Logical Volume Manager (LVM) disk goes offline during datastore expansion

If during the expansion of a datastore on an ESXi host in a cluster you run a storage refresh, a LVM extent might go offline and virtual machines on this volume become unresponsive. The issue occurs because the storage refresh operation triggers a slow refresh of the volume attributes on all the ESXi hosts in the cluster. As a result, LVM metadata on the disk might not match the cached capacity information in the Pluggable Storage Architecture (PSA) layer and ESXi marks the LVM extent offline for the safety of metadata and data.

This issue is resolved in this release.

The durable name of a SCSI LUN might not be set

The durable name property for a SCSI-3 compliant device comes from pages 80h and 83h of the Vital Product Data (VPD) as defined by the T10 and SMI standards. To populate the durable name, ESXi first sends an inquiry command to get a list of VPD pages supported by the device. Then ESXi issues commands to get data for all supported VPD pages. Due to an issue with the target array, the device might fail a command to get VPD page data for a page in the list with a `not supported` error. As a result, ESXi cannot populate the durable name property for the device.

This issue is resolved in this release. The fix ignores the error on command to get VPD page data, except pages 80h and 83h, if that data is not required for the generation of durable name.

Excessive time for path recovery with ESXi High Performance Plug-in (HPP) after a link up event for NVMe over Fibre Channel

In certain scenarios, after a Fibre Channel link up event, it might take as long as 5 minutes for HPP managed NVMe over Fibre Channel paths to recover.

This issue is resolved in this release.

NVMe over Fabrics controller might unexpectedly disconnect during discovery

If discovery and I/O controllers already exist for a NVMe over Fabrics storage target that supports a persistent discovery controller on an ESXi host, a concurrent NVMe over Fabrics controller discovery operation might cause some I/O controllers to disconnect unexpectedly.

This issue is resolved in this release.

If your ESXi PTP service uses hardware timestamping, enabling IGMP snooping on switches might cause synchronization failures

When you enable Internet Group Management Protocol (IGMP) snooping on a connected switch, the Precision Time Protocol (PTP) client needs to send IGMP multicast requests to receive PTP multicast stream from the grandmaster. If the ESXi PTP agent is based on hardware timestamping, the agent might fail to send IGMP join/leave requests to the switch. As a result, the PTP multicast stream cannot go forward to the ESXi host and prevents proper PTP synchronization.

This issue is resolved in this release. For more details, see VMware knowledge base article [92276](#).

If parallel volume expand and volume refresh operations on the same VMFS volume run on two ESXi hosts in the same cluster, the VMFS volume might go offline

While a VMFS volume expand operation is in progress on an ESXi host in a vCenter cluster, if on another host a user or vCenter initiates a refresh of the same VMFS volume capacity, such a volume might go offline. The issue occurs due to a possible mismatch in the device size, which is stamped on the disk in the volume metadata during a device rescan, and the device size value in the Pluggable Storage Architecture (PSA) layer on the host, which might not be updated if the device rescan is not complete.

This issue is resolved in this release. The fix improves the resiliency of the volume manager code to force a consecutive refresh of the device attributes and comparison of the device sizes again if vCenter reports a mismatch in the device size.

Operations with stateless ESXi hosts might not pick the expected remote disk for system cache, which causes remediation or compliance issues

Operations with stateless ESXi hosts, such as storage migration, might not pick the expected remote disk for system cache. For example, you want to keep the new boot LUN as LUN 0, but vSphere Auto Deploy picks LUN 1.

This issue is resolved in this release. The fix provides a consistent way to sort the remote disks and always pick the disk with the lowest LUN ID. To make sure you enable the fix, follow these steps:

1. On the Edit host profile page of the Auto Deploy wizard, select **Advanced Configuration Settings > System Image Cache Configuration**
2. In the **System Image Cache Profile Settings** drop-down menu, select **Enable stateless caching on the host**.
3. Edit **Arguments for first disk** by replacing **remote** with **sortedremote** and/or **remoteesx** with **sortedremoteesx**.

VMware VIB installation might fail during concurrent vendor package installations

When you install update packages from several vendors, such as JetStream Software, Microsoft, and VMware, multiple clients call the same PatchManager APIs and might lead to a race condition. As a result, VMware installation packages (VIBs) might fail to install. In the logs, you see an error such as `vim.fault.PlatformConfigFault`, which is a catch-all fault indicating that some error has occurred regarding the configuration of the ESXi host. In the vSphere Client, you see a message such as `An error occurred during host configuration.`

This issue is resolved in this release. The fix is to return a `TaskInProgress` warning instead of `PlatformConfigFault`, so that you are aware of the actual issue and retry the installation.

Certain applications might take too many ESXi file handles and cause performance aggravation

In very rare cases, applications such as NVIDIA virtual GPU (vGPU) might consume so many file handles that ESXi fails to process other services or VMs. As a result, you might see GPU on some nodes to disappear, or report zero GPU memory, or performance degradation.

This issue is resolved in this release. The fix reduces the number of file handles a vGPU VM can consume.

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

This issue is resolved in this release.

esx-update_8.0.1-0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_loadesx_8.0.1-0.25.22088125 - VMware_bootbank_esx-update_8.0.1-0.25.22088125
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesx` and `esx-update` VIBs.

esxio-update_8.0.1-0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_loadesxio_8.0.1-0.25.22088125 - VMware_bootbank_esxio-update_8.0.1-0.25.22088125
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

Mellanox-nmlx5_4.23.0.36-15vmw.801.0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMW_bootbank_nmlx5-rdma_4.23.0.36-15vmw.801.0.25.22088125 - VMW_bootbank_nmlx5-core_4.23.0.36-15vmw.801.0.25.22088125 - VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-15vmw.801.0.25.22088125 - VMW_bootbank_nmlx5-core-esxio_4.23.0.36-15vmw.801.0.25.22088125

PRs Fixed	3152476
CVE numbers	N/A

Updates the `nmlx5-rdma`, `nmlx5-core`, `nmlx5-rdma-esxio`, and `nmlx5-core-esxio` VIBs to resolve the following issue:

In some cases, you might see low throughput in encapsulated traffic with Mellanox NICs

In non Enhanced Datapath mode, the throughput of encapsulated traffic of NICs with the Mellanox (`nmlx5`) driver might be low and you might see uneven traffic flow across RSS queues.

This issue is resolved in this release.

Broadcom-ntg3_4.1.10.0-5vmw.801.0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	VMW_bootbank_ntg3_4.1.10.0-5vmw.801.0.25.22088125
PRs Fixed	3187446
CVE numbers	N/A

Updates the `ntg3` VIB to resolve the following issue:

After upgrading the ntg3 driver to version 4.1.9.0-4vmw, Broadcom NICs with fiber physical connectivity might lose network

Changes in the `ntg3` driver version `4.1.9.0-4vmw` might cause link issues for the fiber physical layer and connectivity on some NICs, such as Broadcom 1Gb, fails to come up.

This issue is resolved in this release.

VMware-NVMeoF-TCP_1.0.1.7-1vmw.801.0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	VMW_bootbank_nvmetcp_1.0.1.7-1vmw.801.0.25.22088125
PRs Fixed	3219191
CVE numbers	N/A

Updates the `nvmetcp` VIB to resolve the following issue:

Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations

vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

This issue is resolved in this release.

ESXi_8.0.1-0.20.22082334

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_bmcal-esxio_8.0.1-0.20.22082334 - VMware_bootbank_native-misc-drivers_8.0.1-0.20.22082334 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.20.22082334 - VMware_bootbank_vsan_8.0.1-0.20.22082334 - VMware_bootbank_bmcal_8.0.1-0.20.22082334 - VMware_bootbank_esx-base_8.0.1-0.20.22082334 - VMware_bootbank_trx_8.0.1-0.20.22082334 - VMware_bootbank_gc-esxio_8.0.1-0.20.22082334 - VMware_bootbank_gc_8.0.1-0.20.22082334 - VMware_bootbank_esxio_8.0.1-0.20.22082334 - VMware_bootbank_esx-xserver_8.0.1-0.20.22082334 - VMware_bootbank_cpu-microcode_8.0.1-0.20.22082334 - VMware_bootbank_esxio-combiner-esxio_8.0.1-0.20.22082334 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.20.22082334 - VMware_bootbank_vsanhealth_8.0.1-0.20.22082334 - VMware_bootbank_vds-vsip_8.0.1-0.20.22082334 - VMware_bootbank_clusterstore_8.0.1-0.20.22082334 - VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.20.22082334 - VMware_bootbank_esxio-combiner_8.0.1-0.20.22082334 - VMware_bootbank_esxio-base_8.0.1-0.20.22082334 - VMware_bootbank_vdfs_8.0.1-0.20.22082334 - VMware_bootbank_crx_8.0.1-0.20.22082334
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149
CVE numbers	N/A

Updates the bmcal-esxio, native-misc-drivers, esx-dvfilter-generic-fastpath, vsan, bmcal, esx-base, trx, gc-esxio, gc, esxio, esx-xserver, cpu-microcode, esxio-combiner-esxio, esxio-dvfilter-generic-fastpath, vsanhealth, vds-vsip, clusterstore, native-misc-drivers-esxio, esxio-combiner, esxio-base, vdfs, and crx VIBs to resolve the following issues:

ESXi 8.0 Update 1c provides the following security updates:

- The Envoy proxy is updated to version v1.23.9.
- The ESXi userworld libxml2 library is updated to version 2.10.4.
- The cURL library is updated to version 8.0.1.
- The Go library is updated to version 1.19.9.
- The etcd package is updated to 3.4.25.

esx-update_8.0.1-0.20.22082334

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	• VMware_bootbank_esx-update_8.0.1-0.20.22082334 • VMware_bootbank_loadesx_8.0.1-0.20.22082334
PRs Fixed	N/A
CVE numbers	N/A

Updates the `esx-update` and `loadesx` VIBs.

esxio-update_8.0.1-0.20.22082334

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	• VMware_bootbank_loadesxio_8.0.1-0.20.22082334 • VMware_bootbank_esxio-update_8.0.1-0.20.22082334
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

VMware-VM-Tools_12.2.5.21855600-22082334

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	No
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A
Affected Software	N/A

VIBs Included	VMware_locker_tools-light_12.2.5.21855600-22082334
PRs Fixed	3186166
CVE numbers	N/A

Updates the `tools-light` VIB.

- The following VMware Tools ISO images are bundled with ESXi 8.0 Update 1c:
 - windows.iso**: VMware Tools 12.2.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.
- The following VMware Tools ISO images are available for download:
- VMware Tools 11.0.6:
 - windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
 - VMware Tools 10.0.12:
 - winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
 - solaris.iso**: VMware Tools image 10.3.10 for Solaris.
 - darwin.iso**: Supports Mac OS X versions 10.11 and later. VMware Tools 12.1.0 was the last regular release for macOS. Refer VMware knowledge base article [88698](#) for details.
- Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:
- [Updating VMware Tools](#)
 - [VMware Tools for hosts provisioned with Auto Deploy](#)
 - [What Every vSphere Admin Must Know About VMware Tools](#)
 - [Earlier versions of VMware Tools](#)
 - [VMware Tools 12.2.5 Release Notes](#)

ESXi-8.0U1c-22088125-standard

Profile Name	ESXi-8.0U1c-22088125-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_vdfs_8.0.1-0.25.22088125 - VMW_bootbank_nvmetcp_1.0.1.7-1vmw.801.0.25.22088125 - VMW_bootbank_ntg3_4.1.10.0-5vmw.801.0.25.22088125 - VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-15vmw.801.0.25.22088125 - VMware_bootbank_gc_8.0.1-0.25.22088125 - VMware_bootbank_native-misc-drivers_8.0.1-0.25.22088125 - VMware_bootbank_cpu-microcode_8.0.1-0.25.22088125 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.25.22088125 - VMware_bootbank_vds-vsip_8.0.1-0.25.22088125

	• VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.25.22088125 • VMware_bootbank_loadesxio_8.0.1-0.25.22088125 • VMware_bootbank_loadesx_8.0.1-0.25.22088125 • VMware_bootbank_gc-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esxio-combiner_8.0.1-0.25.22088125 • VMW_bootbank_nmlx5-core_4.23.0.36-15vmw.801.0.25.22088125 • VMware_bootbank_esxio-update_8.0.1-0.25.22088125 • VMware_bootbank_esx-update_8.0.1-0.25.22088125 • VMware_bootbank_vsan_8.0.1-0.25.22088125 • VMware_bootbank_crx_8.0.1-0.25.22088125 • VMW_bootbank_nmlx5-rdma_4.23.0.36-15vmw.801.0.25.22088125 • VMware_bootbank_bmcad_8.0.1-0.25.22088125 • VMware_bootbank_esx-xserver_8.0.1-0.25.22088125 • VMware_bootbank_trx_8.0.1-0.25.22088125 • VMware_bootbank_esxio-combiner-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esxio-base_8.0.1-0.25.22088125 • VMware_bootbank_clusterstore_8.0.1-0.25.22088125 • VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esxio_8.0.1-0.25.22088125 • VMware_bootbank_bmcad-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esx-base_8.0.1-0.25.22088125 • VMware_bootbank_vsanhealth_8.0.1-0.25.22088125 • VMW_bootbank_nmlx5-core-esxio_4.23.0.36-15vmw.801.0.25.22088125 • VMware_locker_tools-light_12.2.5.21855600-22082334
PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236, 3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3038908, 3179111, 3152476, 3187446, 3219191
Related CVE numbers	N/A

This patch updates the following issues:

Setting the screen resolution of a powered off, encrypted virtual machine does not always work when manually editing the VMX file

If you manually specify the screen resolution of a powered off, encrypted virtual machine by editing the VMX file, the change might not take effect.

This issue is resolved in this release.

Transient vSAN health check warning: Network configuration is out of sync

vSAN Skyline health might randomly report that network configuration is out of sync. This transient issue occurs when the vSAN health service uses an outdated vCenter configuration to perform unicast check.

This issue is resolved in this release.

vSAN cache overwrite issue might cause inconsistent metadata

In rare cases, racing between unmap and snapshot operations might cause inconsistent vSAN metadata in a vSAN host. Such inconsistency can lead to unpredictable consequences, including VM failures or host failures.

This issue is resolved in this release.

ESX hosts might fail with a purple diagnostic screen and an error NMI IPI: Panic requested by another PCPU

The resource pool cache is a VMFS specific volume level cache that stores the resource clusters corresponding to the VMFS volume. While searching for priority clusters, the cache flusher workflow iterates through a large list of cached resource clusters, which can cause lockup of the physical CPUs. As a result, ESX hosts might fail with a purple diagnostic screen. In the `logDump` file, you see an error such as:

```
^[[7m2022-10-22T07:56:47.322Z cpu13:2101160)WARNING: Heartbeat: 827: PCPU 0 didn't have a
heartbeat for 7 seconds, timeout is 14, 1 IPIs sent; *may* be locked up.^[[0m
^[[31;1m2022-10-22T07:56:47.322Z cpu0:2110633)ALERT: NMI: 710: NMI IPI:
RIPOFF(base):RBP:CS
```

This issue is resolved in this release.

ESXi ConfigStore database fills up and writes fail

Stale data related to block devices might not be deleted in time from the ESXi ConfigStore database and cause an out of space condition. As a result, write operations to ConfigStore start to fail. In the backtrace, you see logs such as:

```
2022-12-19T03:51:42.733Z cpu53:26745174)WARNING: VisorFSRam: 203: Cannot extend visorfs
file /etc/vmware/configstore/current-store-1-journal because its ramdisk (configstore) is
full.
```

This issue is resolved in this release.

Delay before evacuating vSAN storage device in UNHEALTHY state

After detecting devices as UNHEALTHY, the Local Log-Structured Object Manager (LSOM) might wait for 10 minutes (LSOM_DEVICE_MONITORING_INTERVAL) before initiating evacuation on these devices.

This issue is resolved in this release.

vSAN precheck for maintenance mode or disk decommission doesn't list objects that might lose accessibility

This issue affects objects with resyncing components, and some components reside on a device to be removed or placed into maintenance mode. When you run a precheck with the **No-Action** option, the precheck does not evaluate the object correctly to report it in the `inaccessibleObjects` list.

This issue is resolved in this release. Precheck includes all affected objects in the `inaccessibleObjects` list.

The vSAN iSCSI Target service might fail due to a rare race condition

When you run the ESXCLI command `esxcli network firewall load`, the operation does not reload existing dynamic firewall rules and they are lost. When you run the ESXCLI command `esxcli network firewall refresh`, the operation reloads existing dynamic firewall rules, but in some cases a race condition might cause some rules to

be lost. The issue occurs only when multiple firewall refresh commands run at the same time, which leads to the race condition. As a result, the vSAN iSCSI Target service might fail.

This issue is resolved in this release.

If the Internet Control Message Protocol (ICMPA) is not active, ESXi host reboot might take long after upgrading to vSphere 8.0 and later

If ICMPA is not active on the NFS servers in your environment, after upgrading your system to vSphere 8.0 and later, ESXi hosts reboot might take an hour to complete, because restore operations for NFS datastores fail. NFS uses the `vmkping` utility to identify reachable IPs of the NFS servers before executing a mount operation and when ICMP is not active, mount operations fail.

This issue is resolved in this release. To remove dependency on the ICMP protocol to find reachable IPs, the fix adds socket APIs to ensure that IPs on a given NFS server are available.

When you migrate a VM with recently hot-added memory, an ESXi host might repeatedly fail with a purple diagnostic screen

Due to a race condition while the memory hotplug module recomputes the NUMA memory layout of a VM on a destination host after migration, an ESXi host might repeatedly fail with a purple diagnostic screen. In the backtrace, you see errors such as:

```
0x452900262cf0:[0x4200138fee8b]PanicvPanicInt@vmkernel#nover+0x327 stack: 0x452900262dc8,
0x4302f6c06508, 0x4200138fee8b, 0x420013df1300, 0x452900262cf0 0x452900262dc0:
[0x4200138ff43d]Panic_WithBacktrace@vmkernel#nover+0x56 stack: 0x452900262e30,
0x452900262de0, 0x452900262e40, 0x452900262df0, 0x3e7514 0x452900262e30:
[0x4200138fbb90]NMI_Interrupt@vmkernel#nover+0x561 stack: 0x0, 0xf48, 0x0,
0x0, 0x0 0x452900262f00:[0x420013953392]IDTNMIWork@vmkernel#nover+0x7f stack:
0x420049800000, 0x4200139546dd, 0x0, 0x452900262fd0, 0x0 0x452900262f20:
[0x4200139546dc]Int2_NMI@vmkernel#nover+0x19 stack: 0x0, 0x42001394e068,
0xf50, 0xf50, 0x0 0x452900262f40:[0x42001394e067]gate_entry@vmkernel#nover
+0x68 stack: 0x0, 0x43207bc02088, 0xd, 0x0, 0x43207bc02088 0x45397b61bd30:
[0x420013be7514]NUMASched_PageNum2PhysicalDomain@vmkernel#nover+0x58 stack:
0x1, 0x420013be34c3, 0x45396f79f000, 0x1, 0x100005cf757 0x45397b61bd50:
[0x420013be34c2]NUMASched_UpdateAllocStats@vmkernel#nover+0x4b stack: 0x100005cf757, 0x0,
0x0, 0x4200139b36d9, 0x0 0x45397b61bd80:[0x4200139b36d8]VmMem_NodeStatsSub@vmkernel#nover
+0x59 stack: 0x39, 0x45396f79f000, 0xbce0dbf, 0x100005cf757, 0x0 0x45397b61bdc0:
[0x4200139b4372]VmMem_FreePageNoBackmap@vmkernel#nover+0x8b stack: 0x465ec0001be0, 0xa,
0x465ec18748b0, 0x420014e7685f, 0x465ec14437d0
```

This issue is resolved in this release.

After upgrading the ntg3 driver to version 4.1.9.0-4vmw, Broadcom NICs with fiber physical connectivity might lose network

Changes in the ntg3 driver version 4.1.9.0-4vmw might cause link issues for the fiber physical layer and connectivity on some NICs, such as Broadcom 1Gb, fails to come up.

This issue is resolved in this release.

vSAN cluster shutdown fails on a cluster with IPV6 disabled

This issue occurs on vSAN hosts running ESXi 8.0 Update 1 with IPv6 disabled. When you use the vSAN cluster shutdown wizard, the workflow fails with the following error message: 'NoneType' object is not iterable.

This issue is resolved in this release.

You see trap files in a SNMP directory under /var/spool even though SNMP is not enabled

After the `hostd` service starts, for example after an ESXi host reboot, it might create a SNMP directory under `/var/spool` and you see many `.trp` files to pile up in this directory.

This issue is resolved in this release. The fix makes sure that the directory `/var/spool/snmp` exists only when SNMP is enabled.

The `hostd` service repeatedly fails and the ESXi host disconnects from the vCenter system

If for any reason an ESXi host is temporarily in a state of insufficient memory, the `hostd` service might repeatedly fail due to a vSphere Replication filter that prevents the allocation of bitmaps. As a result, the ESXi host disconnects from the vCenter system and cannot connect back.

This issue is resolved in this release.

Performance of certain nested virtual machines on AMD CPUs might degrade

Nested virtual machines on AMD CPUs with operational systems such as Windows with virtualization-based security (VBS) might experience performance degradation, timeouts, or unresponsiveness due to an issue with the virtualization of AMD's Rapid Virtualization Indexing (RVI), also known as Nested Page Tables (NPT).

This issue is resolved in this release.

Changing the mode of the virtual disk on a running virtual machine might cause the VM to fail

If you use the VMware Host Client to edit the disk mode of a running virtual machine, for example from Independent - Nonpersistent to Dependent or Independent - Persistent, the operation fails and might cause the VM to fail. In the `vmware.log`, you see errors such as:

```
msg.disk.notConfigured2] Failed to configure disk 'scsi0:4'. The virtual machine cannot be
powered on with an unconfigured disk.
```

```
[msg.checkpoint.continuesync.error] An operation required the virtual machine to quiesce
and the virtual machine was unable to continue running.
```

This issue is resolved in this release. The fix blocks changing the mode of an Independent - Nonpersistent disk on a running virtual machine by using the VMware Host Client. The vSphere Client already blocks such operations.

ESXi NVMe/TCP initiator fails to recover paths after target failure recovery

When an NVMe/TCP target recovers from a failure, ESXi cannot recover the path.

This issue is resolved in this release.

ESXi host becomes unresponsive and you cannot put the host in Maintenance Mode or migrate VMs from that host

Asynchronous reads of metadata on a VMFS volume attached to an ESXi host might cause a race condition with other threads on the host and make the host unresponsive. As a result, you cannot put the host in Maintenance Mode or migrate VMs from that host.

This issue is resolved in this release.

In some cases, you might see low throughput in encapsulated traffic with Mellanox NICs

In non Enhanced Datapath mode, the throughput of encapsulated traffic of NICs with the Mellanox (nmlx5) driver might be low and you might see uneven traffic flow across RSS queues.

This issue is resolved in this release.

A Logical Volume Manager (LVM) disk goes offline during datastore expansion

If during the expansion of a datastore on an ESXi host in a cluster you run a storage refresh, a LVM extent might go offline and virtual machines on this volume become unresponsive. The issue occurs because the storage refresh operation triggers a slow refresh of the volume attributes on all the ESXi hosts in the cluster. As a result, LVM metadata on the disk

might not match the cached capacity information in the Pluggable Storage Architecture (PSA) layer and ESXi marks the LVM extent offline for the safety of metadata and data.

This issue is resolved in this release.

The durable name of a SCSI LUN might not be set

The durable name property for a SCSI-3 compliant device comes from pages 80h and 83h of the Vital Product Data (VPD) as defined by the T10 and SMI standards. To populate the durable name, ESXi first sends an inquiry command to get a list of VPD pages supported by the device. Then ESXi issues commands to get data for all supported VPD pages. Due to an issue with the target array, the device might fail a command to get VPD page data for a page in the list with a `not supported` error. As a result, ESXi cannot populate the durable name property for the device.

This issue is resolved in this release. The fix ignores the error on command to get VPD page data, except pages 80h and 83h, if that data is not required for the generation of durable name.

Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations

vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

This issue is resolved in this release.

Excessive time for path recovery with ESXi High Performance Plug-in (HPP) after a link up event for NVMe over Fibre Channel

In certain scenarios, after a Fibre Channel link up event, it might take as long as 5 minutes for HPP managed NVMe over Fibre Channel paths to recover.

This issue is resolved in this release.

NVMe over Fabrics controller might unexpectedly disconnect during discovery

If discovery and I/O controllers already exist for a NVMe over Fabrics storage target that supports a persistent discovery controller on an ESXi host, a concurrent NVMe over Fabrics controller discovery operation might cause some I/O controllers to disconnect unexpectedly.

This issue is resolved in this release.

If your ESXi PTP service uses hardware timestamping, enabling IGMP snooping on switches might cause synchronization failures

When you enable Internet Group Management Protocol (IGMP) snooping on a connected switch, the Precision Time Protocol (PTP) client needs to send IGMP multicast requests to receive PTP multicast stream from the grandmaster. If the ESXi PTP agent is based on hardware timestamping, the agent might fail to send IGMP join/leave requests to the switch. As a result, the PTP multicast stream cannot go forward to the ESXi host and prevents proper PTP synchronization.

This issue is resolved in this release. For more details, see VMware knowledge base article [92276](#).

If parallel volume expand and volume refresh operations on the same VMFS volume run on two ESXi hosts in the same cluster, the VMFS volume might go offline

While a VMFS volume expand operation is in progress on an ESXi host in a vCenter cluster, if on another host a user or vCenter initiates a refresh of the same VMFS volume capacity, such a volume might go offline. The issue occurs due to a possible mismatch in the device size, which is stamped on the disk in the volume metadata during a device rescan, and the device size value in the Pluggable Storage Architecture (PSA) layer on the host, which might not be updated if the device rescan is not complete.

This issue is resolved in this release. The fix improves the resiliency of the volume manager code to force a consecutive refresh of the device attributes and comparison of the device sizes again if vCenter reports a mismatch in the device size.

Operations with stateless ESXi hosts might not pick the expected remote disk for system cache, which causes remediation or compliance issues

Operations with stateless ESXi hosts, such as storage migration, might not pick the expected remote disk for system cache. For example, you want to keep the new boot LUN as LUN 0, but vSphere Auto Deploy picks LUN 1.

This issue is resolved in this release. The fix provides a consistent way to sort the remote disks and always pick the disk with the lowest LUN ID. To make sure you enable the fix, follow these steps:

1. On the Edit host profile page of the Auto Deploy wizard, select **Advanced Configuration Settings > System Image Cache Configuration**
2. In the **System Image Cache Profile Settings** drop-down menu, select **Enable stateless caching on the host**.
3. Edit **Arguments for first disk** by replacing **remote** with **sortedremote** and/or **remoteesx** with **sortedremoteesx**.

VMware VIB installation might fail during concurrent vendor package installations

When you install update packages from several vendors, such as JetStream Software, Microsoft, and VMware, multiple clients call the same PatchManager APIs and might lead to a race condition. As a result, VMware installation packages (VIBs) might fail to install. In the logs, you see an error such as `vim.fault.PlatformConfigFault`, which is a catch-all fault indicating that some error has occurred regarding the configuration of the ESXi host. In the vSphere Client, you see a message such as `An error occurred during host configuration.`

This issue is resolved in this release. The fix is to return a `TaskInProgress` warning instead of `PlatformConfigFault`, so that you are aware of the actual issue and retry the installation.

Certain applications might take too many ESXi file handles and cause performance aggravation

In very rare cases, applications such as NVIDIA virtual GPU (vGPU) might consume so many file handles that ESXi fails to process other services or VMs. As a result, you might see GPU on some nodes to disappear, or report zero GPU memory, or performance degradation.

This issue is resolved in this release. The fix reduces the number of file handles a vGPU VM can consume.

ESXi-8.0U1c-22088125-no-tools

Profile Name	ESXi-8.0U1c-22088125-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_vdfs_8.0.1-0.25.22088125 - VMW_bootbank_nvmetcp_1.0.1.7-1vmw.801.0.25.22088125 - VMW_bootbank_ntg3_4.1.10.0-5vmw.801.0.25.22088125 - VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-15vmw.801.0.25.22088125 - VMware_bootbank_gc_8.0.1-0.25.22088125 - VMware_bootbank_native-misc-drivers_8.0.1-0.25.22088125 - VMware_bootbank_cpu-microcode_8.0.1-0.25.22088125 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.25.22088125

	• VMware_bootbank_vds-vsip_8.0.1-0.25.22088125 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.25.22088125 • VMware_bootbank_loadesxio_8.0.1-0.25.22088125 • VMware_bootbank_loadesx_8.0.1-0.25.22088125 • VMware_bootbank_gc-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esxio-combiner_8.0.1-0.25.22088125 • VMW_bootbank_nmlx5-core_4.23.0.36-15vmw.801.0.25.22088125 • VMware_bootbank_esxio-update_8.0.1-0.25.22088125 • VMware_bootbank_esx-update_8.0.1-0.25.22088125 • VMware_bootbank_vsan_8.0.1-0.25.22088125 • VMware_bootbank_crx_8.0.1-0.25.22088125 • VMW_bootbank_nmlx5-rdma_4.23.0.36-15vmw.801.0.25.22088125 • VMware_bootbank_bmcad_8.0.1-0.25.22088125 • VMware_bootbank_esx-xserver_8.0.1-0.25.22088125 • VMware_bootbank_trx_8.0.1-0.25.22088125 • VMware_bootbank_esxio-combiner-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esxio-base_8.0.1-0.25.22088125 • VMware_bootbank_clusterstore_8.0.1-0.25.22088125 • VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esxio_8.0.1-0.25.22088125 • VMware_bootbank_bmcad-esxio_8.0.1-0.25.22088125 • VMware_bootbank_esx-base_8.0.1-0.25.22088125 • VMware_bootbank_vsanhealth_8.0.1-0.25.22088125 • VMW_bootbank_nmlx5-core-esxio_4.23.0.36-15vmw.801.0.25.22088125
PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236, 3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3179111, 3152476, 3187446, 3219191
Related CVE numbers	N/A

This patch updates the following issues:

Setting the screen resolution of a powered off, encrypted virtual machine does not always work when manually editing the VMX file

If you manually specify the screen resolution of a powered off, encrypted virtual machine by editing the VMX file, the change might not take effect.

This issue is resolved in this release.

Transient vSAN health check warning: Network configuration is out of sync

vSAN Skyline health might randomly report that network configuration is out of sync. This transient issue occurs when the vSAN health service uses an outdated vCenter configuration to perform unicast check.

This issue is resolved in this release.

vSAN cache overwrite issue might cause inconsistent metadata

In rare cases, racing between unmap and snapshot operations might cause inconsistent vSAN metadata in a vSAN host. Such inconsistency can lead to unpredictable consequences, including VM failures or host failures.

This issue is resolved in this release.

ESX hosts might fail with a purple diagnostic screen and an error NMI IPI: Panic requested by another PCPU

The resource pool cache is a VMFS specific volume level cache that stores the resource clusters corresponding to the VMFS volume. While searching for priority clusters, the cache flusher workflow iterates through a large list of cached resource clusters, which can cause lockup of the physical CPUs. As a result, ESX hosts might fail with a purple diagnostic screen. In the `logDump` file, you see an error such as:

```
^[[7m2022-10-22T07:56:47.322Z cpu13:2101160)WARNING: Heartbeat: 827: PCPU 0 didn't have a
heartbeat for 7 seconds, timeout is 14, 1 IPIs sent; *may* be locked up.^[[0m
^[[31;1m2022-10-22T07:56:47.322Z cpu0:2110633)ALERT: NMI: 710: NMI IPI:
RIP0FF(base):RBP:CS
```

This issue is resolved in this release.

ESXi ConfigStore database fills up and writes fail

Stale data related to block devices might not be deleted in time from the ESXi ConfigStore database and cause an out of space condition. As a result, write operations to ConfigStore start to fail. In the backtrace, you see logs such as:

```
2022-12-19T03:51:42.733Z cpu53:26745174)WARNING: VisorFSRam: 203: Cannot extend visorfs
file /etc/vmware/configstore/current-store-1-journal because its ramdisk (configstore) is
full.
```

This issue is resolved in this release.

Delay before evacuating vSAN storage device in UNHEALTHY state

After detecting devices as UNHEALTHY, the Local Log-Structured Object Manager (LSOM) might wait for 10 minutes (LSOM_DEVICE_MONITORING_INTERVAL) before initiating evacuation on these devices.

This issue is resolved in this release.

vSAN precheck for maintenance mode or disk decommission doesn't list objects that might lose accessibility

This issue affects objects with resyncing components, and some components reside on a device to be removed or placed into maintenance mode. When you run a precheck with the **No-Action** option, the precheck does not evaluate the object correctly to report it in the `inaccessibleObjects` list.

This issue is resolved in this release. Precheck includes all affected objects in the `inaccessibleObjects` list.

The vSAN iSCSI Target service might fail due to a rare race condition

When you run the ESXCLI command `esxcli network firewall load`, the operation does not reload existing dynamic firewall rules and they are lost. When you run the ESXCLI command `esxcli network firewall refresh`, the operation reloads existing dynamic firewall rules, but in some cases a race condition might cause some rules to

be lost. The issue occurs only when multiple firewall refresh commands run at the same time, which leads to the race condition. As a result, the vSAN iSCSI Target service might fail.

This issue is resolved in this release.

If the Internet Control Message Protocol (ICMPA) is not active, ESXi host reboot might take long after upgrading to vSphere 8.0 and later

If ICMPA is not active on the NFS servers in your environment, after upgrading your system to vSphere 8.0 and later, ESXi hosts reboot might take an hour to complete, because restore operations for NFS datastores fail. NFS uses the `vmkping` utility to identify reachable IPs of the NFS servers before executing a mount operation and when ICMP is not active, mount operations fail.

This issue is resolved in this release. To remove dependency on the ICMP protocol to find reachable IPs, the fix adds socket APIs to ensure that IPs on a given NFS server are available.

When you migrate a VM with recently hot-added memory, an ESXi host might repeatedly fail with a purple diagnostic screen

Due to a race condition while the memory hotplug module recomputes the NUMA memory layout of a VM on a destination host after migration, an ESXi host might repeatedly fail with a purple diagnostic screen. In the backtrace, you see errors such as:

```
0x452900262cf0:[0x4200138fee8b]PanicvPanicInt@vmkernel#nover+0x327 stack: 0x452900262dc8,
0x4302f6c06508, 0x4200138fee8b, 0x420013df1300, 0x452900262cf0 0x452900262dc0:
[0x4200138ff43d]Panic_WithBacktrace@vmkernel#nover+0x56 stack: 0x452900262e30,
0x452900262de0, 0x452900262e40, 0x452900262df0, 0x3e7514 0x452900262e30:
[0x4200138fbb90]NMI_Interrupt@vmkernel#nover+0x561 stack: 0x0, 0xf48, 0x0,
0x0, 0x0 0x452900262f00:[0x420013953392]IDTNMIWork@vmkernel#nover+0x7f stack:
0x420049800000, 0x4200139546dd, 0x0, 0x452900262fd0, 0x0 0x452900262f20:
[0x4200139546dc]Int2_NMI@vmkernel#nover+0x19 stack: 0x0, 0x42001394e068,
0xf50, 0xf50, 0x0 0x452900262f40:[0x42001394e067]gate_entry@vmkernel#nover
+0x68 stack: 0x0, 0x43207bc02088, 0xd, 0x0, 0x43207bc02088 0x45397b61bd30:
[0x420013be7514]NUMASched_PageNum2PhysicalDomain@vmkernel#nover+0x58 stack:
0x1, 0x420013be34c3, 0x45396f79f000, 0x1, 0x100005cf757 0x45397b61bd50:
[0x420013be34c2]NUMASched_UpdateAllocStats@vmkernel#nover+0x4b stack: 0x100005cf757, 0x0,
0x0, 0x4200139b36d9, 0x0 0x45397b61bd80:[0x4200139b36d8]VmMem_NodeStatsSub@vmkernel#nover
+0x59 stack: 0x39, 0x45396f79f000, 0xbce0dbf, 0x100005cf757, 0x0 0x45397b61bdc0:
[0x4200139b4372]VmMem_FreePageNoBackmap@vmkernel#nover+0x8b stack: 0x465ec0001be0, 0xa,
0x465ec18748b0, 0x420014e7685f, 0x465ec14437d0
```

This issue is resolved in this release.

After upgrading the ntg3 driver to version 4.1.9.0-4vmw, Broadcom NICs with fiber physical connectivity might lose network

Changes in the ntg3 driver version 4.1.9.0-4vmw might cause link issues for the fiber physical layer and connectivity on some NICs, such as Broadcom 1Gb, fails to come up.

This issue is resolved in this release.

vSAN cluster shutdown fails on a cluster with IPV6 disabled

This issue occurs on vSAN hosts running ESXi 8.0 Update 1 with IPv6 disabled. When you use the vSAN cluster shutdown wizard, the workflow fails with the following error message: 'NoneType' object is not iterable.

This issue is resolved in this release.

You see trap files in a SNMP directory under /var/spool even though SNMP is not enabled

After the `hostd` service starts, for example after an ESXi host reboot, it might create a SNMP directory under `/var/spool` and you see many `.trp` files to pile up in this directory.

This issue is resolved in this release. The fix makes sure that the directory `/var/spool/snmp` exists only when SNMP is enabled.

The `hostd` service repeatedly fails and the ESXi host disconnects from the vCenter system

If for any reason an ESXi host is temporarily in a state of insufficient memory, the `hostd` service might repeatedly fail due to a vSphere Replication filter that prevents the allocation of bitmaps. As a result, the ESXi host disconnects from the vCenter system and cannot connect back.

This issue is resolved in this release.

Performance of certain nested virtual machines on AMD CPUs might degrade

Nested virtual machines on AMD CPUs with operational systems such as Windows with virtualization-based security (VBS) might experience performance degradation, timeouts, or unresponsiveness due to an issue with the virtualization of AMD's Rapid Virtualization Indexing (RVI), also known as Nested Page Tables (NPT).

This issue is resolved in this release.

Changing the mode of the virtual disk on a running virtual machine might cause the VM to fail

If you use the VMware Host Client to edit the disk mode of a running virtual machine, for example from Independent - Nonpersistent to Dependent or Independent - Persistent, the operation fails and might cause the VM to fail. In the `vmware.log`, you see errors such as:

```
msg.disk.notConfigured2] Failed to configure disk 'scsi0:4'. The virtual machine cannot be
powered on with an unconfigured disk.
```

```
[msg.checkpoint.continuesync.error] An operation required the virtual machine to quiesce
and the virtual machine was unable to continue running.
```

This issue is resolved in this release. The fix blocks changing the mode of an Independent - Nonpersistent disk on a running virtual machine by using the VMware Host Client. The vSphere Client already blocks such operations.

ESXi NVMe/TCP initiator fails to recover paths after target failure recovery

When an NVMe/TCP target recovers from a failure, ESXi cannot recover the path.

This issue is resolved in this release.

ESXi host becomes unresponsive and you cannot put the host in Maintenance Mode or migrate VMs from that host

Asynchronous reads of metadata on a VMFS volume attached to an ESXi host might cause a race condition with other threads on the host and make the host unresponsive. As a result, you cannot put the host in Maintenance Mode or migrate VMs from that host.

This issue is resolved in this release.

In some cases, you might see low throughput in encapsulated traffic with Mellanox NICs

In non Enhanced Datapath mode, the throughput of encapsulated traffic of NICs with the Mellanox (nmlx5) driver might be low and you might see uneven traffic flow across RSS queues.

This issue is resolved in this release.

A Logical Volume Manager (LVM) disk goes offline during datastore expansion

If during the expansion of a datastore on an ESXi host in a cluster you run a storage refresh, a LVM extent might go offline and virtual machines on this volume become unresponsive. The issue occurs because the storage refresh operation triggers a slow refresh of the volume attributes on all the ESXi hosts in the cluster. As a result, LVM metadata on the disk

might not match the cached capacity information in the Pluggable Storage Architecture (PSA) layer and ESXi marks the LVM extent offline for the safety of metadata and data.

This issue is resolved in this release.

The durable name of a SCSI LUN might not be set

The durable name property for a SCSI-3 compliant device comes from pages 80h and 83h of the Vital Product Data (VPD) as defined by the T10 and SMI standards. To populate the durable name, ESXi first sends an inquiry command to get a list of VPD pages supported by the device. Then ESXi issues commands to get data for all supported VPD pages. Due to an issue with the target array, the device might fail a command to get VPD page data for a page in the list with a `not supported` error. As a result, ESXi cannot populate the durable name property for the device.

This issue is resolved in this release. The fix ignores the error on command to get VPD page data, except pages 80h and 83h, if that data is not required for the generation of durable name.

Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations

vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

This issue is resolved in this release.

Excessive time for path recovery with ESXi High Performance Plug-in (HPP) after a link up event for NVMe over Fibre Channel

In certain scenarios, after a Fibre Channel link up event, it might take as long as 5 minutes for HPP managed NVMe over Fibre Channel paths to recover.

This issue is resolved in this release.

NVMe over Fabrics controller might unexpectedly disconnect during discovery

If discovery and I/O controllers already exist for a NVMe over Fabrics storage target that supports a persistent discovery controller on an ESXi host, a concurrent NVMe over Fabrics controller discovery operation might cause some I/O controllers to disconnect unexpectedly.

This issue is resolved in this release.

If your ESXi PTP service uses hardware timestamping, enabling IGMP snooping on switches might cause synchronization failures

When you enable Internet Group Management Protocol (IGMP) snooping on a connected switch, the Precision Time Protocol (PTP) client needs to send IGMP multicast requests to receive PTP multicast stream from the grandmaster. If the ESXi PTP agent is based on hardware timestamping, the agent might fail to send IGMP join/leave requests to the switch. As a result, the PTP multicast stream cannot go forward to the ESXi host and prevents proper PTP synchronization.

This issue is resolved in this release. For more details, see VMware knowledge base article [92276](#).

If parallel volume expand and volume refresh operations on the same VMFS volume run on two ESXi hosts in the same cluster, the VMFS volume might go offline

While a VMFS volume expand operation is in progress on an ESXi host in a vCenter cluster, if on another host a user or vCenter initiates a refresh of the same VMFS volume capacity, such a volume might go offline. The issue occurs due to a possible mismatch in the device size, which is stamped on the disk in the volume metadata during a device rescan, and the device size value in the Pluggable Storage Architecture (PSA) layer on the host, which might not be updated if the device rescan is not complete.

This issue is resolved in this release. The fix improves the resiliency of the volume manager code to force a consecutive refresh of the device attributes and comparison of the device sizes again if vCenter reports a mismatch in the device size.

Operations with stateless ESXi hosts might not pick the expected remote disk for system cache, which causes remediation or compliance issues

Operations with stateless ESXi hosts, such as storage migration, might not pick the expected remote disk for system cache. For example, you want to keep the new boot LUN as LUN 0, but vSphere Auto Deploy picks LUN 1.

This issue is resolved in this release. The fix provides a consistent way to sort the remote disks and always pick the disk with the lowest LUN ID. To make sure you enable the fix, follow these steps:

1. On the Edit host profile page of the Auto Deploy wizard, select **Advanced Configuration Settings > System Image Cache Configuration**
2. In the **System Image Cache Profile Settings** drop-down menu, select **Enable stateless caching on the host**.
3. Edit **Arguments for first disk** by replacing **remote** with **sortedremote** and/or **remoteesx** with **sortedremoteesx**.

VMware VIB installation might fail during concurrent vendor package installations

When you install update packages from several vendors, such as JetStream Software, Microsoft, and VMware, multiple clients call the same PatchManager APIs and might lead to a race condition. As a result, VMware installation packages (VIBs) might fail to install. In the logs, you see an error such as `vim.fault.PlatformConfigFault`, which is a catch-all fault indicating that some error has occurred regarding the configuration of the ESXi host. In the vSphere Client, you see a message such as `An error occurred during host configuration.`

This issue is resolved in this release. The fix is to return a `TaskInProgress` warning instead of `PlatformConfigFault`, so that you are aware of the actual issue and retry the installation.

Certain applications might take too many ESXi file handles and cause performance aggravation

In very rare cases, applications such as NVIDIA virtual GPU (vGPU) might consume so many file handles that ESXi fails to process other services or VMs. As a result, you might see GPU on some nodes to disappear, or report zero GPU memory, or performance degradation.

This issue is resolved in this release. The fix reduces the number of file handles a vGPU VM can consume.

ESXi-8.0U1sc-22082334-standard

Profile Name	ESXi-8.0U1c-22082334-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_bmc-al-esxio_8.0.1-0.20.22082334 - VMware_bootbank_native-misc-drivers_8.0.1-0.20.22082334 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.20.22082334 - VMware_bootbank_vsan_8.0.1-0.20.22082334 - VMware_bootbank_bmc-al_8.0.1-0.20.22082334 - VMware_bootbank_esx-base_8.0.1-0.20.22082334 - VMware_bootbank_trx_8.0.1-0.20.22082334 - VMware_bootbank_gc-esxio_8.0.1-0.20.22082334 - VMware_bootbank_gc_8.0.1-0.20.22082334

	• VMware_bootbank_esxio_8.0.1-0.20.22082334 • VMware_bootbank_esx-xserver_8.0.1-0.20.22082334 • VMware_bootbank_cpu-microcode_8.0.1-0.20.22082334 • VMware_bootbank_esxio-combiner-esxio_8.0.1-0.20.22082334 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.20.22082334 • VMware_bootbank_vsanhealth_8.0.1-0.20.22082334 • VMware_bootbank_vds-vsip_8.0.1-0.20.22082334 • VMware_bootbank_clusterstore_8.0.1-0.20.22082334 • VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.20.22082334 • VMware_bootbank_esxio-combiner_8.0.1-0.20.22082334 • VMware_bootbank_esxio-base_8.0.1-0.20.22082334 • VMware_bootbank_vdfs_8.0.1-0.20.22082334 • VMware_bootbank_crx_8.0.1-0.20.22082334 • VMware_bootbank_esx-update_8.0.1-0.20.22082334 • VMware_bootbank_loadesx_8.0.1-0.20.22082334 • VMware_bootbank_loadesxio_8.0.1-0.20.22082334 • VMware_bootbank_esxio-update_8.0.1-0.20.22082334 • VMware_locker_tools-light_12.2.5.21855600-22082334
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149, 3186166
Related CVE numbers	N/A

This patch updates the following issues:

ESXi 8.0 Update 1c provides the following security updates:

- The Envoy proxy is updated to version v1.23.9.
- The ESXi userworld libxml2 library is updated to version 2.10.4.
- The cURL library is updated to version 8.0.1.
- The Go library is updated to version 1.19.9.
- The etcd package is updated to 3.4.25.

Updates the `tools-light` VIB.

- The following VMware Tools ISO images are bundled with ESXi 8.0 Update 1c:
 - windows.iso**: VMware Tools 12.2.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- darwin.iso**: Supports Mac OS X versions 10.11 and later. VMware Tools 12.1.0 was the last regular release for macOS. Refer VMware knowledge base article [88698](#) for details.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [Updating VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [Earlier versions of VMware Tools](#)
- [VMware Tools 12.2.5 Release Notes](#)

ESXi-8.0U1sc-22082334-no-tools

Profile Name	ESXi-8.0U1c-22082334-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_bmc-al-esxio_8.0.1-0.20.22082334 - VMware_bootbank_native-misc-drivers_8.0.1-0.20.22082334 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.20.22082334 - VMware_bootbank_vsan_8.0.1-0.20.22082334 - VMware_bootbank_bmc-al_8.0.1-0.20.22082334 - VMware_bootbank_esx-base_8.0.1-0.20.22082334 - VMware_bootbank_trx_8.0.1-0.20.22082334 - VMware_bootbank_gc-esxio_8.0.1-0.20.22082334 - VMware_bootbank_gc_8.0.1-0.20.22082334 - VMware_bootbank_esxio_8.0.1-0.20.22082334 - VMware_bootbank_esx-xserver_8.0.1-0.20.22082334 - VMware_bootbank_cpu-microcode_8.0.1-0.20.22082334 - VMware_bootbank_esxio-combiner-esxio_8.0.1-0.20.22082334 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.20.22082334

	- VMware_bootbank_vsanhealth_8.0.1-0.20.22082334 - VMware_bootbank_vds-vsip_8.0.1-0.20.22082334 - VMware_bootbank_clusterstore_8.0.1-0.20.22082334 - VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.20.22082334 - VMware_bootbank_esxio-combiner_8.0.1-0.20.22082334 - VMware_bootbank_esxio-base_8.0.1-0.20.22082334 - VMware_bootbank_vdfs_8.0.1-0.20.22082334 - VMware_bootbank_crx_8.0.1-0.20.22082334 - VMware_bootbank_esx-update_8.0.1-0.20.22082334 - VMware_bootbank_loadesx_8.0.1-0.20.22082334 - VMware_bootbank_loadesxio_8.0.1-0.20.22082334 - VMware_bootbank_esxio-update_8.0.1-0.20.22082334
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149
Related CVE numbers	N/A

This patch updates the following issues:

ESXi 8.0 Update 1c provides the following security updates:

- The Envoy proxy is updated to version v1.23.9.
- The ESXi userworld libxml2 library is updated to version 2.10.4.
- The cURL library is updated to version 8.0.1.
- The Go library is updated to version 1.19.9.
- The etcd package is updated to 3.4.25.

ESXi 8.0 U1c - 22088125

Name	ESXi
Version	ESXi80U1c-22088125
Release Date	July 27, 2023
Category	Bugfix
Affected Components	- ESXi Component - ESXi Install/Upgrade Component - Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers for VMware ESXi - Broadcom NetXtreme I ESX VMKAPI ethernet driver - VMware NVMe over TCP Driver
PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236,

	3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3038908, 3179111, 3152476, 3187446, 3219191
Related CVE numbers	N/A

ESXi 8.0 U1sc - 22082334

Name	ESXi
Version	ESXi-8.0U1sc-22082334
Release Date	July 27, 2023
Category	Security
Affected Components	• ESXi Component • ESXi Install/Upgrade Component • ESXi Tools Component
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149, 3186166
Related CVE numbers	N/A

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
```

```
Validation of vmknics failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') . The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.`

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no`

stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 1a Release Notes

This document contains the following sections

- [Introduction](#)

- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0 Update 1a | JUNE 01 2023 | Build 21813344
Check for additions and updates to these release notes.

What's New

ESXi 8.0 Update 1a is applicable only for vSAN enabled ESXi hosts, configured with deduplication and compression. For more details, see the [Resolved Issues](#) section.

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0 Update 1 Release Notes](#)
- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Patches Contained in This Release

Patch for VMware ESXi 8.0 Update 1a

This release of ESXi 8.0 Update 1a delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U1a-21813344-depot.zip
Build:	21813344
Download Size:	583.2 MB
sha256checksum:	52dd56f65367afdc0503419370cee7bf62965940cd9cca3af6b9c2409e615950
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.1-0.10.21813344	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.1-0.10.21813344	Bugfix	Critical

Component	Bulletin	Category	Severity
ESXi Install/Upgrade Component	esxio-update_8.0.1-0.10.21813344	Bugfix	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi80U1a-21813344	Bugfix	Critical

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U1a-21813344-standard
ESXi-8.0U1a-21813344-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi-8.0U1a-21813344	06/01/2023	Bugfix	Bugfix image

For information about the individual components and bulletins, see the [Product Patches](#) page and the [Resolved Issues](#) section.

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

Resolved Issues

ESXi_8.0.1-0.10.21813344

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_trx_8.0.1-0.10.21813344 - VMware_bootbank_bmcad-esxio_8.0.1-0.10.21813344 - VMware_bootbank_bmcad_8.0.1-0.10.21813344 - VMware_bootbank_vsanhealth_8.0.1-0.10.21813344 - VMware_bootbank_clusterstore_8.0.1-0.10.21813344 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.10.21813344

	• VMware_bootbank_native-misc-drivers_8.0.1-0.10.21813344 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.10.21813344 • VMware_bootbank_esxio-combiner_8.0.1-0.10.21813344 • VMware_bootbank_crx_8.0.1-0.10.21813344 • VMware_bootbank_gc_8.0.1-0.10.21813344 • VMware_bootbank_vdfs_8.0.1-0.10.21813344 • VMware_bootbank_gc-esxio_8.0.1-0.10.21813344 • VMware_bootbank_esx-base_8.0.1-0.10.21813344 • VMware_bootbank_esxio-combiner-esxio_8.0.1-0.10.21813344 • VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.10.21813344 • VMware_bootbank_esxio-base_8.0.1-0.10.21813344 • VMware_bootbank_vds-vsip_8.0.1-0.10.21813344 • VMware_bootbank_vsan_8.0.1-0.10.21813344 • VMware_bootbank_esx-xserver_8.0.1-0.10.21813344 • VMware_bootbank_cpu-microcode_8.0.1-0.10.21813344 • VMware_bootbank_esxio_8.0.1-0.10.21813344
PRs Fixed	3217456
CVE numbers	N/A

This patch updates the `vsan` VIB. Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `trx`, `bmc-al-esxio`, `bmc-al`, `vsanhealth`, `clusterstore`, `esx-dvfilter-generic-fastpath`, `native-misc-drivers`, `esxio-dvfilter-generic-fastpath`, `esxio-combiner`, `crx`, `gc`, `vdfs`, `gc-esxio`, `esx-base`, `esxio-combiner-esxio`, `native-misc-drivers-esxio`, `esxio-base`, `vds-vsip`, `esx-xserver`, `cpu-microcode`, and `esxio`.

vSAN hosts with deduplication and compression might fail after upgrade to ESXi 8.0 Update 1

This issue occurs only on vSAN all-flash hosts with deduplication and compression enabled and only when:

1. You have an existing disk group created on a vSAN version earlier than vSAN 7.0 Update 3.
2. You have added new capacity disks to such a disk group.
3. You upgrade the vSAN host to ESXi 8.0 Update 1.

The issue does not impact vSAN hosts with compression-only or checksum-only disk groups.

ESXi 8.0 Update 1 includes a background scanner to collect metadata. This background scanner might not recognize new disks added to an existing disk group. This issue can cause the ESXi host to fail with a purple diagnostic screen. Because the host fails and reboots before the background scan completes, repeated failures might occur.

The following details appear in the backtrace:

Panic Details: Crash at 2023-05-01T11:31:12.957Z on CPU 51 running world

2019239 - VSAN_0x450080017780_PLOG. VMK Uptime:0:09:54:11.108

Panic Message: @BlueScreen: Failed at bora/modules/vmkernel/plog/dedup/dedup_util.h:231 -- VMK_ASSERT(blkNum < DDPGetEndBlkNumber(&ddCtx->devices[devIndex].sb, blkType, blkNumCheck, 1))

For more information, see VMware knowledge base article [92458](#).

This issue is resolved in this release.

esx-update_8.0.1-0.10.21813344

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-update_8.0.1-0.10.21813344 - VMware_bootbank_loadesx_8.0.1-0.10.21813344
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `esx-update` and `loadesx`.

esxio-update_8.0.1-0.10.21813344

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-update_8.0.1-0.10.21813344 - VMware_bootbank_loadesxio_8.0.1-0.10.21813344
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `esxio-update` and `loadesxio`.

ESXi-8.0U1a-21813344-standard

Profile Name	ESXi-8.0U1a-21813344-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	June 1, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_trx_8.0.1-0.10.21813344 - VMware_bootbank_bmcal-esxio_8.0.1-0.10.21813344 - VMware_bootbank_bmcal_8.0.1-0.10.21813344 - VMware_bootbank_vsanhealth_8.0.1-0.10.21813344

	• VMware_bootbank_clusterstore_8.0.1-0.10.21813344 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.10.21813344 • VMware_bootbank_native-misc-drivers_8.0.1-0.10.21813344 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.10.21813344 • VMware_bootbank_esxio-combiner_8.0.1-0.10.21813344 • VMware_bootbank_crx_8.0.1-0.10.21813344 • VMware_bootbank_gc_8.0.1-0.10.21813344 • VMware_bootbank_vdfs_8.0.1-0.10.21813344 • VMware_bootbank_gc-esxio_8.0.1-0.10.21813344 • VMware_bootbank_esx-base_8.0.1-0.10.21813344 • VMware_bootbank_esxio-combiner-esxio_8.0.1-0.10.21813344 • VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.10.21813344 • VMware_bootbank_esxio-base_8.0.1-0.10.21813344 • VMware_bootbank_vds-vsip_8.0.1-0.10.21813344 • VMware_bootbank_vsan_8.0.1-0.10.21813344 • VMware_bootbank_esx-xserver_8.0.1-0.10.21813344 • VMware_bootbank_cpu-microcode_8.0.1-0.10.21813344 • VMware_bootbank_esxio_8.0.1-0.10.21813344 • VMware_bootbank_esx-update_8.0.1-0.10.21813344 • VMware_bootbank_loadesx_8.0.1-0.10.21813344 • VMware_bootbank_esxio-update_8.0.1-0.10.21813344 • VMware_bootbank_loadesxio_8.0.1-0.10.21813344
PRs Fixed	3154255
Related CVE numbers	N/A

This patch updates the following issue:

This issue occurs only on vSAN all-flash hosts with deduplication and compression enabled and only when:

1. You have an existing disk group created on a vSAN version earlier than vSAN 7.0 Update 3.
2. You have added new capacity disks to such a disk group.
3. You upgrade the vSAN host to ESXi 8.0 Update 1.

The issue does not impact vSAN hosts with compression-only or checksum-only disk groups.

ESXi 8.0 Update 1 includes a background scanner to collect metadata. This background scanner might not recognize new disks added to an existing disk group. This issue can cause the ESXi host to fail with a purple diagnostic screen. Because the host fails and reboots before the background scan completes, repeated failures might occur.

The following details appear in the backtrace:

Panic Details: Crash at 2023-05-01T11:31:12.957Z on CPU 51 running world

2019239 - VSAN_0x450080017780_PLOG. VMK Uptime:0:09:54:11.108

Panic Message: @BlueScreen: Failed at bora/modules/vmkernel/plog/dedup/dedup_util.h:231 -- VMK_ASSERT(blkNum < DDPGetEndBlkNumber(&ddCtx->devices[devIndex].sb, blkType, blkNumCheck, 1))

For more information, see VMware knowledge base article [92458](#).

ESXi-8.0U1a-21813344-no-tools

Profile Name	ESXi-8.0U1a-21813344-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	June 1, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_trx_8.0.1-0.10.21813344 - VMware_bootbank_bmcal-esxio_8.0.1-0.10.21813344 - VMware_bootbank_bmcal_8.0.1-0.10.21813344 - VMware_bootbank_vsanhealth_8.0.1-0.10.21813344 - VMware_bootbank_clusterstore_8.0.1-0.10.21813344 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.10.21813344 - VMware_bootbank_native-misc-drivers_8.0.1-0.10.21813344 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.10.21813344 - VMware_bootbank_esxio-combiner_8.0.1-0.10.21813344 - VMware_bootbank_crx_8.0.1-0.10.21813344 - VMware_bootbank_gc_8.0.1-0.10.21813344 - VMware_bootbank_vdfs_8.0.1-0.10.21813344 - VMware_bootbank_gc-esxio_8.0.1-0.10.21813344 - VMware_bootbank_esx-base_8.0.1-0.10.21813344 - VMware_bootbank_esxio-combiner-esxio_8.0.1-0.10.21813344 - VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.10.21813344 - VMware_bootbank_esxio-base_8.0.1-0.10.21813344 - VMware_bootbank_vds-vsip_8.0.1-0.10.21813344 - VMware_bootbank_vsan_8.0.1-0.10.21813344 - VMware_bootbank_esx-xserver_8.0.1-0.10.21813344 - VMware_bootbank_cpu-microcode_8.0.1-0.10.21813344 - VMware_bootbank_esxio_8.0.1-0.10.21813344 - VMware_bootbank_esx-update_8.0.1-0.10.21813344 - VMware_bootbank_loadesx_8.0.1-0.10.21813344 - VMware_bootbank_esxio-update_8.0.1-0.10.21813344 - VMware_bootbank_loadesxio_8.0.1-0.10.21813344
PRs Fixed	3154255
Related CVE numbers	N/A

This patch updates the following issue:

This issue occurs only on vSAN all-flash hosts with deduplication and compression enabled and only when:

1. You have an existing disk group created on a vSAN version earlier than vSAN 7.0 Update 3.
2. You have added new capacity disks to such a disk group.
3. You upgrade the vSAN host to ESXi 8.0 Update 1.

The issue does not impact vSAN hosts with compression-only or checksum-only disk groups.

ESXi 8.0 Update 1 includes a background scanner to collect metadata. This background scanner might not recognize new disks added to an existing disk group. This issue can cause the ESXi host to fail with a purple diagnostic screen. Because the host fails and reboots before the background scan completes, repeated failures might occur.

The following details appear in the backtrace:

```
Panic Details: Crash at 2023-05-01T11:31:12.957Z on CPU 51 running world
```

```
2019239 - VSAN_0x450080017780_PLOG. VMK Uptime:0:09:54:11.108
```

```
Panic Message: @BlueScreen: Failed at bora/modules/vmkernel/plog/dedup/dedup_util.h:231 --
VMK_ASSERT(blkNum < DDPGetEndBlkNumber(&ddCtx->devices[devIndex].sb, blkType, blkNumCheck,
1))
```

For more information, see VMware knowledge base article [92458](#).

ESXi-8.0U1a-21813344

Name	ESXi
Version	ESXi-8.0U1a-21813344
Release Date	June 01, 2023
Category	Bugfix
Affected Components	- ESXi Component - core ESXi VIBs - ESXi Install/Upgrade Component
PRs Fixed	3217456
Related CVE numbers	N/A

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is deactivated, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
```

```
Validation of vmknics failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0 Update 1 Release Notes

This document contains the following sections

- [Introduction](#)
- [General Availability](#)
- [What's New](#)
- [Earlier Releases of ESXi 8.0](#)
- [Patches Contained in This Release](#)
- [Installation and Upgrade Notes for This Release](#)
- [Compatibility](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0 Update 1 | 18 APR 2023 | GA ISO Build 21495797
Check for additions and updates to these release notes.

General Availability

This [ESXi 8.0 Update 1](#) release is a General Availability (GA) designation. For more information on the vSphere 8.0 IA/GA Release Model of vSphere Update releases, see [The vSphere 8 Release Model Evolves](#).

What's New

- **vSphere Configuration Profiles:** vSphere 8.0 Update 1 officially launches vSphere Configuration Profiles, which allow you to manage ESXi cluster configurations by specifying a desired host configuration at the cluster level, automate the scanning of ESXi hosts for compliance to the specified Desired Configuration and remediate any host that is not compliant. vSphere Configuration Profiles require that you use vSphere Lifecycle Manager images to manage your cluster lifecycle, a vSphere 8.0 Update 1 environment, and Enterprise Plus or vSphere+ license. For more information, see [Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level](#).
- **With vSphere 8.0 Update 1, vSphere Distributed Services Engine adds support for:**
 - NVIDIA BlueField-2 DPUs to server designs from Lenovo (Lenovo ThinkSystem SR650 V2).
 - 100G NVIDIA BlueField-2 DPUs to server designs from Dell.
 - UPTv2 for NVIDIA BlueField-2 DPUs.
 - AMD Genoa CPU based server designs from Dell.
- **Recorder for required privileges on vCenter Server workflows:** With vSphere 8.0 Update 1, to help you identify the minimal set of privileges that are required to run a vCenter Server workflow, you can record workflows and store the

required privileges. You can use the stored data to create specific roles for users that run the same workflows, similar to templates. Searching through vCenter Server workflow recordings allows you to retrieve lists of privilege checks along with the corresponding sessions, users, managed objects, and operation IDs (opIDs). For more information, see [Performing Privilege Checks Operations](#).

- **Support for heterogeneous virtual graphics processing unit (vGPU) profiles on the same GPU hardware:** vSphere 8.0 Update 1 removes the requirement that all vGPUs on a physical GPU must be of the same type and you can set different vGPU profiles, such as compute, graphics, or Virtual Desktop Infrastructure workload, on one GPU to save cost by higher GPU utilization and reduced workload fragmentation.
- **Integration of VMware Skyline™ Health Diagnostics™ with vCenter:** Starting with vSphere 8.0 Update 1, you can detect and remediate issues in your vSphere environment by using the VMware Skyline Health Diagnostics self-service diagnostics platform, which is integrated with the vSphere Client. For more information, see [VMware Skyline Health Diagnostics for vSphere Documentation](#).
- **VM-level power consumption metrics:** Starting with vSphere 8.0 Update 1, you as a vSphere admin can track power consumption at a VM level to support the environmental, social, and governance goals of your organization.
- **NVSwitch support:** vSphere 8.0 Update 1 adds support to up to 8 GPUs with NVSwitch connections between them on a single ESXi host. With the NVSwitch technology, you can run high-performance computing (HPC) and AI applications such as deep learning, scientific simulations, and big data analytics, which require multiple GPUs working together in parallel.
- **Okta Identity Federation for vCenter:** With vSphere 8.0 Update 1, you can configure vCenter Server Identity Provider Federation for Okta as an external identity provider. For more information, see [Configure vCenter Server Identity Provider Federation for Okta](#).
- **Support for Fault Tolerance of virtual machines that use a virtual TPM (vTPM) module:** With vSphere 8.0 Update 1, you can use FT for VMs with a vTPM module to ensure continuous availability and security for your mission critical VMs.
- **Quick Boot support for servers with TPM 2.0 chips:** With vSphere 8.0 Update 1, you no longer need to disable TPM 2.0 for Quick Boot, which allows you to save time and eliminate security gaps.
- **vSphere API for Storage Awareness (VASA) version 5 for vSphere Virtual Volumes:** VASA version 5 for vSphere Virtual Volumes enhances security, certification management for multi vCenter deployments, and usability support. VASA version 5 deprecates self-signed certificates but provides backward compatibility support for earlier VASA versions. Before upgrading to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#) to ensure continued availability of vSphere Virtual Volumes datastores.
- **Sidecar files become regular files in Config-vVol instead of vSphere Virtual Volumes objects:** To improve scalability and performance of vSphere Virtual Volumes, starting with vSphere 8.0 Update 1, sidecar files are no longer vSphere Virtual Volumes objects but regular files in Config-vVol. When sidecar files are created as vSphere Virtual Volumes objects, for solutions such as First Class Disk (FCD) that create a large number of small sidecar files, they can cause an overhead of VASA operations such as binding and unbinding the volumes to a protocol endpoint. Creating sidecar files under the Config-vVol home directory removes such overhead from VASA and storage. After updating to ESXi 8.0 Update 1, new virtual machines, or virtual disks in the Config-vVol namespace are not supported on ESXi hosts of earlier versions. For more information and resolution, see VMware knowledge base article [90791](#).
- **Increased default capacity for vSphere Virtual Volumes objects of type Config-vVol:** Starting with vSphere 8.0 Update 1, to allow the use of folders over vSphere Virtual Volumes datastores as content repositories, vSphere Virtual Volumes objects of type Config-vVol are created by default as thin-provisioned 255 GB, up from 4 GB in earlier releases. Also, starting with ESXi 8.0 Update 1, Config-vVol objects use VMFS-6 format instead of VMFS-5.
- **NVMe over TCP support for vSphere Virtual Volumes:** vSphere 8.0 Update 1 adds NVMe over TCP support for vSphere Virtual Volumes.
- **Extended XCOPY support:** With vSphere 8.0 Update 1, you can use `EXTENDED COPY` (`XCOPY`) commands to optimize the data copy between VMFS datastores that are in different storage arrays, not only between datastores within the same array. `XCOPY` operations across storage arrays are supported only for VMFS 6 datastores and allow you to offload, migrate and clone workloads. While ESXi enables this feature, the actual data migration across the

arrays must be implemented on the storage array side. For more information on how to use SCSI T10 commands, see [Managing Hardware Acceleration on Block Storage Devices](#).

- **New file type for OSDATA volumes on SSD devices:** vSphere 8.0 Update 1 adds a new file system type, VMFSOS, specifically for the ESX-OSData system partition on local SSD devices, which allows you to continue using virtual flash resources on other devices. The new file type prevents cases when you format an ESX-OSData volume on a local SSD device, and `fsType` returns a file of type Virtual Flash File System (VFFS). As a result, the disk backing of the ESX-OSData volume is listed under the Virtual Flash resources in vCenter, but such a disk belongs to the ESX-OSData volume and is not a part of the Virtual Flash resource pool.
- **NFS traffic isolation enhancements:** With vSphere 8.0 Update 1, to isolate NFS traffic, you can bind an NFS3 datastore on an NFS share over the network to a VM kernel adaptor on an ESXi host from a cluster. VMKNIC port binding is not supported in this release for vSphere Virtual Volumes datastores backed with NFS3 configurations. For more information, see [Configure VMkernel Binding for NFS 3 Datastore](#).
- **Fourfold increase in NVMe-oF namespace capacity:** Starting with vSphere 8.0 Update 1, you can use 32 paths to a NVMe-oF namespace from 8 in earlier releases.
- **Support for end-to-end NVMe stack without protocol translation:** With vSphere 8.0 Update 1, the Pluggable Storage Architecture (PSA) extends NVMe capabilities to support end-to-end NVMe stack without any protocol translation in any of the layers.
vSphere 8.0 Update 1 also extends NVMe functionality by adding support for third-party multi-pathing plug-ins to control and manage NVMe arrays.
- **Increased maximum for Windows Server Failover Clusters (WSFC):** vSphere 8.0 Update 1 increases the maximum WSFC clusters per ESXi host to 16 from 3.
- **Scale the number of NVMe over TCP adapters:** vSphere 8.0 Update 1 increases the number of NVMe over TCP storage adapters per ESXi host to 8 from 2 to allow more redundancy and better bandwidth.
- **Support hot-add and hot-remove for VMDirectPath I/O devices:** With vSphere 8.0 Update 1, by using vSphere API you can add or remove a VMDirectPath I/O device without powering off VMs. For more information, see [Hot-add Support for VMDirectPath I/O Devices](#).
- **Increased number of PCI passthrough devices per VM:** With vSphere 8.0 Update 1, you can add up to 64 PCI passthrough devices per VM, up from 32 in previous releases, to allow better support for Virtualized Radio Access Network (vRAN) use cases.
- **Local depot overrides for Remote Office and Branch Office (ROBO) standalone ESXi hosts:** Starting with vSphere 8.0 Update 1, to meet demand of certain business use cases, you can manage local depot overrides for Remote Office and Branch Office (ROBO) standalone ESXi hosts in the vSphere Client and by using the vSphere Automation API. For more information, see [Manage Depot Overrides for a Cluster or a Standalone Host](#).
- **Advanced filters in the vSphere Client:** With vSphere 8.0 Update 1, in addition to the existing quick filter option, you can set advanced filters on most inventory objects in the vSphere Client, such as VM lists, Hosts and Clusters, Datastores, and linked vCenters. You can use strings, such as Name, Status, and Guest OS, and enumerations, such as IP address. To refine your filters, you can use various logical operators and 2 conditions.
- **Increased maximum number of NFSv3 datastores that can be mounted with multiple connections:** Starting with vSphere 8.0 Update 1, you can configure up to 8 connections to NFSv3 datastores on an ESXi host, from 1 in earlier releases, and have up to 256 connections used by the NFSv3 datastores on the host, depending on the combination of the number of datastores and the number of connection configured for each datastore. For more information, see VMware knowledge base article [91481](#).
- **HTTP/JSON-based wire protocol as an alternative to SOAP/XML:** vSphere 8.0 Update 1 adds support for a new HTTP/JSON-based wire protocol as an alternative to SOAP/XML. The new protocol is described using the industry standard OpenAPI specification, version 3.0, and you can use it to access the popular VIM APIs. For more information, see the OpenAPI schema and Web Services Programming Guide, part of the [vSphere Management SDK](#).
- This release resolves CVE-2023-1017. VMware has evaluated the severity of this issue to be in the low severity range with a maximum CVSSv3 base score of [3.3](#).
- This release resolves CVE-2023-1018. VMware has evaluated the severity of this issue to be in the low severity range with a maximum CVSSv3 base score of [3.3](#).

Earlier Releases of ESXi 8.0

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- [VMware ESXi 8.0c Release Notes](#)
- [VMware ESXi 8.0b Release Notes](#)
- [VMware ESXi 8.0a Release Notes](#)

For internationalization, compatibility, and open source components, see the [VMware vSphere 8.0 Release Notes](#).

Patches Contained in This Release

Patch for VMware ESXi 8.0 Update 1

This release of ESXi 8.0 Update 1 delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U1-21495797-depot.zip
Build:	21495797
Download Size:	583.5 MB
sha256checksum:	68d9b58c9a00ec63fd899cda1a3d044d1fb0f2ee353bcf9c8e163483977f5d7c
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi80U1-21495797	Bugfix	Critical

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U1-21495797-standard-7698281098777227235
ESXi-8.0U1-21495797-no-tools1278187871193924799

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi_8.0.1-0.0.21495797	04/18/2023	General	Bugfix image

For information about the individual components and bulletins, see the [Resolved Issues](#) section.

Patch Download and Installation

You can download ESXi 8.0 Update 2b from the [Broadcom Support Portal](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Installation and Upgrade Notes for This Release

Before upgrading to VASA version 5, refer to VMware knowledge base article [91387](#) to ensure continued availability of your vSphere Virtual Volumes datastores.

Compatibility

Newly created virtual machines or virtual disks of version ESXi 8.0 Update 1 in the Config-vVol namespace of vSphere Virtual Volumes are not supported on ESXi hosts of earlier versions. For more information and resolution, see VMware knowledge base article [90791](#).

Resolved Issues

Networking Issues

First Class Disk (FCD) sync might fail on NFSv3 datastores and you see inconsistent values in Govc reports

A rare file corruption issue might occur when multiple ESXi hosts do a simultaneous file append operation on an NFSv3 datastore. As a result, if you use the Govc open source command-line utility to perform administrative actions on a vCenter instance, you might see reports with inconsistent values. For example, if you run the command `govc disk.ls` to list all FCDs in a newly created datastore, you see a different number of disks from what you expect.

This issue is resolved in this release.

Installation, Upgrade, and Migration Issues

If the Internet Control Message Protocol (ICMPA) is not active, ESXi host reboot might take long after upgrading to vSphere 8.0 and later

If ICMPA is not active on the NFS servers in your environment, after upgrading your system to vSphere 8.0 and later, ESXi hosts reboot might take an hour to complete, because restore operations for NFS datastores fail. NFS uses the `vmkping` utility to identify reachable IPs of the NFS servers before executing a mount operation and when ICMP is not active, mount operations fail.

This issue is resolved in this release. To remove dependency on the ICMP protocol to find reachable IPs, the fix adds socket APIs to ensure that IPs on a given NFS server are available.

Miscellaneous Issues

The autocomplete attribute in forms for web access is not disabled on password and username fields

Your remote web server might contain an HTML form field that has an input of type `password` and `username` where the `autocomplete` attribute is not set to `off`. As a result, you might see a warning by security scanners such as

AutoComplete Attribute Not Disabled for Password in Form Based Authentication. The setting of the `autocomplete` attribute to `on` does not directly put a risk on web servers, but in certain browsers, user credentials in such forms might be saved and potentially lead to a loss of confidentiality.

This issue is resolved in this release. The fix makes sure the attribute `autocomplete` is set to `off` in password and username fields to prevent browsers from caching credentials.

Known Issues

Installation, Upgrade, and Migration Issues

If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible

Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#).

Miscellaneous Issues

You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands

When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

```
[::~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1
Validation of vmknict failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1) Input(): Not found:
```

The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)

The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS +VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

Workaround: None

Networking Issues

Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations

vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

Workaround: Use the command `esxcli nvme fabrics discover -a vmhbaX -D` to manually generate the mDNS packets.

ESXi reboot takes long due to NFS server mount timeout

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is not active, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS.`

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed in the backtrace. The snapshot creation task completes upon retry.`

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') . The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.`

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_XXX. Extraction of image from host XXX.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware ESXi 8.0f Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Upgrade Notes](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0f | JUL 15 2025 | ISO Build 24814147

Check for additions and updates to these release notes.

IMPORTANT: This release is specifically for compliance with Common Criteria assurance components and is not intended for general use. It provides only an ISO file as an offline deliverable, which you must download manually from the [Broadcom Support Portal](#).

What's New

- This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, and CVE-2025-41227.

Upgrade Notes

The recommended upgrade path from ESXi 8.0f is to ESXi 8.0 Update 3f or later.

Patches Contained in This Release

VMware ESXi 8.0f

Build Details

VMware vSphere Hypervisor (ESXi ISO) image

Download Filename:	VMware-VMvisor-Installer-8.0f-24814147.x86_64.iso
Build:	24814147
Download Size:	621 MB
SHA256 checksum:	0e9a00c534edcb2f9dd39472f650784419c4e99012091d12baef3f97109d4e3d
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

Resolved Issues

Security Issues

This release resolves CVE-2025-41236, CVE-2025-41237, CVE-2025-41238, CVE-2025-41239, and CVE-2025-41227.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch `vSwitchBMC` with uplink `vusb0`

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, `vUSB` or `vNIC` respectively, an additional standard switch `vSwitchBMC` with

uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled`.

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') .` The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size of at least 2048 bits for signature generation.

Guest OS Issues

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

VMware ESXi 8.0e Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Upgrade Notes](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0e | MAR 11 2025 | ISO Build 24569005

Check for additions and updates to these release notes.

IMPORTANT: This release is specifically for compliance with Common Criteria assurance components and is not intended for general use. It provides only an ISO file as an offline deliverable, which you must download manually from the [Broadcom Support Portal](#).

What's New

- This release resolves CVE-2025-22224 and CVE-2025-22226.

Upgrade Notes

The recommended upgrade path from ESXi 8.0e is to ESXi 8.0 Update 3d or later.

Patches Contained in This Release

VMware ESXi 8.0e

Build Details

VMware vSphere Hypervisor (ESXi ISO) image

Download Filename:	VMware-VMvisor-Installer-8.0e-24569005.x86_64.iso
Build:	24569005
Download Size:	621 MB
SHA256 checksum:	ec5ab1fa77814416f1d997be48faf06749b82bdea8cd72d9eb88f70e5e33aaf0
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

Security Issues

This release resolves CVE-2025-22224 and CVE-2025-22226.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliance$` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the ntg3 driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on ntg3 drivers of versions earlier than 4.1.3 or the tg3 driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a ntg3 driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Guest OS Issues

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

VMware ESXi 8.0d Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Upgrade Notes](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0d | OCT 7 2024 | ISO Build 24118393

Check for additions and updates to these release notes.

IMPORTANT: This release is specifically for compliance with Common Criteria assurance components and is not intended for general use. It provides only an ISO file as an offline deliverable, which you must download manually from the [Broadcom Support Portal](#).

What's New

- This release resolves CVE-2024-22253, CVE-2024-22255, CVE-2024-22252, CVE-2024-22254, CVE-2024-37086, and CVE-2024-22273.
- ESXi 8.0d adds the following audit events:
 - service.access.denied
 - service.apikey.get
 - service.apikey.use
 - service.exttoken.use
 - service.password.use
 - service.lcticket.get
 - service.lcticket.use
- ESXi 8.0d removes the following audit events:
 - vapi.connect
 - vapi.disconnect

Upgrade Notes

The recommended upgrade path from ESXi 8.0d is to ESXi 8.0 Update 3 or later.

Patches Contained in This Release

VMware ESXi 8.0d

Build Details

VMware vSphere Hypervisor (ESXi ISO) image

Download Filename:	VMware-VMvisor-Installer-8.0d-24118393.x86_64.iso
Build:	24118393
Download Size:	621 MB
SHA256 checksum:	d3a445abab476a308df04d9608a7484d4bc3503a0e58d78a51c1605c78a9779a
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

For details on updates and upgrades by using vSphere Lifecycle Manager, see [About vSphere Lifecycle Manager](#) and [vSphere Lifecycle Manager Baselines and Images](#). You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually [download](#) the patch offline bundle ZIP file.

For more information, see the [Upgrading Hosts by Using ESXCLI Commands](#) and the [VMware ESXi Upgrade](#) guide.

Resolved Issues

Miscellaneous Issues

This release resolves CVE-2024-22253, CVE-2024-22255, CVE-2024-22252, CVE-2024-22254, and CVE-2024-22273.

ESXi 8.0d adds the following audit events:

- service.access.denied
- service.apikey.get
- service.apikey.use
- service.exttoken.use
- service.password.use
- service.lcticket.get
- service.lcticket.use

ESXi 8.0d removes the following audit events:

- vapi.connect
- vapi.disconnect

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server

upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Guest OS Issues

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

VMware ESXi 8.0c Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Patch Download and Installation](#)
- [Resolved Issues](#)

- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0c | MAR 30 2023 | Build 21493926
Check for additions and updates to these release notes.

What's New

- ESXi 8.0c delivers a fix for a licensing issue on vSphere systems with the vSphere Distributed Services Engine feature enabled.

NOTE: This patch is applicable only for vSphere Distributed Services Engine enabled ESXi hosts, configured with data processing units (DPU). For more details, see the [Resolved Issues](#) section.

Patches Contained in This Release

This release of ESXi 8.0c delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0c-21493926-depot
Build:	21493926
Download Size:	604.0 MB
md5sum:	53d1368d0ed6c7b15bf89f0f22e9f617
sha256checksum:	aabf02f457ce0db951f31f6a7f9ac6a5fc049fff094e5b9a59b310548602afe6
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.0-1.25.21493926	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.0-1.25.21493926	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.0-1.25.21493926	Bugfix	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi-80c-21493926	Bugfix	Critical

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0c-21493926-standard
ESXi-8.0c-21493926-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi-80c-21493926	03/30/2023	Bugfix	Bugfix image

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

Resolved Issues

ESXi_8.0.0-1.25.21493926

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.25.21493926 - VMware_bootbank_esxio-base_8.0.0-1.25.21493926 - VMware_bootbank_esx-xserver_8.0.0-1.25.21493926 - VMware_bootbank_esxio_8.0.0-1.25.21493926 - VMware_bootbank_drivervm-gpu_8.0.0-1.25.21493926 - VMware_bootbank_cpu-microcode_8.0.0-1.25.21493926 - VMware_bootbank_vsan_8.0.0-1.25.21493926 - VMware_bootbank_trx_8.0.0-1.25.21493926 - VMware_bootbank_bmcad_8.0.0-1.25.21493926 - VMware_bootbank_clusterstore_8.0.0-1.25.21493926 - VMware_bootbank_native-misc-drivers_8.0.0-1.25.21493926 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.25.21493926 - VMware_bootbank_vsanhealth_8.0.0-1.25.21493926 - VMware_bootbank_vdfs_8.0.0-1.25.21493926 - VMware_bootbank_bmcad-esxio_8.0.0-1.25.21493926 - VMware_bootbank_gc_8.0.0-1.25.21493926 - VMware_bootbank_crx_8.0.0-1.25.21493926 - VMware_bootbank_esxio-combiner_8.0.0-1.25.21493926 - VMware_bootbank_gc-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.25.21493926

	VMware_bootbank_esx-base_8.0.0-1.25.21493926
PRs Fixed	3154255
CVE numbers	N/A

This patch updates the `esxio-dvfilter-generic-fastpath`, `esxio-base`, `esx-xserver`, `esxio`, `drivervm-gpu`, `cpu-microcode`, `vsan`, `trx`, `bmcal`, `clusterstore`, `native-misc-drivers`, `native-misc-drivers-esxio`, `esx-dvfilter-generic-fastpath`, `vsanhealth`, `vdfs`, `bmcal-esxio`, `gc`, `crx`, `esxio-combiner`, `gc-esxio`, `esxio-combiner-esxio`, and `esx-base` VIBs.

On vSphere systems with VMWare NSX installed, and the vSphere Distributed Services Engine feature enabled, some NSX config operations fail with an error for an expired license

vSphere Distributed Services Engine feature enabled ESXi hosts are servers configured with data processing units (DPUs). In NSX Manager, Transport Node configuration operations on such hosts fail with an error status of `Partial Success`. Error details point to the underlying license expiry with a message such as `current license or ESXi version prohibits execution of the requested operation`. This error prevents any further NSX configuration operations on the nodes with DPU. The issue occurs 60 days after the first installation of ESXi on DPU, which is the default evaluation license. vSphere Distributed Services Engine systems with DPUs do not require a separate ESXi license for ESXi on DPU, but in certain scenarios, a license version validator on the DPU might trigger an error.

This issue is resolved in this release. The fix makes sure that license checks in vSphere systems with DPUs occur on the ESXi host, not on DPUs.

esx-update_8.0.0-1.25.21493926

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-update_8.0.0-1.25.21493926 - VMware_bootbank_loadesx_8.0.0-1.25.21493926
PRs Fixed	N/A
CVE numbers	N/A

This patch updates the `esx-update`, and `loadesx` VIBs.

esxio-update_8.0.0-1.25.21493926

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-update_8.0.0-1.25.21493926 - Mware_bootbank_loadesxio_8.0.0-1.25.21493926
PRs Fixed	N/A

CVE numbers	N/A
--------------------	-----

This patch updates the `esxio-update`, and `loadesxio` VIBs.

ESXi-8.0c-21493926-standard

Profile Name	ESXi-8.0c-21493926-standard
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	March 30, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.25.21493926 - VMware_bootbank_esxio-base_8.0.0-1.25.21493926 - VMware_bootbank_esx-xserver_8.0.0-1.25.21493926 - VMware_bootbank_esxio_8.0.0-1.25.21493926 - VMware_bootbank_drivervm-gpu_8.0.0-1.25.21493926 - VMware_bootbank_cpu-microcode_8.0.0-1.25.21493926 - VMware_bootbank_vsan_8.0.0-1.25.21493926 - VMware_bootbank_trx_8.0.0-1.25.21493926 - VMware_bootbank_bmcad_8.0.0-1.25.21493926 - VMware_bootbank_clusterstore_8.0.0-1.25.21493926 - VMware_bootbank_native-misc-drivers_8.0.0-1.25.21493926 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.25.21493926 - VMware_bootbank_vsanhealth_8.0.0-1.25.21493926 - VMware_bootbank_vdfs_8.0.0-1.25.21493926 - VMware_bootbank_bmcad-esxio_8.0.0-1.25.21493926 - VMware_bootbank_gc_8.0.0-1.25.21493926 - VMware_bootbank_crx_8.0.0-1.25.21493926 - VMware_bootbank_esxio-combiner_8.0.0-1.25.21493926 - VMware_bootbank_gc-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esx-base_8.0.0-1.25.21493926 - VMware_bootbank_esx-update_8.0.0-1.25.21493926 - VMware_bootbank_loadesx_8.0.0-1.25.21493926 - VMware_bootbank_esxio-update_8.0.0-1.25.21493926 - Mware_bootbank_loadesxio_8.0.0-1.25.21493926
PRs Fixed	3154255
Related CVE numbers	N/A

This patch updates the following issue:

- vSphere Distributed Services Engine feature enabled ESXi hosts are servers configured with data processing units (DPUs). In NSX Manager, Transport Node configuration operations on such hosts fail with an error status of `Partial`

Success. Error details point to the underlying license expiry with a message such as `current license or ESXi version prohibits execution of the requested operation`. This error prevents any further NSX configuration operations on the nodes with DPU. The issue occurs 60 days after the first installation of ESXi on DPU, which is the default evaluation license. vSphere Distributed Services Engine systems with DPUs do not require a separate ESXi license for ESXi on DPU, but in certain scenarios, a license version validator on the DPU might trigger an error.

ESXi-8.0c-21493926-no-tools

Profile Name	ESXi-8.0c-21493926-no-tools
Build	For build information, see Patches Contained in This Release .
Vendor	VMware, Inc.
Release Date	March 30, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.25.21493926 - VMware_bootbank_esxio-base_8.0.0-1.25.21493926 - VMware_bootbank_esx-xserver_8.0.0-1.25.21493926 - VMware_bootbank_esxio_8.0.0-1.25.21493926 - VMware_bootbank_drivervm-gpu_8.0.0-1.25.21493926 - VMware_bootbank_cpu-microcode_8.0.0-1.25.21493926 - VMware_bootbank_vsan_8.0.0-1.25.21493926 - VMware_bootbank_trx_8.0.0-1.25.21493926 - VMware_bootbank_bmcsl_8.0.0-1.25.21493926 - VMware_bootbank_clusterstore_8.0.0-1.25.21493926 - VMware_bootbank_native-misc-drivers_8.0.0-1.25.21493926 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.25.21493926 - VMware_bootbank_vsanhealth_8.0.0-1.25.21493926 - VMware_bootbank_vdfs_8.0.0-1.25.21493926 - VMware_bootbank_bmcsl-esxio_8.0.0-1.25.21493926 - VMware_bootbank_gc_8.0.0-1.25.21493926 - VMware_bootbank_crx_8.0.0-1.25.21493926 - VMware_bootbank_esxio-combiner_8.0.0-1.25.21493926 - VMware_bootbank_gc-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.25.21493926 - VMware_bootbank_esx-base_8.0.0-1.25.21493926 - VMware_bootbank_esx-update_8.0.0-1.25.21493926 - VMware_bootbank_loadesx_8.0.0-1.25.21493926 - VMware_bootbank_esxio-update_8.0.0-1.25.21493926 - Mware_bootbank_loadesxio_8.0.0-1.25.21493926
PRs Fixed	3154255
Related CVE numbers	N/A

This patch updates the following issue:

- vSphere Distributed Services Engine feature enabled ESXi hosts are servers configured with data processing units (DPUs). In NSX Manager, Transport Node configuration operations on such hosts fail with an error status of `Partial Success`. Error details point to the underlying license expiry with a message such as `current license or ESXi version prohibits execution of the requested operation`. This error prevents any further NSX configuration operations on the nodes with DPU. The issue occurs 60 days after the first installation of ESXi on DPU, which is the default evaluation license. vSphere Distributed Services Engine systems with DPUs do not require a separate ESXi license for ESXi on DPU, but in certain scenarios, a license version validator on the DPU might trigger an error.

ESXi-80c-21493926

Name	ESXi
Version	ESXi-80c-21493926
Release Date	March 30, 2023
Category	Bugfix
Affected Components	- ESXi Component - core ESXi VIBs - ESXi Install/Upgrade Component
PRs Fixed	3154255
Related CVE numbers	N/A

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

- Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings > <specific broadcom NIC> > Device Configuration Menu > Media Auto Detect**. Reboot the host.
- Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss` : Use the `DRSS` parameter.
- `drss` : Use the `DRSS` parameter.
- `rss` : Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS.`

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed in the backtrace. The snapshot creation task completes upon retry.`

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') . The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.`

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

Guest OS Issues

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

VMware ESXi 8.0b Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

ESXi 8.0b | FEB 14 2023 | Build 21203435
Check for additions and updates to these release notes.

What's New

ESXi 8.0.b supports vSphere Quick Boot on the following servers:

- HPE:
 - ProLiant DL110 Gen10
 - ProLiant DL160 Gen10
 - ProLiant ML350 Gen10
- Lenovo:
 - ThinkSystem SR 665

Patches Contained in This Release

This release of ESXi 8.0b delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0b-21203435-depot.zip
Build:	21203435
Download Size:	993.8 MB
md5sum:	3311032dada17653a5b017b3f2f0d3e4
sha256checksum:	94c4e8fddaf8da2b127ba891b82fefbad6d672e7ac160477a159f462de29c5b6
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.0-1.20.21203435	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.0-1.20.21203435 esxio-update_8.0.0-1.20.21203435	Bugfix	Critical
Broadcom Emulex Connectivity Division lpfc driver for FC adapters	Broadcom-ELX-lpfc_14.0.635.4-14vmw.800.1.20.21203435	Bugfix	Critical
Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers for VMware ESXi	Mellanox-nmlx5_4.23.0.36-10.2vmw.800.1.20.21203435	Bugfix	Critical
VMware DesignWare I2C Driver	VMware-dwi2c_0.1-7vmw.800.1.20.21203435	Bugfix	Critical
Network driver for Intel(R) X710/XL710/XXV710/X722 Adapters	Intel-i40en_1.11.2.6-1vmw.800.1.20.21203435	Bugfix	Critical
Pensando Systems Native Ethernet Driver	Pensando-ionic-en_20.0.0-30vmw.800.1.20.21203435	Bugfix	Critical
ESXi Component - core ESXi VIBs	ESXi_8.0.0-1.15.21203431	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.0-1.15.21203431	Security	Critical

Component	Bulletin	Category	Severity
	esxio-update_8.0.0-1.15.21203431		
ESXi Tools Component	VMware-VM-Tools_12.1.5.20735119-21203431	Security	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi80b-21203435	Bugfix	Critical
ESXi80sb-21203431	Security	Critical

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0b-21203435-standard
ESXi-8.0b-21203435-no-tools
ESXi-8.0sb-21203431-standard
ESXi-8.0sb-21203431-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi_8.0.0-1.20.21203435	FEB 14 2023	Bugfix	Bugfix image
ESXi_8.0.0-1.15.21203431	FEB 14 2023	Security	Enhancement

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

Resolved Issues

ESXi_8.0.0-1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_esx-xserver_8.0.0-1.20.21203435 - VMware_bootbank_gc-esxio_8.0.0-1.20.21203435 - VMware_bootbank_vdfs_8.0.0-1.20.21203435

	- VMware_bootbank_esxio-combiner-esxio_8.0.0-1.20.21203435 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.20.21203435 - VMware_bootbank_crx_8.0.0-1.20.21203435 - VMware_bootbank_trx_8.0.0-1.20.21203435 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.20.21203435 - VMware_bootbank_esxio-base_8.0.0-1.20.21203435 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.20.21203435 - VMware_bootbank_esxio_8.0.0-1.20.21203435 - VMware_bootbank_cpu-microcode_8.0.0-1.20.21203435 - VMware_bootbank_clusterstore_8.0.0-1.20.21203435 - VMware_bootbank_native-misc-drivers_8.0.0-1.20.21203435 - VMware_bootbank_esxio-combiner_8.0.0-1.20.21203435 - VMware_bootbank_gc_8.0.0-1.20.21203435 - VMware_bootbank_bmcals_8.0.0-1.20.21203435 - VMware_bootbank_esx-base_8.0.0-1.20.21203435 - VMware_bootbank_vsanhealth_8.0.0-1.20.21203435 - VMware_bootbank_vsan_8.0.0-1.20.21203435 - VMware_bootbank_drivervm-gpu_8.0.0-1.20.21203435 - VMware_bootbank_bmcals-esxio_8.0.0-1.20.21203435 - VMware_bootbank_esx-ui_2.10.1-20925449
PRs Fixed	3053428, 3051651, 3057027, 3035652, 3065063, 3051004, 3062043, 3070831, 3065990, 3065990, 3041101, 2988813, 3047637, 3044476, 3070865, 3078203, 3055341, 3081041, 3035652, 3051140, 3041537, 3050240, 3044476, 3051933, 3057524, 3048660, 3053504, 3026095, 3026096, 3064350, 3064715, 3058714, 3031741, 3064723, 3064425, 3055431, 3064226, 3027341, 3063928, 3064308, 3064605, 3053812, 3053705, 3050596, 3062043, 3050878, 3030641, 3031898, 3057619, 3064334, 3051004, 3047546, 3064068, 3064055, 3055374, 3064059, 3064063, 3064062, 3064054, 3064053, 3062062, 3064096, 3066138, 3065063, 3064848, 3068605, 3064060, 3034879, 3051651, 3081884, 3078334, 3073025, 3070865, 3073230, 3071612, 3070831, 3068028, 3071614, 3071981, 3072190, 3065990, 3081686, 3062907, 3055084, 3069024, 3069113, 3061447, 3065694
CVE numbers	N/A

Updates the **esx-xserver**, **gc-esxio**, **vdfs**, **esxio-combiner-esxio**, **native-misc-drivers-esxio**, **trx**, **esx-dvfilter-generic-fastpath**, **esxio-base**, **esxio-dvfilter-generic-fastpath**, **esxio**, **cpu-microcode**, **clusterstore**, **native-misc-drivers**, **esxio-combiner**, **gc**, **bmcals**, **esx-base**, **vsanhealth**, **vsan**, **drivervm-gpu**, **bmcals-esxio**, **esx-ui**, and **crx** VIBs to resolve the following issues:

- PR 3053428: Direct Console User Interface (DCUI) does not display network adapters**

When you select **Network Adapters** in the DCUI, you see no output and cannot modify the network configuration. In the **/var/log/syslog.log** file, you see an error such as **2022-xx-xx:02:02.xxxx Er(11) DCUI[265583]: Scanning dvswitches failed: Unable to get the dvs name: Status(bad0001)= Failure**

This issue is resolved in this release.

- **PR 3051651: Virtual machines might intermittently experience soft lockups inside the guest kernel on Intel machines**

Starting with 7.0 Update 2, ESXi supports Posted Interrupts (PI) on Intel CPUs for PCI passthrough devices to improve the overall system performance. In some cases, a race between PIs and the VMkernel scheduling might occur. As a result, virtual machines that are configured with PCI passthrough devices with normal or low latency sensitivity might experience soft lockups.

This issue is resolved in this release.

- **PR 3057027: When you change the NUMA nodes Per Socket (NPS) setting on an AMD EPYC 9004 'Genoa' server from the default of 1, you might see incorrect number of sockets**

In the vSphere Client and the ESXi Host Client, you might see incorrect number of sockets and respectively cores per socket when you change the NPS setting on an AMD EPYC 9004 server from the default of **1**, or **Auto**, to another value, such as **2** or **4**.

This issue is resolved in this release.

- **PR 3035652: ESXi hosts might fail with a purple diagnostic screen when a VM with bus sharing set to Physical issues a SCSI-2 reserve command**

Windows 2012 and later use SCSI-3 reservation for resource arbitration to support Windows failover clustering (WSFC) on ESXi for cluster-across-box (CAB) configurations. However, if you configure the bus sharing of the SCSI controller on that VM to Physical, the **SCSI RESERVE** command causes the ESXi host to fail with a purple diagnostic screen. **SCSI RESERVE** is SCSI-2 semantic and is not supported with WSFC clusters on ESXi.

This issue is resolved in this release.

- **PR 3065063: Virtual machines might take longer to respond or become temporary unresponsive when vSphere DRS triggers multiple migration operations**

In rare cases, multiple concurrent open and close operations on datastores, for example when vSphere DRS triggers multiple migration tasks, might cause some VMs to take longer to respond or become unresponsive for a short time. The issue is more likely to occur on VMs that have multiple virtual disks over different datastores.

This issue is resolved in this release.

- **PR 3051004: You might see incorrect facility and severity values in ESXi syslog log messages**

The logger command of some logger clients might not pass the right facility code to the ESXi syslog daemon and cause incorrect facility or severity values in ESXi syslog log messages.

This issue is resolved in this release.

- **PR 3062043: vSAN File Service fails to enable when an isolated witness network is configured**

vSAN File Service requires hosts to communicate with each other. File Service might incorrectly use an IP address in the witness network for inter-communication. If you have configured an isolated witness network for vSAN, the host can communicate with a witness node over the witness network, but hosts cannot communicate with each other over the witness network. Communication between hosts for vSAN File Service cannot be established.

This issue is resolved in this release.

- **PR 3070831: ESXi hosts might fail with a purple diagnostic screen after a migration operation**

A rare race condition caused by memory reclamation for some performance optimization processes in vSphere vMotion might cause ESXi hosts to fail with a purple diagnostic screen.

This issue is resolved in this release.

- **PR 3065990: On vSphere systems with number of CPUs more than 448, ESXi hosts might fail with a purple diagnostic screen upon upgrade or reboot**

In rare cases, on vSphere systems with number of CPUs more than 448, the calculation of memory utilization overflows the 32 bit value which rounds up to a value lower than the expected to be allocated. As a result, ESXi hosts might fail with a purple diagnostic screen upon upgrade or reboot.

This issue is resolved in this release.

- **PR 3041101: When TPM 2.0 is enabled with TXT on an ESXi host, attempts to power-on a virtual machine might fail**

When TXT is enabled on an ESX host, attempts to power-on a VM might fail with an error. In the vSphere Client, you see a message such as **This host supports Intel VT-x, but Intel VT-x is restricted. Intel VT-x might be restricted**

because 'trusted execution' has been enabled in the BIOS/firmware settings or because the host has not been power-cycled since changing this setting.

This issue is resolved in this release.

- **PR 2988813: If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation**

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

This issue is resolved in this release.

- **PR 3047637: ESXi hosts might fail with a purple diagnostic screen during a Fast Suspend Resume (FSR) operation on a VM with Transparent Page Sharing (TPS) enabled**

Due to a rare race condition during a Fast Suspend Resume (FSR) operation on a VM with shared pages between VMs, ESXi hosts might fail with a purple diagnostic screen with an error such as **PFrame_IsBackedByLPage** in the backtrace.

This issue is resolved in this release.

- **PR 3044476: Virtual machines might fail with ESX unrecoverable error due to a rare issue with handling AVX2 instructions**

Due to a rare issue with handling AVX2 instructions, a virtual machine of version ESX 8.0 might fail with **ESX unrecoverable** error. In the **vmware.log** file, you see a message such as: **MONITOR PANIC: vcpu-0:VMM fault 6: src=MONITOR ...**. The issue is specific for virtual machines with hardware versions 12 or earlier.

This issue is resolved in this release.

- **PR 3041057: The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)**

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

This issue is resolved in this release.

- **PR 3041057: Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled**

If the advanced processor setting **Enable IOMMU in this virtual machine** is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

This issue is resolved in this release.

- **PR 3070865: When you upgrade an ESXi host in a vSAN cluster of version earlier than vSAN 6.2 to ESXi 8.0, the vSAN network configuration might be lost**

For ESXi hosts in vSAN clusters of version earlier than vSAN 6.2, upgrades to ESXi 8.0 might cause loss of the VMkernel port tagging for vSAN. As a result, the vSAN networking configuration does not exist after the upgrade.

This issue is resolved in this release. For earlier releases, you can manually fix the issue by enabling VMkernel traffic for vSAN on the ESXi host after the upgrade.

- **PR 3078203: Storage I/O Control quickly generates a large volume of logs marked as critical that might also fill up the datastore**

In VMware Aria Operations for Logs, formerly vRealize Log Insight, you might see a large volume of logs generated by Storage I/O Control such as **Invalid share value: 0. Using default.** and **Skipping device naa.xxxx either due to VSI read error or abnormal state.** The volume of logs varies depending on the number of ESXi hosts in a cluster and the number of devices in switched off state. When the issue occurs, the log volume generates quickly, within 24 hours, and VMware Aria Operations for Logs might classify the messages as critical. However, such logs are harmless and do not impact the operations on other datastores that are online.

This issue is resolved in this release. The fix moves such logs from error to trivia to prevent misleading logging.

- **PR 3055341: When the ToolsRamdisk advanced option is active on an ESXi host, after a VM Tools upgrade, new versions are not available for VMs**

When you enable the **ToolsRamdisk** advanced option that makes sure the **/vmtools** partition is always on a RAM disk rather than on a USB or SD-card, installing or updating the **tools-light** VIB does not update the content of the RAM disk automatically. Instead, you must reboot the host so that the RAM disk updates and the new VM Tools version is available for the VMs.

This issue is resolved in this release. The fix makes sure that when **ToolsRamdisk** is enabled, VM Tools updates automatically, without the need of manual steps or restart.

- **PR 3081041: You cannot export lists of Inventory objects**

When you try to export a file with a list of Inventory objects, such as VMs, Hosts, and Datastores, the task fails with an error **Export Data Failure**.

This issue is resolved in this release.

- **PR 3068605: Memory problems might cause vSAN File Service errors**

If the vSAN sockrelay runs out of memory, it can cause spurious errors in the vSAN File Service. The vSAN health service shows numerous file service health check errors. The following entry appears in the **/var/run/log/sockrelay.log**: **Failed to create thread: No space left on device.**

This issue is resolved in this release.

- **PR 3058913: vSAN iSCSI firewall rule is lost after network firewall rules are refreshed**

If the network firewall rules are refreshed on the ESXi host, vSAN iSCSI firewall rules might be lost. This problem can impact vSAN iSCSI network connections.

This issue is resolved in this release.

- **PR 3093182: VMware NSX installation and upgrade in a vSphere environment with DPUs on ESXi 8.0 hosts might fail**

The installation and upgrade of VMware NSX by using the NSX Manager in a vSphere environment with DPUs on ESXi 8.0 hosts might fail with a **Configuration State** error and a message similar to:

Host configuration: Failed to send the HostConfig message. [TN=TransportNode/ce1dba65-7fa5-4a47-bf84-e897859fe0db]. Reason: Failed to send HostConfig RPC to MPA TN:ce1dba65-7fa5-4a47-bf84-e897859fe0db. Error: Unable to reach client ce1dba65-7fa5-4a47-bf84-e897859fe0db, application SwitchingVertical. LogicalSwitch full-sync: LogicalSwitch full-sync realization query skipped.

The issue occurs because a DPU scan prevents the ESXi 8.0 host to reboot after installing the NSX image on the DPU.

This issue is resolved in this release. If you do not update to ESXi 8.0b, do the following:

- From the NSX Manager, remove the existing NSX image.
- Upload the ESXi 8.0a offline depot **VMware-ESXi-8.0a-20842819-depot.zip**.
- From the NSX Manager, retry the NSX installation.

- **PR 3060276: After upgrade to ESXi 8.0b, ESXi hosts might fail with a purple diagnostic screen due to legacy I/O scheduler**

Starting from ESXi 6.0, mClock is the default I/O scheduler for ESXi, but some environments might still use legacy schedulers of ESXi versions earlier than 6.0. As a result, upgrades of such hosts to ESXi 7.0 Update 3 and later might fail with a purple diagnostic screen.

This issue is resolved in this release.

• **PR 3092270: VMware Host Client fails to attach an existing virtual disk to a virtual machine**

After upgrading to ESXi 8.0, you might not be able to attach an existing virtual machine disk (VMDK) to a VM by using the VMware Host Client, which is used to connect to and manage single ESXi hosts. The issue does not affect the vSphere Client.

This issue is resolved in this release.

• **PR 3069024: VM with multiple fractional vGPU devices fails to power on**

NVIDIA vGPU 15.0 drivers supports multiple fractional vGPU devices in a single VM, but this does not work on non SR-IOV GPUs such as NVIDIA T4.

This issue is resolved in this release. The fix ensures multiple fractional vGPU devices can work on all supported GPU hardware.

• **PR 3066138: vSAN cluster with Data-in-Transit encryption loses connection to HCI mesh datastore**

Data-in-Transit encryption is not supported on clusters with HCI Mesh configuration. If you deactivate vSAN with DIT encryption, and then reenables vSAN without DIT encryption, the RDT counter for DIT encryption is not decremented correctly. When you enable HCI mesh, the RDT code assumes that traffic needs to be encrypted. The unencrypted traffic is dropped, causing loss of connectivity to the HCI mesh datastore.

This issue is resolved in this release.

esx-update_8.0.0-1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_esx-update_8.0.0-1.20.21203435 - VMware_bootbank_loadesx_8.0.0-1.20.21203435
PRs Fixed	3045268
CVE numbers	N/A

Updates the **loadesx** and **esx-update** VIBs.

esxio-update_8.0.0-1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_esxio-update_8.0.0-1.20.21203435 - VMware_bootbank_loadesxio_8.0.0-1.20.21203435
PRs Fixed	N/A
CVE numbers	N/A

Updates the **esxio-update** and **loadesxio** VIBs.

Mellanox-nmlx5_4.23.0.36-10.2vmw.800.1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMW_bootbank_nmlx5-rdma_4.23.0.36-10.2vmw.800.1.20.21203435 - VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-10.2vmw.800.1.20.21203435 - VMW_bootbank_nmlx5-core-esxio_4.23.0.36-10.2vmw.800.1.20.21203435 - VMW_bootbank_nmlx5-core_4.23.0.36-10.2vmw.800.1.20.21203435
PRs Fixed	3059592, 3054305, 3048127, 3064828, 3060495
CVE numbers	N/A

Updates the **nmlx5-rdma**, **nmlx5-rdma-esxio**, **nmlx5-core-esxio**, and **nmlx5-core** VIBs to resolve the following issue:

- PR 3059592: You might see a reduction in TCP throughput in vSphere environments with Nvidia Bluefield 2 DPUs**

If your underlying hardware allows devices to support a TCP window scale option higher than 7, you might see a reduction in TCP throughput in vSphere environments with Nvidia Bluefield 2 DPUs, because the Nvidia Bluefield 2 scale limit is 7.

This issue is resolved in this release.

VMware-dwi2c_0.1-7vmw.800.1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMW_bootbank_dwi2c-esxio_0.1-7vmw.800.1.20.21203435 - VMW_bootbank_dwi2c_0.1-7vmw.800.1.20.21203435
PRs Fixed	3057364
CVE numbers	N/A

Updates the **dwi2c-esxio** and **dwi2c** VIBs.

Intel-i40en_1.11.2.6-1vmw.800.1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical

Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	VMW_bootbank_i40en_1.11.2.6-1vmw.800.1.20.21203435
PRs Fixed	3046937
CVE numbers	N/A

Updates the **i40en** VIB.

Pensando-ionic-en_20.0.0-30vmw.800.1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMW_bootbank_ionic-en-esxio_20.0.0-30vmw.800.1.20.21203435 - VMW_bootbank_ionic-en_20.0.0-30vmw.800.1.20.21203435
PRs Fixed	3041427
CVE numbers	N/A

Updates the **ionic-en-esxio** and **ionic-en** VIBs to resolve the following issue:

- PR 3039496: Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode**
 Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example **vmnic6** and **vmnic7**) and one management controller (for example **vmnic8**):

```

[~] esxcli nics -l
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando Systems DSC-25
10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando Systems DSC-25
10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
[~] esxcli nics -ls
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando Systems DSC-25
10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller

```

 The high-speed ethernet controllers **vmnic6** and **vmnic7** register first and operate with RSS set to 16 receive queues.

```

[~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }

```

 However, in rare cases, if the management controller **vmnic8** registers first with the vSphere Distributed Switch, the high-speed ethernet controllers **vmnic6** or **vmnic7** uplink might end up operating with RSS set to 1 receive queue.

```

[~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }

```

 As a result, you might see slower performance in native mode.
 This issue is resolved in this release.

Broadcom-ELX-lpfc_14.0.635.4-14vmw.800.1.20.21203435

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	VMW_bootbank_lpfc_14.0.635.4-14vmw.800.1.20.21203435
PRs Fixed	3063528
CVE numbers	N/A

Updates the **lpfc** VIB to resolve the following issue:

- PR 3063528: On a server with tens of devices attached, an ESXi host might fail with a purple diagnostic screen due to exceeding the dump files limit**

On servers with tens of devices attached, dump files for each of the devices might exceed the total limit. As a result, when you reboot an ESXi host or unload the lpfc driver, the host might fail with a purple diagnostic screen and an error **#PF exception vmk_DumpAddFileCallback()** identifying that the dump file resources are used up.

This issue is resolved in this release.

ESXi_8.0.0-1.15.21203431

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_esxio_8.0.0-1.15.21203431 - VMware_bootbank_drivervm-gpu_8.0.0-1.15.21203431 - VMware_bootbank_esxio-base_8.0.0-1.15.21203431 - VMware_bootbank_clusterstore_8.0.0-1.15.21203431 - VMware_bootbank_gc_8.0.0-1.15.21203431 - VMware_bootbank_esx-base_8.0.0-1.15.21203431 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.15.21203431 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.15.21203431 - VMware_bootbank_vdfs_8.0.0-1.15.21203431 - VMware_bootbank_crx_8.0.0-1.15.21203431 - VMware_bootbank_esxio-combiner_8.0.0-1.15.21203431 - VMware_bootbank_bmcad-esxio_8.0.0-1.15.21203431 - VMware_bootbank_native-misc-drivers_8.0.0-1.15.21203431 - VMware_bootbank_esx-xserver_8.0.0-1.15.21203431 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.15.21203431 - VMware_bootbank_bmcad_8.0.0-1.15.21203431 - VMware_bootbank_vsanhealth_8.0.0-1.15.21203431

	- VMware_bootbank_vsan_8.0.0-1.15.21203431 - VMware_bootbank_cpu-microcode_8.0.0-1.15.21203431 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.15.21203431 - VMware_bootbank_trx_8.0.0-1.15.21203431 - VMware_bootbank_gc-esxio_8.0.0-1.15.21203431
PRs Fixed	3049562, 3049613, 3049619, 3049640, 3050030, 3050032, 3050037, 3050512, 3054127, 3068093, 3030692
CVE numbers	CVE-2020-28196

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

Updates the **esx-ui**, **esx-xserver**, **cpu-microcode**, **trx**, **vsanhealth**, **esx-base**, **esx-dvfilter-generic-fastpath**, **gc**, **esxio-combiner**, **native-misc-drivers**, **bmcal**, **vsan**, **vdbs**, and **crx** VIBs to resolve the following issues:

- ESXi 8.0b provides the following security updates:**
 - The SQLite database is updated to version 3.39.2.
 - The Python library is updated to version 3.8.15.
 - OpenSSL is updated to version 3.0.7.
 - The libxml2 library is updated to version 2.10.2.
 - The Expat XML parser is updated to version 2.5.0.
 - The lxml XML toolkit is updated to version 4.5.2.
 - cURL is updated to version 7.84.
 - The Go library is updated to version 1.18.6.
 - The tcpdump package is updated to version 4.99.1.
- The cpu-microcode VIB includes the following Intel microcode:

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	0x0000001d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	0x00000011	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	0x00000007	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	0x0000002f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	0x0000001f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series

Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	0x00000621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	0x0000071a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	0x0000000d	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	0x0000003b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	0x00000021	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	0x00000028	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	0x0000042e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	0x00000715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series

Haswell EP	0x306f2 (06/3f/2)	0x6f	0x00000049	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	0x0000001a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	0x00000022	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	0x0000012d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	0x0b000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	0x02006e05	3/8/2022	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	0x04003302	12/10/2021	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	0x05003302	12/10/2021	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	0x07002501	11/19/2021	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300

Broadwell DE	0x50662 (06/56/2)	0x10	0x0000001c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	0x0700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	0x0f00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	0x0e000014	9/18/2021	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	0x000000f0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	0x00000038	12/2/2021	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	0x0d000375	4/7/2022	Intel Xeon Silver 4300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Platinum 8300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	0x010001f0	6/24/2022	Intel Xeon D Series
Snow Ridge	0x80665 (06/86/5)	0x01	0x4c000020	5/10/2022	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	0x4c000020	5/10/2022	Intel Atom P5000 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	0x000000f0	11/12/2021	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	0x000000f0	11/15/2021	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	0x000000f0	11/12/2021	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	0x000000f0	11/15/2021	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	0x000000f4	7/31/2022	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	0x00000056	8/2/2022	Intel Xeon E-2300 Series

esx-update_8.0.0-1.15.21203431

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	VMware_bootbank_loadesx_8.0.0-1.15.21203431

	VMware_bootbank_esx-update_8.0.0-1.15.21203431
PRs Fixed	N/A
CVE numbers	N/A

Updates the **loadesx** and **esx-update** VIBs.

esxio-update_8.0.0-1.15.21203431

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_esxio-update_8.0.0-1.15.21203431 - VMware_bootbank_loadesxio_8.0.0-1.15.21203431
PRs Fixed	N/A
CVE numbers	N/A

Updates the **esxio-update** and **loadesxio** VIBs.

VMware-VM-Tools_12.1.5.20735119-21203431

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	VMware_locker_tools-light_12.1.5.20735119-21203431
PRs Fixed	3070447
CVE numbers	N/A

Updates the **tools-light** VIB.

- The following VMware Tools ISO images are bundled with ESXi 8.0b:
 - windows.iso**: VMware Tools 12.1.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

- VMware Tools 11.0.6:
 - windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
- VMware Tools 10.0.12:
 - winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
- solaris.iso**: VMware Tools image 10.3.10 for Solaris.
- darwin.iso**: Supports Mac OS X versions 10.11 and later. VMware Tools 12.1.0 was the last regular release for macOS. Refer VMware knowledge base article [88698](#) for details.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [Updating VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [Earlier versions of VMware Tools](#)
- [VMware Tools 12.1.5 Release Notes](#)

ESXi-8.0b-21203435-standard

Profile Name	ESXi-8.0b-21203435-standard
Build	For build information, see Patches Contained in this Release .
Vendor	VMware, Inc.
Release Date	February 14, 2023
Acceptance Level	PartnerSupported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esx-xserver_8.0.0-1.20.21203435 - VMware_bootbank_gc-esxio_8.0.0-1.20.21203435 - VMware_bootbank_vdfs_8.0.0-1.20.21203435 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.20.21203435 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.20.21203435 - VMware_bootbank_crx_8.0.0-1.20.21203435 - VMware_bootbank_trx_8.0.0-1.20.21203435 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.20.21203435 - VMware_bootbank_esxio-base_8.0.0-1.20.21203435 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.20.21203435 - VMware_bootbank_esxio_8.0.0-1.20.21203435 - VMware_bootbank_cpu-microcode_8.0.0-1.20.21203435 - VMware_bootbank_clusterstore_8.0.0-1.20.21203435 - VMware_bootbank_native-misc-drivers_8.0.0-1.20.21203435

	- VMware_bootbank_esxio-combiner_8.0.0-1.20.21203435 - VMware_bootbank_gc_8.0.0-1.20.21203435 - VMware_bootbank_bmcad_8.0.0-1.20.21203435 - VMware_bootbank_esx-base_8.0.0-1.20.21203435 - VMware_bootbank_vsanhealth_8.0.0-1.20.21203435 - VMware_bootbank_vsan_8.0.0-1.20.21203435 - VMware_bootbank_drivervm-gpu_8.0.0-1.20.21203435 - VMware_bootbank_bmcad-esxio_8.0.0-1.20.21203435 - VMware_bootbank_esx-ui_2.10.1-20925449 - VMware_bootbank_esx-update_8.0.0-1.20.21203435 - VMware_bootbank_loadesx_8.0.0-1.20.21203435 - VMware_bootbank_esxio-update_8.0.0-1.20.21203435 - VMware_bootbank_loadesxio_8.0.0-1.20.21203435 - VMW_bootbank_nmlx5-rdma_4.23.0.36-10.2vmw.800.1.20.21203435 - VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-10.2vmw.800.1.20.21203435 - VMW_bootbank_nmlx5-core-esxio_4.23.0.36-10.2vmw.800.1.20.21203435 - VMW_bootbank_nmlx5-core_4.23.0.36-10.2vmw.800.1.20.21203435 - VMW_bootbank_dwi2c-esxio_0.1-7vmw.800.1.20.21203435 - VMW_bootbank_dwi2c_0.1-7vmw.800.1.20.21203435 - VMW_bootbank_i40en_1.11.2.6-1vmw.800.1.20.21203435 - VMW_bootbank_ionic-en-esxio_20.0.0-30vmw.800.1.20.21203435 - VMW_bootbank_ionic-en_20.0.0-30vmw.800.1.20.21203435 - VMW_bootbank_lpf_14.0.635.4-14vmw.800.1.20.21203435 - VMware_locker_tools-light_12.1.5.20735119-21203431
PRs Fixed	3053428, 3051651, 3057027, 3035652, 3065063, 3051004, 3062043, 3070831, 3065990, 3065990, 3041101, 2988813, 3047637, 3044476, 3070865, 3078203, 3055341, 3081041, 3064063, 3064062, 3064054, 3064053, 3062062, 3064096, 3066138, 3065063, 3064848, 3068605, 3064060, 3034879, 3051651, 3081884, 3078334, 3073025, 3070865, 3073230, 3071612, 3070831, 3068028, 3071614, 3071981, 3072190, 3065990, 3081686, 3062907, 3055084, 3069024
Related CVE numbers	N/A

• **This patch updates the following issues:**

- When you select **Network Adapters** in DCUI, you see no output and cannot modify the network configuration. In the `/var/log/syslog.log` file, you see an error such as **2022-xx-xx:02:02.xxxx Er(11) DCUI[265583]: Scanning dvswitches failed: Unable to get the dvs name: Status(bad0001)= Failure**
- Starting with 7.0 Update 2, ESXi supports Posted Interrupts (PI) on Intel CPUs for PCI passthrough devices to improve the overall system performance. In some cases, a race between posted interrupts and the VMkernel scheduling might occur. As a result, virtual machines that are configured with PCI passthrough devices with normal or low latency sensitivity, might experience soft lockups.
- In the vSphere Client and the ESXi Host Client, you might see incorrect number of sockets and respectively cores per socket when you change the NPS setting on an AMD EPYC 9004 server from the default of 1, or Auto, to another value, such as 2 or 4.
- Windows 2012 and later use SCSI-3 reservation for resource arbitration to support Windows failover clustering (WSFC) on ESXi for cluster-across-box (CAB) configurations. However, if you configure the bus sharing of the

SCSI controller on that VM to Physical, the **SCSI RESERVE** command causes the ESXi host to fail with a purple diagnostic screen. **SCSI RESERVE** is SCSI-2 semantic and is not supported with WSFC clusters on ESXi.

- In rare cases, multiple concurrent open and close operations on datastores, for example when vSphere DRS triggers multiple migration tasks, might cause some VMs to take longer to respond or become unresponsive for a short time. The issue is more likely to occur on VMs that have multiple virtual disks over different datastores.
- The logger command of some logger clients might not pass the right facility code to the ESXi syslog daemon and cause incorrect facility or severity values in ESXi syslog log messages.
- vSAN File Service requires hosts to communicate with each other. File Service might incorrectly use an IP address in the witness network for inter-communication. If you have configured an isolated witness network for vSAN, the host can communicate with a witness node over the witness network, but hosts cannot communicate with each other over the witness network. Communication between hosts for vSAN File Service cannot be established.
- A rare race condition caused by memory reclamation for some performance optimization processes in vSphere vMotion might cause ESXi hosts to fail with a purple diagnostic screen.
- In rare cases, on vSphere systems with number of CPUs more than 448, the calculation of memory utilization overflows the 32 bit value which rounds up to a value lower than the expected to be allocated. As a result, ESXi hosts might fail with a purple diagnostic screen upon upgrade or reboot.
This issue is resolved in this release.
- When TXT is enabled on an ESX host, attempts to power-on a VM might fail with an error. In the vSphere Client, you see a message such as **This host supports Intel VT-x, but Intel VT-x is restricted. Intel VT-x might be restricted because 'trusted execution' has been enabled in the BIOS/firmware settings or because the host has not been power-cycled since changing this setting.**
- With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.
- Due to a rare race condition during a Fast Suspend Resume (FSR) operation on a VM with shared pages between VMs, ESXi hosts might fail with a purple diagnostic screen with an error such as **PFrame_IsBackedByLPage** in the backtrace.
- Due to a rare issue with handling AVX2 instructions, a virtual machine of version ESX 8.0 might fail with **ESX unrecoverable** error. In the **vmware.log** file, you see a message such as: **MONITOR PANIC: vcpu-0:VMM fault 6: src=MONITOR** The issue is specific for virtual machines with hardware versions 12 or earlier.
- For ESXi hosts in vSAN clusters of version earlier than vSAN 6.2, upgrades to ESXi 8.0 might cause loss of the VMkernel port tagging for vSAN. As a result, the vSAN networking configuration does not exist after the upgrade.
- In VMware Aria Operations for Logs, formerly vRealize Log Insight, you might see a large volume of logs generated by Storage I/O Control such as **Invalid share value: 0. Using default.** and **Skipping device naa.xxxx either due to VSI read error or abnormal state.** The volume of logs varies depending on the number of ESXi hosts in a cluster and the number of devices in switched off state. When the issue occurs, the log volume generates quickly, within 24 hours, and VMware Aria Operations for Logs might classify the messages as critical. However, such logs are harmless and do not impact the operations on other datastores that are online.
- When you enable the **ToolsRamdisk** advanced option that makes sure the **/vmtools** partition is always on a RAM disk rather than on a USB or SD-card, installing or updating the **tools-light** VIB does not update the content of the

RAM disk automatically. Instead, you must reboot the host so that the RAM disk updates and the new VM Tools version is available for the VMs.

- When you try to export a file with a list of Inventory objects, such as VMs, Hosts, and Datastores, the task fails with an error **Export Data Failure**.
- If the vSAN sockrelay runs out of memory, it can cause spurious errors in the vSAN File Service. The vSAN health service shows numerous file service health check errors. The following entry appears in the `/var/run/log/sockrelay.log`: **Failed to create thread: No space left on device**.
- If the network firewall rules are refreshed on the ESXi host, vSAN iSCSI firewall rules might be lost. This problem can impact vSAN iSCSI network connections.
- The installation and upgrade of VMware NSX by using the NSX Manager in a vSphere environment with DPUs on ESXi 8.0 hosts might fail with a **Configuration State** error and a message similar to:

Host configuration: Failed to send the HostConfig message. [TN=TransportNode/ce1dba65-7fa5-4a47-bf84-e897859fe0db]. Reason: Failed to send HostConfig RPC to MPA TN:ce1dba65-7fa5-4a47-bf84-e897859fe0db. Error: Unable to reach client ce1dba65-7fa5-4a47-bf84-e897859fe0db, application SwitchingVertical. LogicalSwitch full-sync: LogicalSwitch full-sync realization query skipped.

The issue occurs because a DPU scan prevents the ESXi 8.0 host to reboot after installing the NSX image on the DPU.

- Starting from ESXi 6.0, mClock is the default I/O scheduler for ESXi, but some environments might still use legacy schedulers of ESXi versions earlier than 6.0. As a result, upgrades of such hosts to ESXi 7.0 Update 3 and later might fail with a purple diagnostic screen.
- After upgrading to ESXi 8.0, you might not be able to attach an existing virtual machine disk (VMDK) to a VM by using the VMware Host Client, which is used to connect to and manage single ESXi hosts. The issue does not affect the vSphere Client.
- NVIDIA vGPU 15.0 drivers supports multiple fractional vGPU devices in a single VM, but this does not work on non SR-IOV GPUs such as NVIDIA T4.
- Data-in-Transit encryption is not supported on clusters with HCI Mesh configuration. If you deactivate vSAN with DIT encryption, and then reenables vSAN without DIT encryption, the RDT counter for DIT encryption is not decremented correctly. When you enable HCI mesh, the RDT code assumes that traffic needs to be encrypted. The unencrypted traffic is dropped, causing loss of connectivity to the HCI mesh datastore.
- If your underlying hardware allows devices to support a TCP window scale option higher than 7, you might see a reduction in TCP throughput in vSphere environments with Nvidia Bluefield 2 DPUs, because the Nvidia Bluefield 2 scale limit is 7.
- Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example **vmnic6** and **vmnic7**) and one management controller (for example **vmnic8**):

```
~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando Systems
DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando Systems
DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
~] esxcfg-nics -lS
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando Systems
DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers **vmnic6** and **vmnic7** register first and operate with RSS set to 16 receive queues.

```
~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num
of RSS-Q=16, ntxq_descs=2048, nrqx_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1,
geneve_offload=1 }
```

However, in rare cases, if the management controller **vmnic8** registers first with the vSphere Distributed Switch, the high-speed ethernet controllers **vmnic6** or **vmnic7** uplink might end up operating with RSS set to 1 receive queue.:

```
~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...
```

Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }

As a result, you might see slower performance in native mode.

- On servers with tens of devices attached, dump files for each of the devices might exceed the total limit. As a result, when you reboot an ESXi host or unload the lpfc driver, the host might fail with a purple diagnostic screen and an error **#PF exception vmk_DumpAddFileCallback()** identifying that the dump file resources are used up.
- In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.
- If the advanced processor setting **Enable IOMMU in this virtual machine** is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

ESXi-8.0b-21203435-no-tools

Profile Name	ESXi-8.0b-21203435-no-tools
Build	For build information, see Patches Contained in this Release .
Vendor	VMware, Inc.
Release Date	February 14, 2023
Acceptance Level	PartnerSupported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	• VMware_bootbank_esx-xserver_8.0.0-1.20.21203435 • VMware_bootbank_gc-esxio_8.0.0-1.20.21203435 • VMware_bootbank_vdfs_8.0.0-1.20.21203435 • VMware_bootbank_esxio-combiner-esxio_8.0.0-1.20.21203435 • VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.20.21203435 • VMware_bootbank_crx_8.0.0-1.20.21203435 • VMware_bootbank_trx_8.0.0-1.20.21203435 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.20.21203435 • VMware_bootbank_esxio-base_8.0.0-1.20.21203435 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.20.21203435 • VMware_bootbank_esxio_8.0.0-1.20.21203435 • VMware_bootbank_cpu-microcode_8.0.0-1.20.21203435 • VMware_bootbank_clusterstore_8.0.0-1.20.21203435 • VMware_bootbank_native-misc-drivers_8.0.0-1.20.21203435 • VMware_bootbank_esxio-combiner_8.0.0-1.20.21203435 • VMware_bootbank_gc_8.0.0-1.20.21203435 • VMware_bootbank_bmcal_8.0.0-1.20.21203435

	• VMware_bootbank_esx-base_8.0.0-1.20.21203435 • VMware_bootbank_vsanhealth_8.0.0-1.20.21203435 • VMware_bootbank_vsan_8.0.0-1.20.21203435 • VMware_bootbank_drivervm-gpu_8.0.0-1.20.21203435 • VMware_bootbank_bmcad-esxio_8.0.0-1.20.21203435 • VMware_bootbank_esx-ui_2.10.1-20925449 • VMware_bootbank_esx-update_8.0.0-1.20.21203435 • VMware_bootbank_loadesx_8.0.0-1.20.21203435 • VMware_bootbank_esxio-update_8.0.0-1.20.21203435 • VMware_bootbank_loadesxio_8.0.0-1.20.21203435 • VMW_bootbank_nmlx5-rdma_4.23.0.36-10.2vmw.800.1.20.21203435 • VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-10.2vmw.800.1.20.21203435 • VMW_bootbank_nmlx5-core-esxio_4.23.0.36-10.2vmw.800.1.20.21203435 • VMW_bootbank_nmlx5-core_4.23.0.36-10.2vmw.800.1.20.21203435 • VMW_bootbank_dwi2c-esxio_0.1-7vmw.800.1.20.21203435 • VMW_bootbank_dwi2c_0.1-7vmw.800.1.20.21203435 • VMW_bootbank_i40en_1.11.2.6-1vmw.800.1.20.21203435 • VMW_bootbank_ionic-en-esxio_20.0.0-30vmw.800.1.20.21203435 • VMW_bootbank_ionic-en_20.0.0-30vmw.800.1.20.21203435 • VMW_bootbank_lpf_14.0.635.4-14vmw.800.1.20.21203435
PRs Fixed	3053428, 3051651, 3057027, 3035652, 3065063, 3051004, 3062043, 3070831, 3065990, 3065990, 3041101, 2988813, 3047637, 3044476, 3070865, 3078203, 3055341, 3081041, 3064063, 3064062, 3064054, 3064053, 3062062, 3064096, 3066138, 3065063, 3064848, 3068605, 3064060, 3034879, 3051651, 3081884, 3078334, 3073025, 3070865, 3073230, 3071612, 3070831, 3068028, 3071614, 3071981, 3072190, 3065990, 3081686, 3062907, 3055084, 3069024, 3061447, 3065694, 3045268, 3059592, 3054305, 3048127, 3041427, 3063528
Related CVE numbers	N/A

- **This patch updates the following issues:**
 - When you select **Network Adapters** in DCUI, you see no output and cannot modify the network configuration. In the `/var/log/syslog.log` file, you see an error such as **2022-xx-xx:02:02.xxxx Er(11) DCUI[265583]: Scanning dvswitches failed: Unable to get the dvs name: Status(bad0001)= Failure**
 - Starting with 7.0 Update 2, ESXi supports Posted Interrupts (PI) on Intel CPUs for PCI passthrough devices to improve the overall system performance. In some cases, a race between posted interrupts and the VMkernel scheduling might occur. As a result, virtual machines that are configured with PCI passthrough devices with normal or low latency sensitivity, might experience soft lockups.
 - In the vSphere Client and the ESXi Host Client, you might see incorrect number of sockets and respectively cores per socket when you change the NPS setting on an AMD EPYC 9004 server from the default of 1, or Auto, to another value, such as 2 or 4.
 - Windows 2012 and later use SCSI-3 reservation for resource arbitration to support Windows failover clustering (WSFC) on ESXi for cluster-across-box (CAB) configurations. However, if you configure the bus sharing of the

SCSI controller on that VM to Physical, the **SCSI RESERVE** command causes the ESXi host to fail with a purple diagnostic screen. **SCSI RESERVE** is SCSI-2 semantic and is not supported with WSFC clusters on ESXi.

- In rare cases, multiple concurrent open and close operations on datastores, for example when vSphere DRS triggers multiple migration tasks, might cause some VMs to take longer to respond or become unresponsive for a short time. The issue is more likely to occur on VMs that have multiple virtual disks over different datastores.
- The logger command of some logger clients might not pass the right facility code to the ESXi syslog daemon and cause incorrect facility or severity values in ESXi syslog log messages.
- vSAN File Service requires hosts to communicate with each other. File Service might incorrectly use an IP address in the witness network for inter-communication. If you have configured an isolated witness network for vSAN, the host can communicate with a witness node over the witness network, but hosts cannot communicate with each other over the witness network. Communication between hosts for vSAN File Service cannot be established.
- A rare race condition caused by memory reclamation for some performance optimization processes in vSphere vMotion might cause ESXi hosts to fail with a purple diagnostic screen.
- In rare cases, on vSphere systems with number of CPUs more than 448, the calculation of memory utilization overflows the 32 bit value which rounds up to a value lower than the expected to be allocated. As a result, ESXi hosts might fail with a purple diagnostic screen upon upgrade or reboot.
This issue is resolved in this release.
- When TXT is enabled on an ESX host, attempts to power-on a VM might fail with an error. In the vSphere Client, you see a message such as **This host supports Intel VT-x, but Intel VT-x is restricted. Intel VT-x might be restricted because 'trusted execution' has been enabled in the BIOS/firmware settings or because the host has not been power-cycled since changing this setting.**
- With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.
- Due to a rare race condition during a Fast Suspend Resume (FSR) operation on a VM with shared pages between VMs, ESXi hosts might fail with a purple diagnostic screen with an error such as **PFrame_IsBackedByLPage** in the backtrace.
- Due to a rare issue with handling AVX2 instructions, a virtual machine of version ESX 8.0 might fail with **ESX unrecoverable** error. In the **vmware.log** file, you see a message such as: **MONITOR PANIC: vcpu-0:VMM fault 6: src=MONITOR** The issue is specific for virtual machines with hardware versions 12 or earlier.
- For ESXi hosts in vSAN clusters of version earlier than vSAN 6.2, upgrades to ESXi 8.0 might cause loss of the VMkernel port tagging for vSAN. As a result, the vSAN networking configuration does not exist after the upgrade.
- In VMware Aria Operations for Logs, formerly vRealize Log Insight, you might see a large volume of logs generated by Storage I/O Control such as **Invalid share value: 0. Using default.** and **Skipping device naa.xxxx either due to VSI read error or abnormal state.** The volume of logs varies depending on the number of ESXi hosts in a cluster and the number of devices in switched off state. When the issue occurs, the log volume generates quickly, within 24 hours, and VMware Aria Operations for Logs might classify the messages as critical. However, such logs are harmless and do not impact the operations on other datastores that are online.
- When you enable the **ToolsRamdisk** advanced option that makes sure the **/vmtools** partition is always on a RAM disk rather than on a USB or SD-card, installing or updating the **tools-light** VIB does not update the content of the

RAM disk automatically. Instead, you must reboot the host so that the RAM disk updates and the new VM Tools version is available for the VMs.

- When you try to export a file with a list of Inventory objects, such as VMs, Hosts, and Datastores, the task fails with an error **Export Data Failure**.
- If the vSAN sockrelay runs out of memory, it can cause spurious errors in the vSAN File Service. The vSAN health service shows numerous file service health check errors. The following entry appears in the `/var/run/log/sockrelay.log`: **Failed to create thread: No space left on device**.
- If the network firewall rules are refreshed on the ESXi host, vSAN iSCSI firewall rules might be lost. This problem can impact vSAN iSCSI network connections.
- The installation and upgrade of VMware NSX by using the NSX Manager in a vSphere environment with DPUs on ESXi 8.0 hosts might fail with a **Configuration State** error and a message similar to:

Host configuration: Failed to send the HostConfig message. [TN=TransportNode/ce1dba65-7fa5-4a47-bf84-e897859fe0db]. Reason: Failed to send HostConfig RPC to MPA TN:ce1dba65-7fa5-4a47-bf84-e897859fe0db. Error: Unable to reach client ce1dba65-7fa5-4a47-bf84-e897859fe0db, application SwitchingVertical. LogicalSwitch full-sync: LogicalSwitch full-sync realization query skipped.

The issue occurs because a DPU scan prevents the ESXi 8.0 host to reboot after installing the NSX image on the DPU.

- Starting from ESXi 6.0, mClock is the default I/O scheduler for ESXi, but some environments might still use legacy schedulers of ESXi versions earlier than 6.0. As a result, upgrades of such hosts to ESXi 7.0 Update 3 and later might fail with a purple diagnostic screen.
- After upgrading to ESXi 8.0, you might not be able to attach an existing virtual machine disk (VMDK) to a VM by using the VMware Host Client, which is used to connect to and manage single ESXi hosts. The issue does not affect the vSphere Client.
- NVIDIA vGPU 15.0 drivers supports multiple fractional vGPU devices in a single VM, but this does not work on non SR-IOV GPUs such as NVIDIA T4.
- Data-in-Transit encryption is not supported on clusters with HCI Mesh configuration. If you deactivate vSAN with DIT encryption, and then reenables vSAN without DIT encryption, the RDT counter for DIT encryption is not decremented correctly. When you enable HCI mesh, the RDT code assumes that traffic needs to be encrypted. The unencrypted traffic is dropped, causing loss of connectivity to the HCI mesh datastore.
- If your underlying hardware allows devices to support a TCP window scale option higher than 7, you might see a reduction in TCP throughput in vSphere environments with Nvidia Bluefield 2 DPUs, because the Nvidia Bluefield 2 scale limit is 7.
- Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example **vmnic6** and **vmnic7**) and one management controller (for example **vmnic8**):

```
~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando Systems
DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando Systems
DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
~] esxcfg-nics -lS
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando Systems
DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers **vmnic6** and **vmnic7** register first and operate with RSS set to 16 receive queues.

```
~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num
of RSS-Q=16, ntxq_descs=2048, nrqx_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1,
geneve_offload=1 }
```

However, in rare cases, if the management controller **vmnic8** registers first with the vSphere Distributed Switch, the high-speed ethernet controllers **vmnic6** or **vmnic7** uplink might end up operating with RSS set to 1 receive queue.:

```
~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...
```

Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }

As a result, you might see slower performance in native mode.

- On servers with tens of devices attached, dump files for each of the devices might exceed the total limit. As a result, when you reboot an ESXi host or unload the lpfc driver, the host might fail with a purple diagnostic screen and an error **#PF exception vmk_DumpAddFileCallback()** identifying that the dump file resources are used up.
- In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.
- If the advanced processor setting **Enable IOMMU in this virtual machine** is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

ESXi-8.0sb-21203431-standard

Profile Name	ESXi-8.0sb-21203431-standard
Build	For build information, see Patches Contained in this Release .
Vendor	VMware, Inc.
Release Date	February 14, 2023
Acceptance Level	PartnerSupported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	• VMware_bootbank_esxio_8.0.0-1.15.21203431 • VMware_bootbank_drivervm-gpu_8.0.0-1.15.21203431 • VMware_bootbank_esxio-base_8.0.0-1.15.21203431 • VMware_bootbank_clusterstore_8.0.0-1.15.21203431 • VMware_bootbank_gc_8.0.0-1.15.21203431 • VMware_bootbank_esx-base_8.0.0-1.15.21203431 • VMware_bootbank_esxio-combiner-esxio_8.0.0-1.15.21203431 • VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.15.21203431 • VMware_bootbank_vdfs_8.0.0-1.15.21203431 • VMware_bootbank_crx_8.0.0-1.15.21203431 • VMware_bootbank_esxio-combiner_8.0.0-1.15.21203431 • VMware_bootbank_bmcal-esxio_8.0.0-1.15.21203431 • VMware_bootbank_native-misc-drivers_8.0.0-1.15.21203431 • VMware_bootbank_esx-xserver_8.0.0-1.15.21203431 • VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.15.21203431 • VMware_bootbank_bmcal_8.0.0-1.15.21203431 • VMware_bootbank_vsanhealth_8.0.0-1.15.21203431 • VMware_bootbank_vsan_8.0.0-1.15.21203431

	• VMware_bootbank_cpu-microcode_8.0.0-1.15.21203431 • VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.15.21203431 • VMware_bootbank_trx_8.0.0-1.15.21203431 • VMware_bootbank_gc-esxio_8.0.0-1.15.21203431 • VMware_bootbank_loadesx_8.0.0-1.15.21203431 • VMware_bootbank_esx-update_8.0.0-1.15.21203431 • VMware_bootbank_esxio-update_8.0.0-1.15.21203431 • VMware_bootbank_loadesxio_8.0.0-1.15.21203431 • VMware_locker_tools-light_12.1.5.20735119-21203431
PRs Fixed	3049562, 3049613, 3049619, 3049640, 3050030, 3050032, 3050037, 3050512, 3054127, 3068093, 3030692, 3070447
Related CVE numbers	CVE-2020-28196

This patch updates the following issues:

- The following VMware Tools ISO images are bundled with ESXi 8.0b:
 - **windows.iso**: VMware Tools 12.1.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - **linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.

The following VMware Tools ISO images are available for download:

 - VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
 - VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
 - **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
 - **darwin.iso**: Supports Mac OS X versions 10.11 and later. VMware Tools 12.1.0 was the last regular release for macOS. Refer VMware knowledge base article [88698](#) for details.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

 - [Updating VMware Tools](#)
 - [VMware Tools for hosts provisioned with Auto Deploy](#)
 - [What Every vSphere Admin Must Know About VMware Tools](#)
 - [Earlier versions of VMware Tools](#)
 - [VMware Tools 12.1.5 Release Notes](#)
 - **ESXi 8.0b provides the following security updates:**
 - The SQLite database is updated to version 3.39.2.
 - The Python library is updated to version 3.8.15.
 - OpenSSL is updated to version 3.0.7.
 - The libxml2 library is updated to version 2.10.2.
 - The Expat XML parser is updated to version 2.5.0.
 - The lxml XML toolkit is updated to version 4.5.2.
 - cURL is updated to version 7.84.
 - The Go library is updated to version 1.18.6.
 - The tcpdump package is updated to version 4.99.1.
 - The cpu-microcode VIB includes the following Intel microcode:

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	0x0000001d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	0x00000011	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	0x00000007	4/23/2018	Intel Core i7-620LE Processor
Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	0x0000002f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	0x0000001f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	0x00000621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	0x0000071a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	0x0000000d	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	0x0000003b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	0x00000021	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C

Haswell DT	0x306c3 (06/3c/3)	0x32	0x00000028	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series
Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	0x0000042e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	0x00000715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	0x00000049	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	0x0000001a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	0x00000022	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	0x0000012d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	0x0b000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	0x02006e05	3/8/2022	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series

Cascade Lake B-0	0x50656 (06/55/6)	0xbf	0x04003302	12/10/2021	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cascade Lake	0x50657 (06/55/7)	0xbf	0x05003302	12/10/2021	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	0x07002501	11/19/2021	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	0x0000001c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	0x0700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	0x0f00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	0x0e000014	9/18/2021	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	0x000000f0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	0x00000038	12/2/2021	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	0x0d000375	4/7/2022	Intel Xeon Silver 4300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Platinum 8300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	0x010001f0	6/24/2022	Intel Xeon D Series
Snow Ridge	0x80665 (06/86/5)	0x01	0x4c000020	5/10/2022	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	0x4c000020	5/10/2022	Intel Atom P5000 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	0x000000f0	11/12/2021	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	0x000000f0	11/15/2021	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	0x000000f0	11/12/2021	Intel Xeon E-2100 Series

Coffee Lake	0x906ec (06/9e/c)	0x22	0x000000f0	11/15/2021	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	0x000000f4	7/31/2022	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	0x00000056	8/2/2022	Intel Xeon E-2300 Series

ESXi-8.0sb-21203431-no-tools

Profile Name	ESXi-8.0sb-21203431-no-tools
Build	For build information, see Patches Contained in this Release .
Vendor	VMware, Inc.
Release Date	February 14, 2023
Acceptance Level	PartnerSupported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_esxio_8.0.0-1.15.21203431 - VMware_bootbank_drivervm-gpu_8.0.0-1.15.21203431 - VMware_bootbank_esxio-base_8.0.0-1.15.21203431 - VMware_bootbank_clusterstore_8.0.0-1.15.21203431 - VMware_bootbank_gc_8.0.0-1.15.21203431 - VMware_bootbank_esx-base_8.0.0-1.15.21203431 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.15.21203431 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.15.21203431 - VMware_bootbank_vdfs_8.0.0-1.15.21203431 - VMware_bootbank_crx_8.0.0-1.15.21203431 - VMware_bootbank_esxio-combiner_8.0.0-1.15.21203431 - VMware_bootbank_bmcad-esxio_8.0.0-1.15.21203431 - VMware_bootbank_native-misc-drivers_8.0.0-1.15.21203431 - VMware_bootbank_esx-xserver_8.0.0-1.15.21203431 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.15.21203431 - VMware_bootbank_bmcad_8.0.0-1.15.21203431 - VMware_bootbank_vsanhealth_8.0.0-1.15.21203431 - VMware_bootbank_vsan_8.0.0-1.15.21203431 - VMware_bootbank_cpu-microcode_8.0.0-1.15.21203431 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.15.21203431 - VMware_bootbank_trx_8.0.0-1.15.21203431 - VMware_bootbank_gc-esxio_8.0.0-1.15.21203431 - VMware_bootbank_loadesx_8.0.0-1.15.21203431 - VMware_bootbank_esx-update_8.0.0-1.15.21203431 - VMware_bootbank_esxio-update_8.0.0-1.15.21203431 - VMware_bootbank_loadesxio_8.0.0-1.15.21203431
PRs Fixed	3049562, 3049613, 3049619, 3049640, 3050030, 3050032, 3050037, 3050512, 3054127, 3068093, 3030692, 3070447
Related CVE numbers	CVE-2020-28196

This patch updates the following issues:

- The following VMware Tools ISO images are bundled with ESXi 8.0b:
 - **windows.iso**: VMware Tools 12.1.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
 - **linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.
 - The following VMware Tools ISO images are available for download:
 - VMware Tools 11.0.6:
 - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
 - VMware Tools 10.0.12:
 - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
 - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
 - **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
 - **darwin.iso**: Supports Mac OS X versions 10.11 and later. VMware Tools 12.1.0 was the last regular release for macOS. Refer VMware knowledge base article [88698](#) for details.

Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

- [Updating VMware Tools](#)
- [VMware Tools for hosts provisioned with Auto Deploy](#)
- [What Every vSphere Admin Must Know About VMware Tools](#)
- [Earlier versions of VMware Tools](#)
- [VMware Tools 12.1.5 Release Notes](#)
- **ESXi 8.0b provides the following security updates:**
 - The SQLite database is updated to version 3.39.2.
 - The Python library is updated to version 3.8.15.
 - OpenSSL is updated to version 3.0.7.
 - The libxml2 library is updated to version 2.10.2.
 - The Expat XML parser is updated to version 2.5.0.
 - The lxml XML toolkit is updated to version 4.5.2.
 - cURL is updated to version 7.84.
 - The Go library is updated to version 1.18.6.
 - The tcpdump package is updated to version 4.99.1.
- The cpu-microcode VIB includes the following Intel microcode:

Code Name	FMS	Pit ID	MCU Rev	MCU Date	Brand Names
Nehalem EP	0x106a5 (06/1a/5)	0x03	0x0000001d	5/11/2018	Intel Xeon 35xx Series; Intel Xeon 55xx Series
Clarkdale	0x20652 (06/25/2)	0x12	0x00000011	5/8/2018	Intel i3/i5 Clarkdale Series; Intel Xeon 34xx Clarkdale Series
Arrandale	0x20655 (06/25/5)	0x92	0x00000007	4/23/2018	Intel Core i7-620LE Processor

Sandy Bridge DT	0x206a7 (06/2a/7)	0x12	0x0000002f	2/17/2019	Intel Xeon E3-1100 Series; Intel Xeon E3-1200 Series; Intel i7-2655-LE Series; Intel i3-2100 Series
Westmere EP	0x206c2 (06/2c/2)	0x03	0x0000001f	5/8/2018	Intel Xeon 56xx Series; Intel Xeon 36xx Series
Sandy Bridge EP	0x206d6 (06/2d/6)	0x6d	0x00000621	3/4/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Sandy Bridge EP	0x206d7 (06/2d/7)	0x6d	0x0000071a	3/24/2020	Intel Pentium 1400 Series; Intel Xeon E5-1400 Series; Intel Xeon E5-1600 Series; Intel Xeon E5-2400 Series; Intel Xeon E5-2600 Series; Intel Xeon E5-4600 Series
Nehalem EX	0x206e6 (06/2e/6)	0x04	0x0000000d	5/15/2018	Intel Xeon 65xx Series; Intel Xeon 75xx Series
Westmere EX	0x206f2 (06/2f/2)	0x05	0x0000003b	5/16/2018	Intel Xeon E7-8800 Series; Intel Xeon E7-4800 Series; Intel Xeon E7-2800 Series
Ivy Bridge DT	0x306a9 (06/3a/9)	0x12	0x00000021	2/13/2019	Intel i3-3200 Series; Intel i7-3500-LE/UE; Intel i7-3600-QE; Intel Xeon E3-1200-v2 Series; Intel Xeon E3-1100-C-v2 Series; Intel Pentium B925C
Haswell DT	0x306c3 (06/3c/3)	0x32	0x00000028	11/12/2019	Intel Xeon E3-1200-v3 Series; Intel i7-4700-EQ Series; Intel i5-4500-TE Series; Intel i3-4300 Series

Ivy Bridge EP	0x306e4 (06/3e/4)	0xed	0x0000042e	3/14/2019	Intel Xeon E5-4600-v2 Series; Intel Xeon E5-2600-v2 Series; Intel Xeon E5-2400-v2 Series; Intel Xeon E5-1600-v2 Series; Intel Xeon E5-1400-v2 Series
Ivy Bridge EX	0x306e7 (06/3e/7)	0xed	0x00000715	3/14/2019	Intel Xeon E7-8800/4800/2800-v2 Series
Haswell EP	0x306f2 (06/3f/2)	0x6f	0x00000049	8/11/2021	Intel Xeon E5-4600-v3 Series; Intel Xeon E5-2600-v3 Series; Intel Xeon E5-2400-v3 Series; Intel Xeon E5-1600-v3 Series; Intel Xeon E5-1400-v3 Series
Haswell EX	0x306f4 (06/3f/4)	0x80	0x0000001a	5/24/2021	Intel Xeon E7-8800/4800-v3 Series
Broadwell H	0x40671 (06/47/1)	0x22	0x00000022	11/12/2019	Intel Core i7-5700EQ; Intel Xeon E3-1200-v4 Series
Avoton	0x406d8 (06/4d/8)	0x01	0x0000012d	9/16/2019	Intel Atom C2300 Series; Intel Atom C2500 Series; Intel Atom C2700 Series
Broadwell EP/EX	0x406f1 (06/4f/1)	0xef	0x0b000040	5/19/2021	Intel Xeon E7-8800/4800-v4 Series; Intel Xeon E5-4600-v4 Series; Intel Xeon E5-2600-v4 Series; Intel Xeon E5-1600-v4 Series
Skylake SP	0x50654 (06/55/4)	0xb7	0x02006e05	3/8/2022	Intel Xeon Platinum 8100 Series; Intel Xeon Gold 6100/5100, Silver 4100, Bronze 3100 Series; Intel Xeon D-2100 Series; Intel Xeon D-1600 Series; Intel Xeon W-3100 Series; Intel Xeon W-2100 Series
Cascade Lake B-0	0x50656 (06/55/6)	0xbf	0x04003302	12/10/2021	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200

Cascade Lake	0x50657 (06/55/7)	0xbf	0x05003302	12/10/2021	Intel Xeon Platinum 9200/8200 Series; Intel Xeon Gold 6200/5200; Intel Xeon Silver 4200/ Bronze 3200; Intel Xeon W-3200
Cooper Lake	0x5065b (06/55/b)	0xbf	0x07002501	11/19/2021	Intel Xeon Platinum 8300 Series; Intel Xeon Gold 6300/5300
Broadwell DE	0x50662 (06/56/2)	0x10	0x0000001c	6/17/2019	Intel Xeon D-1500 Series
Broadwell DE	0x50663 (06/56/3)	0x10	0x0700001c	6/12/2021	Intel Xeon D-1500 Series
Broadwell DE	0x50664 (06/56/4)	0x10	0x0f00001a	6/12/2021	Intel Xeon D-1500 Series
Broadwell NS	0x50665 (06/56/5)	0x10	0x0e000014	9/18/2021	Intel Xeon D-1600 Series
Skylake H/S	0x506e3 (06/5e/3)	0x36	0x000000f0	11/12/2021	Intel Xeon E3-1500-v5 Series; Intel Xeon E3-1200-v5 Series
Denverton	0x506f1 (06/5f/1)	0x01	0x00000038	12/2/2021	Intel Atom C3000 Series
Ice Lake SP	0x606a6 (06/6a/6)	0x87	0x0d000375	4/7/2022	Intel Xeon Silver 4300 Series; Intel Xeon Gold 6300/5300 Series; Intel Xeon Platinum 8300 Series
Ice Lake D	0x606c1 (06/6c/1)	0x10	0x010001f0	6/24/2022	Intel Xeon D Series
Snow Ridge	0x80665 (06/86/5)	0x01	0x4c000020	5/10/2022	Intel Atom P5000 Series
Snow Ridge	0x80667 (06/86/7)	0x01	0x4c000020	5/10/2022	Intel Atom P5000 Series
Kaby Lake H/S/X	0x906e9 (06/9e/9)	0x2a	0x000000f0	11/12/2021	Intel Xeon E3-1200-v6 Series; Intel Xeon E3-1500-v6 Series
Coffee Lake	0x906ea (06/9e/a)	0x22	0x000000f0	11/15/2021	Intel Xeon E-2100 Series; Intel Xeon E-2200 Series (4 or 6 core)
Coffee Lake	0x906eb (06/9e/b)	0x02	0x000000f0	11/12/2021	Intel Xeon E-2100 Series
Coffee Lake	0x906ec (06/9e/c)	0x22	0x000000f0	11/15/2021	Intel Xeon E-2100 Series
Coffee Lake Refresh	0x906ed (06/9e/d)	0x22	0x000000f4	7/31/2022	Intel Xeon E-2200 Series (8 core)
Rocket Lake S	0xa0671 (06/a7/1)	0x02	0x00000056	8/2/2022	Intel Xeon E-2300 Series

ESXi 8.0.0-1.20.21203435

Name	ESXi
Version	ESXi_8.0.0-1.20.21203435
Release Date	February 14, 2023
Category	Bugfix
Affected Components	• ESXi Component • ESXi Install/Upgrade Component • Broadcom Emulex Connectivity Division Ipfc driver for FC adapters • Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers for VMware ESXi • VMware DesignWare I2C Driver • Network driver for Intel(R) X710/XL710/XXV710/X722 Adapters • Pensando Systems Native Ethernet Driver
PRs Fixed	3003171, 3023438
Related CVE numbers	N/A

ESXi 8.0.0-1.15.21203431

Name	ESXi
Version	ESXi_8.0.0-1.15.21203431
Release Date	February 14, 2023
Category	Bugfix
Affected Components	• ESXi Component - core ESXi VIBs • ESXi Install/Upgrade Component • ESXi Tools Component
PRs Fixed	3003171, 3023438
Related CVE numbers	CVE-2020-28196

Known Issues**Miscellaneous Issues**

In a vSphere environment with DPUs, if ESXi selects a DPU NIC during installation by using a x86 kickstart file, the installation fails

In a vSphere environment with DPUs, when you use a kickstart file to automate ESXi installation, ESXi might select a DPU NIC for the default **vSwitch0** during the installation for the management network, which leads to networking failures. For example, fetching a network IP address for the Management Network portgroup might fail. As a result, the ESXi installation also fails.

Workaround: Remove the line **network --bootproto=dhcp** from the kickstart file.

Known Issues from Previous Releases**Installation, Upgrade, and Migration Issues**

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is deactivated, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings > <specific broadcom NIC> > Device Configuration Menu > Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

Miscellaneous Issues

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reinstall ESXi. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In the vSphere API Explorer, VMware Datacenter CLI (DCLI) and PowerCLI, you see an API option "contentinternal" that is not functional

You see an API option **contentinternal** in the metadata of either the vSphere API Explorer, DCLI and PowerCLI. For example, when you open `https://<your vCenter IP>/ui/app/devcenter/api-explorer`, you see the option in the **select API** drop-down menu. This option is not functional.

Workaround: Ignore the **contentinternal** API option and do not use it.

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client.

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled`.

Workaround: None.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example vmnic6 and vmnic7) and one management controller (for example vmnic8):

```
:~] esxcfg-nics -l

vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller

vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller

:~] esxcfg-nics -ls

vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers vmnic6 and vmnic7 register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller vmnic8 registers first with the vSphere Distributed Switch, the high-speed ethernet controllers vmnic6 or vmnic7 uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the ionic_en driver on ESXi by using the following commands:

```
:~] esxcfg-module -u ionic_en
:~] esxcfg-module ionic_en
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind.
```

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough (SR-IOV or UPT) to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

Networking Issues

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

You see link flapping on NICs that use the `ntg3` driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Storage Issues

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vCenter Server and vSphere Client Issues

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade.`

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [90259](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files.`

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_XXX. Extraction of image from host XXX.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Guest OS Issues

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of

communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

VMware ESXi 8.0a Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware ESXi 8.0a | 08 DEC 2022 | Build 20842819
Check for additions and updates to these release notes.

What's New

This patch addresses TPM-related security enhancements. These enhancements are necessary only for the upcoming CPU support in newer Lenovo and HPE platforms.

Patches Contained in This Release

Download Filename:	VMware-ESXi-8.0a-20842819-depot
Build:	20842819
Download Size:	603.9 MB
sha256checksum:	7d67771e856dab1f64c966748322576ddfd3c5a5bd293265367851a21e67f8ee
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi	ESXi_8.0.0-1.10.20842819	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.0-1.10.20842819	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.0-1.10.20842819	Bugfix	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes since the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi80a-20842819	Bugfix	Critical

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0a-20842819-standard
ESXi-8.0a-20842819-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi80a-20842819	12/08/2022	Enhancement	Bugfix image

For information about the individual components and bulletins, see the [Build Details](#) and the [Resolved Issues](#) section.

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

Patch Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

Resolved Issues

ESXi_8.0.0-1.10.20842819

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_vsan_8.0.0-1.10.20842819 - VMware_bootbank_esx-xserver_8.0.0-1.10.20842819 - VMware_bootbank_crx_8.0.0-1.10.20842819 - VMware_bootbank_clusterstore_8.0.0-1.10.20842819 - VMware_bootbank_gc-esxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-base_8.0.0-1.10.20842819 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.10.20842819 - VMware_bootbank_vdfs_8.0.0-1.10.20842819 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.10.20842819 - VMware_bootbank_drivervm-gpu_8.0.0-1.10.20842819

	- VMware_bootbank_esxio-combiner_8.0.0-1.10.20842819 - VMware_bootbank_trx_8.0.0-1.10.20842819 - VMware_bootbank_bmc-al-esxio_8.0.0-1.10.20842819 - VMware_bootbank_bmc-al_8.0.0-1.10.20842819 - VMware_bootbank_gc_8.0.0-1.10.20842819 - VMware_bootbank_esxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.10.20842819 - VMware_bootbank_native-misc-drivers_8.0.0-1.10.20842819 - VMware_bootbank_cpu-microcode_8.0.0-1.10.20842819 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.10.20842819 - VMware_bootbank_esx-base_8.0.0-1.10.20842819 - VMware_bootbank_vsanhealth_8.0.0-1.10.20842819
PRs Fixed	3041101
CVE numbers	N/A

The ESXi and esx-update bulletins are dependent on each other. Always include both in a single ESXi host patch baseline or include the rollup bulletin in the baseline to avoid failure during host patching.

Updates the **esx-ui**, **esx-xserver**, **cpu-microcode**, **trx**, **vsanhealth**, **esx-base**, **esx-dvfilter-generic-fastpath**, **esxio-dvfilter-generic-fastpath**, **gc**, **esxio-combiner**, **native-misc-drivers**, **native-misc-drivers-esxio**, **bmc-al**, **vsan**, **vdfs**, **bmc-al-esxio**, **clusterstore**, **gc-esxio**, **esxio-base**, **drivervm-gpu**, and **crx** VIBs to resolve the following issue:

When TPM 2.0 is enabled with TXT on an ESXi host, attempts to power-on a virtual machine might fail

When TXT is enabled on an ESX host, attempts to power-on a VM might fail with an error. In the vSphere Client, you see a message such as **This host supports Intel VT-x, but Intel VT-x is restricted. Intel VT-x might be restricted because 'trusted execution' has been enabled in the BIOS/firmware settings or because the host has not been power-cycled since changing this setting.**

This issue is resolved in this release.

esx-update_8.0.0-1.10.20842819

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_loadesx_8.0.0-1.10.20842819 - VMware_bootbank_esx-update_8.0.0-1.10.20842819
PRs Fixed	N/A
CVE numbers	N/A

Updates the **loadesx** and **esx-update** VIBs.

esxio-update_8.0.0-1.10.20842819

Patch Category	Bugfix
-----------------------	--------

Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	- VMware_bootbank_loadesxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-update_8.0.0-1.10.20842819
PRs Fixed	N/A
CVE numbers	N/A

Updates the **loadesxio** and **esxio-update** VIBs.

ESXi-8.0a-20842819-standard

Profile Name	ESXi-8.0a-20842819-standard
Build	For build information, see Build Details .
Vendor	VMware, Inc.
Release Date	December 8, 2022
Acceptance Level	PartnerSupported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_vsan_8.0.0-1.10.20842819 - VMware_bootbank_esx-xserver_8.0.0-1.10.20842819 - VMware_bootbank_crx_8.0.0-1.10.20842819 - VMware_bootbank_clusterstore_8.0.0-1.10.20842819 - VMware_bootbank_gc-esxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-base_8.0.0-1.10.20842819 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.10.20842819 - VMware_bootbank_vdfs_8.0.0-1.10.20842819 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.10.20842819 - VMware_bootbank_drivervm-gpu_8.0.0-1.10.20842819 - VMware_bootbank_esxio-combiner_8.0.0-1.10.20842819 - VMware_bootbank_trx_8.0.0-1.10.20842819 - VMware_bootbank_bmcal-esxio_8.0.0-1.10.20842819 - VMware_bootbank_bmcal_8.0.0-1.10.20842819 - VMware_bootbank_gc_8.0.0-1.10.20842819 - VMware_bootbank_esxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.10.20842819 - VMware_bootbank_native-misc-drivers_8.0.0-1.10.20842819 - VMware_bootbank_cpu-microcode_8.0.0-1.10.20842819 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.10.20842819 - VMware_bootbank_esx-base_8.0.0-1.10.20842819 - VMware_bootbank_vsanhealth_8.0.0-1.10.20842819 - VMware_bootbank_loadesx_8.0.0-1.10.20842819

	- VMware_bootbank_esx-update_8.0.0-1.10.20842819 - VMware_bootbank_loadesxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-update_8.0.0-1.10.20842819
PRs Fixed	3041101
Related CVE numbers	N/A

This patch updates the following issue:

• **When TPM 2.0 is enabled with TXT on an ESXi host, attempts to power-on a virtual machine might fail**

When TXT is enabled on an ESX host, attempts to power-on a VM might fail with an error. In the vSphere Client, you see a message such as **This host supports Intel VT-x, but Intel VT-x is restricted. Intel VT-x might be restricted because 'trusted execution' has been enabled in the BIOS/firmware settings or because the host has not been power-cycled since changing this setting.**

This issue is resolved in this release.

ESXi-8.0a-20842819-no-tools

Profile Name	ESXi-8.0a-20842819-no-tools
Build	For build information, see Build Details .
Vendor	VMware, Inc.
Release Date	December 8, 2022
Acceptance Level	PartnerSupported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	- VMware_bootbank_vsan_8.0.0-1.10.20842819 - VMware_bootbank_esx-xserver_8.0.0-1.10.20842819 - VMware_bootbank_crx_8.0.0-1.10.20842819 - VMware_bootbank_clusterstore_8.0.0-1.10.20842819 - VMware_bootbank_gc-esxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-base_8.0.0-1.10.20842819 - VMware_bootbank_esxio-combiner-esxio_8.0.0-1.10.20842819 - VMware_bootbank_vdfs_8.0.0-1.10.20842819 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.0-1.10.20842819 - VMware_bootbank_drivervm-gpu_8.0.0-1.10.20842819 - VMware_bootbank_esxio-combiner_8.0.0-1.10.20842819 - VMware_bootbank_trx_8.0.0-1.10.20842819 - VMware_bootbank_bmcal-esxio_8.0.0-1.10.20842819 - VMware_bootbank_bmcal_8.0.0-1.10.20842819 - VMware_bootbank_gc_8.0.0-1.10.20842819 - VMware_bootbank_esxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.0-1.10.20842819 - VMware_bootbank_native-misc-drivers_8.0.0-1.10.20842819 - VMware_bootbank_cpu-microcode_8.0.0-1.10.20842819 - VMware_bootbank_native-misc-drivers-esxio_8.0.0-1.10.20842819 - VMware_bootbank_esx-base_8.0.0-1.10.20842819 - VMware_bootbank_vsanhealth_8.0.0-1.10.20842819

	- VMware_bootbank_loadesx_8.0.0-1.10.20842819 - VMware_bootbank_esx-update_8.0.0-1.10.20842819 - VMware_bootbank_loadesxio_8.0.0-1.10.20842819 - VMware_bootbank_esxio-update_8.0.0-1.10.20842819
PRs Fixed	3041101
Related CVE numbers	N/A

This patch updates the following issue:

• **When TPM 2.0 is enabled with TXT on an ESXi host, attempts to power-on a virtual machine might fail**

When TXT is enabled on an ESX host, attempts to power-on a VM might fail with an error. In the vSphere Client, you see a message such as **This host supports Intel VT-x, but Intel VT-x is restricted. Intel VT-x might be restricted because 'trusted execution' has been enabled in the BIOS/firmware settings or because the host has not been power-cycled since changing this setting.**

This issue is resolved in this release.

ESXi80a-20842819

Name	ESXi
Version	ESXi80a-20842819
Release Date	December 8, 2022
Category	Bugfix
Affected Components	- ESXi Component - core ESXi VIBs - ESXi Install/Upgrade Component
PRs Fixed	3041101
Related CVE numbers	N/A

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliance$` shell. If you use VAMI, refresh your browser.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is deactivated, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

Miscellaneous Issues

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reinstall ESXi. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In the vSphere API Explorer, VMware Datacenter CLI (DCLI) and PowerCLI, you see an API option "contentinternal" that is not functional

You see an API option **contentinternal** in the metadata of either the vSphere API Explorer, DCLI and PowerCLI. For example, when you open `https://<your vCenter IP>/ui/app/devcenter/api-explorer`, you see the option in the **select API** drop-down menu. This option is not functional.

Workaround: Ignore the **contentinternal** API option and do not use it.

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client.

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: None.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example `vmnic6` and `vmnic7`) and one management controller (for example `vmnic8`):

```

[~] esxcfg-nics -l

vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller

vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller

[~] esxcfg-nics -ls

vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller

```

The high-speed ethernet controllers `vmnic6` and `vmnic7` register first and operate with RSS set to 16 receive queues.

```

[~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -
n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }

```

However, in rare cases, if the management controller `vmnic8` registers first with the vSphere Distributed Switch, the high-speed ethernet controllers `vmnic6` or `vmnic7` uplink might end up operating with RSS set to 1 receive queue.:

```

[~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }

```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:

```

[~] esxcfg-module -u ionic_en
[~] esxcfg-module ionic_en
[~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind.

```

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```

~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable

```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```

~# reboot

```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough (SR-IOV or UPT) to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

Networking Issues

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Storage Issues

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vCenter Server and vSphere Client Issues

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [90259](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_XXX. Extraction of image from host XXX.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Guest OS Issues

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

vCenter Server Update and Patch Release Notes

VMware vCenter 8.0 Update 3g Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)

- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 3g | 29 JUL 2025 | ISO Build 24853646
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2025-41241 and CVE-2025-41250. For more information on this vulnerability and its impact on Broadcom products, see [VMSA-2025-0014](#) and [VMSA-2025-0016](#) respectively.
- vCenter 8.0 Update 3g introduces 5 differentiated alarms for certificates expiration, replacing the current generic one, which does not provide details on the specific VECS repository. The new alarms identify the certificate that is about to expire and provide a link to a KB with remediation steps.
- This release contains bugfixes as listed in the Resolved Issues section.
- For vSphere Supervisor updates, see the [vSphere Supervisor 8.0 Release Notes](#).
- For Photon OS updates, see [VMware vCenter Server Appliance Photon OS Security Patches](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 3g

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00600-24853646-patch-FP.iso
Build	24853646
Download Size	8281.3 MB
sha256checksum	327fa456fb2883b3550f1fdd46a51ef52e2289c1510165741c00fab b0163e4ce
PRs fixed	3120158, 3040952, 3532388, 3514882, 3511057, 3488172, 3493240, 3521167, 3526879, 3526961, 3497083, 3526895, 3522903, 3325035, 3270623, 3033184, 3510273, 3530994, 3514558, 3516568, 3498829, 3295816, 3285774, 3270684, 3254294, 3262844, 3151016, 3501926, 3526913, 3276369, 3243942, 3026630, 3119837, 3494317, 3522443, 3530648, 3521452, 3523522, 3494312, 3269706, 3155558, 3525528, 2916365, 3323116, 3304747, 3507227, 3512385, 3504192
CVEs	CVE-2025-41241

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

See these KBs for details on the authenticated online and offline download: [390121](#), [390098](#), and [389276](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Miscellaneous Issues

PR 3120158: In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

This issue is resolved in this release.

PR 3040952: In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

This issue is resolved in this release.

vSphere Client and vCenter Issues

PR 3532388: In the vSphere Client, you cannot set the datetime format in Browser Default

When you select a datetime format from **My Preferences > Language > Browser Default**, for example English (New Zealand) with format DD/MM/yyyy, the datetime format in the vSphere Client continues to display in English (US) with format MM/DD/yyyy.

This issue is resolved in this release.

PR 3514882: When deploying a virtual machine with EFI firmware from a content library, the chipset property is always set to acpi as default value

Due to a configuration issue, virtual machine with EFI firmware deployed from Content Library might always take `acpi` as the default value for the `chipset.motherboardLayout` property.

This issue is resolved in this release.

PR 3511057: In a rare race condition, the PlaceVM API might cause the VPXD service to fail if the candidate host set is empty

In rare cases, while the Place VM API finds a candidate host for the Placement Spec, the task might race with a parallel task on that host for shutting down or entering maintenance mode. As a result, since PlaceVM API does not handle this corner case correctly, the vpxd service fails and might cause various issues such as failing SDDC manager prechecks.

This issue is resolved in this release.

PR 3488172: Login instructions for RSA SecurID Authentication in the vSphere Client might display HTML tags as plain text

In the vSphere Client, the login instructions for RSA SecurID Authentication might display HTML tags such as `<href>` and `<br>` as plain text. This issue does not affect the RSA functionality.

This issue is resolved in this release.

PR 3493240: You see no indication in the vSphere Client when an ESXi host is in standby mode

In the vSphere Client, when you put an ESXi host in standby mode, you might see no indication of the changed status, because the standby mode icon is not visible.

This issue is resolved in this release.

PR 3521167: If a vCenter task to move a file times out, the file might be deleted from both the destination and source datastores as part of the rollback

When you start a task to move a file from the vSphere Client, the file is copied from the source datastore to the target datastore and then deleted from the source. If the task times out, vCenter rolls back the operation, but in some cases the file in the target datastore might not be restored.

This issue is resolved in this release.

PR 3526879: Applying multiple tags to a Content Library overrides previous tags even if the cardinality of the respective category is Multiple

If you assign tags from the same category to a template in a local Content Library and publish the template to a subscribed library, the sync operation for the subscribed library might remove all other tags except one of the new tags, even if the cardinality of the respective category is Multiple.

This issue is resolved in this release.

PR 3526961: vCenter alarms for child objects in the Summary tab of parent objects are not cleared

In the vSphere Client, when you click **Reset to Green** or **Acknowledge** on an alarm for a child object in the **Summary** tab of a parent object, the alarm is not cleared. For example, if you click **Reset to Green** or **Acknowledge** on an alarm for an ESXi host in the **Summary** tab of a cluster, you still see the alarm on the cluster level.

This issue is resolved in this release.

PR 3497083: You see large threadmonitor<num>.log files in the vsphere-ui logs directory

Due to a configuration issue, the threadmonitor logs might not be properly compressed and can reach up to 10 files, 100 MB each.

This issue is resolved in this release.

PR 3526895: GPU compute and memory usage data bars in the vSphere Client have no color

In the vSphere Client, you do not see the bars for compute and memory usage of vSphere objects in the **GPU Details** card under the **Summary** tab in color as expected.

This issue is resolved in this release.

PR 3522903: You cannot cancel tasks from the list under Monitor > Tasks in the vSphere Client

If you try to cancel a task in progress from the list under **Monitor > Tasks** in the vSphere Client, the task might continue to run. The cancel option works for tasks in the **Recent Tasks** panel.

This issue is resolved in this release.

PR 3325035: You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

This issue is resolved in this release.

PR 3270623: You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

This issue is resolved in this release.

PR 3033184: If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

This issue is resolved in this release.

Migration Issues

PR 3510273: Cross-cluster migration of virtual machines might fail if the source cluster is of type vSAN ESA

Cross-cluster migration for a VM might fail when the source VM on vSAN ESA has a snapshot in native delta format and the destination is either vSAN ESA or vSAN OSA.

This issue is resolved in this release.

Networking Issues

PR 3530994: Network configuration on virtual machines with Red Hat Enterprise Linux 8 (RHEL 8) and earlier might remain unchanged after the guest customization

For VMs with RHEL 8 and earlier, the guest customization engine might only clean configuration files at `/etc/sysconfig/network-scripts/ifcfg-eth*` and leaves files that do not have the prefix `ifcfg-eth`. As a result, such files might continue to take effect, leading to inconsistent results in network connectivity for the VM.

This issue is resolved in this release.

PR 3514558: Virtual machine network adapter is not connected after a VM powers on

Ephemeral distributed virtual port groups create and delete ports dynamically when a VM powers on or off. When a VM is powered on from the ESXi host, a host-assigned DVPort is created and later replaced with a regular DVPort by a rekey operation from vCenter. If an ESXi host sync occurs before the VM powers on, vCenter might reclaim the host-assigned port and initiate a rekey while the VM is still powered off. As a result, the virtual NIC disconnects after the VM powers on.

This issue is resolved in this release.

Upgrade Issues

PR 3516568: After an update to vCenter 8.0 Update 3, you see a warning alarm for root disk exhaustion

When updating vCenter to version 8.0 Update 3 and later, an issue with the snapshot path might cause the root partition / overlay to reach almost 100% due to an increase in the size of container snapshots.

This issue is resolved in this release.

PR 3498829: vCenter upgrade might fail due to undecodable characters in the vCenter certificates

If the human-readable decoded form of vCenter certificates such as VECS MACHINE_SSL_CERT, TRUSTED_ROOTS, or TRUSTED_ROOT_CRLS, contains a string that cannot be decoded with UTF-8, such as "ñ", the first boot of the VMware Authentication Framework Service Daemon (VMAFD) might fail. As a result, you cannot complete an upgrade of your vCenter system.

This issue is resolved in this release.

PR 3295816: You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLS>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

This issue is resolved in this release.

PR 3285774: A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: Update 8.0.2.00000 for component vlcm is not found.

This issue is resolved in this release.

PR 3270684: You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot /storage/core unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as No such file or directory: '/storage/core/software-update/install_operation'. In the software-packages.logs, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
```

```
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
```

```
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
```

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
```

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
```

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state file from State._get

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf

2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed. Please collect the VC support bundle.

This issue is resolved in this release.

PR 3254294: Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

This issue is resolved in this release.

PR 3262844: During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

This issue is resolved in this release.

PR 3151016: Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a vSphere Lifecycle Manager image in the vSphere Client and then navigate away.
2. You trigger a page refresh.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session.

If you use API, firmware compliance details are missing from the response.

This issue is resolved in this release.

vSAN Issues

PR 3501926: In the vSphere Client, during data migration precheck on vSAN clusters, you see an error message that the requested data cannot be extracted

Due to a rare null pointer error, during a data migration precheck on vSAN clusters, you might see the error `Unable to extract requested data`. Check vSphere Client logs for details in the vSphere Client.

This issue is resolved in this release.

vSphere Lifecycle Manager Issues

PR 3526913: Cluster images with removed components might show ESX hosts as non-compliant after activating or deactivating vSphere solutions

You can remove some components from the vSphere Lifecycle Manager image of a cluster, like VMware-VM-Tools, for some reason, such as security consideration, and if you use it for remediation, ESX hosts are still compliant to the cluster image. However, if you activate or deactivate a solution such as vSphere High Availability on that cluster, consequent cluster compliance scans might show ESX hosts in the cluster as noncompliant.

This issue is resolved in this release.

PR 3276369: VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

This issue is resolved in this release.

PR 3243942: You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

This issue is resolved in this release.

PR 3026630: You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

This issue is resolved in this release.

PR 3119837: Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

This issue is resolved in this release.

API Issues

PR 3494317: The vpxd service fails because the Guest OS customization API requires IPv4 settings

When you use API to create a Guest Customization Specification and customize the guest OS of a virtual machine, you might need to enter IPv4 settings in the specification or the vpxd service fails with a core dump.

This issue is resolved in this release.

PR 3522443: The vpxd service might fail when invoking the CompatibilityChecker.checkCompatibility() API

Invoking the `CompatibilityChecker.checkCompatibility()` API to evaluate the compatibility of a virtual machine connected to an NSX-backed distributed port group against a destination ESXi host might cause the vCenter vpxd service to fail.

This issue is resolved in this release.

Guest OS Customization Issues**PR 3530648: You cannot customize a Linux virtual machine with installed SUSE Linux Enterprise Server 16 (SLES 16)**

Guest OS customization for SLES 16 VMs is not supported and customization tasks fail.

This issue is resolved in this release. The fix adds guest customization support to SLES 16.

PR 3521452: Guest OS customization might fail on Ubuntu cloud image virtual machines

When you deploy a virtual machine from an Ubuntu cloud image and try to customize the VM, the customization task might fail with an error that the VM network adapter is in disconnected state. In the `/var/log/vmware-imc/toolsDeployPkg.log`, you see that the task stops at Command: `"/usr/sbin/netplan apply 2>&1`.

This issue is resolved in this release.

Security Issues**PR 3523522: Adding a Standard Key Provider to your vCenter system from the vSphere Client fails with Vim.fault.DatabaseError**

Starting with vCenter 8.0 Update 1c, you must use an RSA certificate when adding a Standard Key Provider to vCenter otherwise in the vSphere Client you see a `Vim.fault.DatabaseError` error. In the `vmafdd.log` file, you see the following line:

```
<Timestamp> [vmafdd][ERROR] Certificate uses an unsupported signature algorithm
(NID=ecdsa-with-SHA256). Only SHA-2 RSA algorithms are supported on the vCenter Server.
```

This issue is resolved in this release.

PR 3494312: Licensing info in a custom sysprep answer file for Windows might be in plain text

When you use the VM Customization Specifications workflow to create a custom sysprep answer file for Windows, if the file contains licensing info, such as a key, the key might display in plain text in the file.

This issue is resolved in this release.

PR 3269706: You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

This issue is resolved in this release.

Virtual Machine Management Issues**PR 3155558: In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted**

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

This issue is resolved in this release.

Storage Issues

PR 3525528: Hot extension of vSphere Virtual Volumes disk can fail due to a datastore accessibility check

During a hot-extend operation on a vSphere Virtual Volumes disk, the datastore accessibility check might fail even if another ESXi host on the datastore, not the one where the VM is located, is inaccessible. As a result, you see the error `The operation is not supported on the object. The operation can not be performed because datastore is inaccessible.`

In the `vpzd` logs, you see an error such as:

```
2025-05-06T15:36:31.731Z warning vpzd[07478] [Originator@6876 sub=provisioning
opID=mabmw1q-73097-auto-1kej-h5:70002251-2a] Datastore is not accessible from host
```

This issue is resolved in this release. The fix makes sure the check includes only datastore accessibility on the VM-specific hosts where VM sharing the disk are located.

PR 2916365: VASA API version does not automatically refresh after upgrade to vCenter 8.0

vCenter 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of vCenter, the VASA API version remains unchanged after you upgrade to vCenter 8.0. For example, if you upgrade a vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

This issue is resolved in this release.

PR 3323116: Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState.`

This issue is resolved in this release.

vSphere Configuration Profiles Issues

PR 3304747: You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

This issue is resolved in this release.

vSphere High Availability Issues

PR 3507227: A faulty health status of a microservice under VMware Identity Services might trigger an unexpected vCenter High Availability failover

A false negative health status of a microservice under VMware Identity Services that occurs in very specific conditions might trigger an unexpected vCenter High Availability failover.

This issue is resolved in this release.

Auto Deploy Issues

PR 3512385: vCenter might become unresponsive due to overloading of the /tmp folder space with VIB files

While populating the Auto Deploy VIB cache, vCenter might initiate multiple requests to get the same VIB file which might lead to an overload of the /tmp folder space. As a result, vCenter becomes unresponsive and regular tasks such as deploying a virtual machine or creating a backup might fail.

This issue is resolved in this release.

PR 3504192: Booting a stateless host by using UEFI might fail with a "crypto64.efi: Not found" error

Booting a stateless host by using UEFI might fail with iPXE version 1.21.1-vmw. In the vSphere Client, you see an error such as `crypto64.efi: Not found`.

This issue is resolved in this release. The fix rolls back the iPXE version to 1.0.0-vmw.

Known Issues

Miscellaneous Issues

Revert snapshot operation fails with the error 'The operation is not allowed in the current state.'

If a virtual machine has a chain of snapshots, which contains both VM and Improved Virtual Disk (IVD), also known as First Class Disk (FCD), snapshots, reverting to a VM snapshot is not supported. If you try such an operation, in the logs you see the message `The virtual machine cannot be reverted when crossing an Improved-Virtual-Disk snapshot.` or an error such as `The operation is not allowed in the current state.`

Workaround: None. The operation is not supported.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you do not open port 9087 in your firewall between ESXi hosts and vCenter, compliance checks and vSphere HA might fail after vCenter update to 8.0 Update 3

Starting from 8.0 Update 3, the vSphere Lifecycle Manager downloads updates for ESXi hosts by a HTTPS connection to the vCenter instance on port 9087. If you do not open port 9087 in your firewall between ESXi hosts and vCenter, you might see compliance check errors. For example, in the `lifecycle.log` you see messages such as:

```
<Timestamp> In(14) lifecycle[2112988]: Downloader:373 Opening https://<VC-
FQDN>:9087/vum/repository/hostupdate/__micro-depot__vendor-DEL__DEL-ESXi-8.0-Addon-
cumulative_metadata__index__.xml for download
```



```
<Timestamp> Wa(12) lifecycle[2112988]: Downloader:210 Download failed: <urlopen error
timed out>, 9 retry left...
```

In addition, vSphere HA might fail to start.

Workaround: Open port 9087 in your firewall between ESXi hosts and vCenter by using the TCP protocol.

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system

Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as `The host returns esxupdate error codes: -1.`

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:
Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails

In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to

detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

You see a "vc.health.error.dbjob3" warning after a vCenter upgrade

In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

Workaround: See [Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x](#) how to clear stats data, if not relevant anymore.

Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"

Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists`.

Workaround: Delete all files under `/storage/log/vmware/analytics/stage`, `/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Verify following URLs and their respective statuses and follow KB 93526.
```

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

You cannot upload logs packages to a VMware service request by using the vSphere Client

You cannot upload logs packages to a VMware service request by navigating in the vSphere Client to **Administration > Support > Upload File to Service Request**. Instead, you see a link to a KB article that provides steps to upload log bundles in the Broadcom Support Portal. This is expected.

Workaround: Follow the steps in [Uploading files to cases on the Broadcom Support Portal](#).

When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client

If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

You see the same license product name displayed multiple times in the vSphere Client

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

Workaround: None

An external gateway firewall might block vSphere Pod traffic to clusterIPs

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to ClusterIPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see [vSphere Pod Traffic to ClusterIP Time-outs](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

The vpxd agent might fail after expanding the disk size of a virtual machine with a NVMe controller in a vCenter 8.0 Update 3 mixed environment

If you have ESXi hosts of version 8.0 Update 1 and earlier in your vCenter 8.0 Update 3 system, expanding the disk of a VM with a NVMe controller might cause the vpxd service to fail with `core.vpxd-worker coredump`.

Workaround: Upgrade your ESXi hosts of version 8.0 Update 1 and earlier to 8.0 Update 3. Alternatively, use the VMware Host Client to extend the disk size from **Edit settings** of the VM.

vSphere Lifecycle Manager Issues

In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image

When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components**

and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0 Update 3e Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 3e | 10 APR 2025 | ISO Build 24674346
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2025-41225 and CVE-2025-41228. For more information on these vulnerabilities and their impact on Broadcom products, see [VMSA-2025-0010](#).
- This release contains bugfixes as listed in the Resolved Issues section.
- For vSphere Supervisor updates, see the [vSphere Supervisor 8.0 Release Notes](#).
- For Photon OS updates, see [VMware vCenter Server Appliance Photon OS Security Patches](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 3e

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00500-24674346-patch-FP.iso
Build	24674346
Download Size	7925 MB
sha256checksum	44774f4b874d669f3ed5ecba3828da8909dc9dfe5ce219e6d48f75581c40c161
PRs fixed	3442316, 3455754, 3442983, 3476963, 3470631, 3478686, 3484244, 3466450, 3488834, 3426784, 3483510, 3430738, 3442875, 3467472, 3483567, 3432060, 3488092, 3488098, 3470361, 3465133, 3474651, 3452081, 3442763, 3467808, 3465275, 3469515, 3452906, 3493822, 3470090, 3491095, 3456483, 3486630, 3472936, 3459747, 3454123, 3478625, 3467612, 3490881, 3439353

CVEs	CVE-2025-41225, CVE-2025-41228
------	--------------------------------

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

See these KBs for details on the authenticated online and offline download: [390121](#), [390098](#), and [389276](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Miscellaneous Issues

PR 3434710: After upgrading to vSphere 8.0 Update 3, Single Sign-On (SSO) users in Azure Active Directory groups get a "namespaces is forbidden" error when using kubectl

Due to the way vCenter handles User Principal Name (UPN) formatting of just-in-time (JiT) groups of external identity providers such as Azure Active Directory, you might see the error `namespaces is forbidden: User` when using the `kubectl` tool.

This issue is resolved in this release.

PR 3442316: A very rare issue with security tokens might trigger excessive warnings in the apigw.log

In extremely rare cases, when two or more threads in the Security Token Service (STS) encounter an expired token at exactly the same time, the first thread might trigger the acquisition of a new token, while the other threads enter an endless loop, because they continue using the expired token. As a result, you see warnings flooding the `apigw.log` of your syslog server.

This issue is resolved in this release. For more information, see KB [378091](#).

PR 3455754: Guest OS customization writes network configuration to both NetworkManager keyfile and deprecated ifcfg-rh profiles

Guest OS customization tasks on RHEL 9.x, AlmaLinux 9.x, Rocky Linux 9.x, and Oracle Linux 9.x virtual machines write the network configuration to both `NetworkManager` keyfile and the deprecated `ifcfg-rh` profiles, and cause the `NetworkManager` to load mixed profiles, especially on VMs with multiple network adapters.

This issue is resolved in this release. With the fix, the network configuration after a guest OS customization is located only in the `NetworkManager` keyfile profile at `/etc/NetworkManager/system-connections/VMware-customization-{NAME}.nmconnection`.

PR 3442983: IPv6 DNS server does not work after guest OS customization on RHEL 9.x, AlmaLinux 9.x, Rocky Linux 9.x, and Oracle Linux 9.x virtual machines

For RHEL 9.x, AlmaLinux 9.x, Rocky Linux 9.x, and Oracle Linux 9.x virtual machines, during guest OS customization, the IPv6 DNS server settings might be placed under the IPv4 section in the `NetworkManager` keyfile located at `/etc/NetworkManager/system-connections/VMware-customization-{NAME}.nmconnection`. As a result, the IPv6 DNS server does not work on the customized virtual machine.

This issue is resolved in this release.

vSphere Client and vCenter Issues

PR 3454774: If you try to re-log in to an existing vCenter session with wrong password, the vCenter vpxd service fails

If you try to re-log in to an existing vCenter session with an invalid password, the vpxd service fails. As a result, immediate attempts to log in to vCenter fail with an error `503 Service Unavailable`. Log in succeeds after the vpxd service restarts.

This issue is resolved in this release. The fix prevents the vpxd service from failure in such cases.

PR 3478332: When deleting a virtual machine from the list of a Storage Distributed Resource Scheduler (Storage DRS) rule, a random entry in the list might be deleted instead of the selected

In the vSphere Client, when you select to delete a specific VM from the list of a Storage DRS rule, a different entry might be deleted instead of the selected one from the list. The issue occurs intermittently.

This issue is resolved in this release.

PR 3461082: The monitoring query for storage consumption by Alarm data type in vCenter returns an error

Due to a mismatch between the actual table name and the query table name, the monitoring query for storage consumption by Alarm data type in vCenter, `storage.totalsize.directory.vcdb_alarms`, returns an error.

This issue is resolved in this release.

PR 3470631: If vCenter is configured behind a https proxy, you cannot configure an external identity provider such as EntraID

DNS resolution for external identity providers such as EntraID or Azure might fail on vCenter when a https proxy is configured on the vCenter instance.

This issue is resolved in this release.

PR 3478686: Smart card authentication on vCenter intermittently fails with the error "Unable to validate submitted credentials"

The Certificate Revocation List (CRL) cache updates hourly and if you log in to vCenter by using a smart card during such an hourly cycle in which the old CRL is expired but no new CRL is downloaded, the login fails. Since this is a corner case, smart card authentication on vCenter intermittently fails and in the vSphere Client, you see the error `Unable to validate submitted credentials`.

This issue is resolved in this issue.

PR 3484244: A rare race condition might cause repeating failovers due to a failure of the vCenter vpxd service or vSphere Client becoming unresponsive

A rare race condition during a VAPI call from the vpxd service to a virtual machine that coincides with the deletion of a `VirtualMachine` object might cause the vpxd service to fail. As a result, you see repeating failovers or the vSphere Client might become unresponsive.

This issue is resolved in this release.

PR 3466450: You don't see the Subject Key Identifier field after an update of the machine SSL certificate

When you use the Certificate Management workflow in the vSphere Client to update the machine SSL certificate of a vCenter instance, the new certificate lacks the mandatory subject key identifier `x509` field. As a result, the updated certificate becomes non-compliant with RFC 5280. The issue might also occur by using API to update the machine SSL certificate.

This issue is resolved in this release.

PR 3488834: In the vSphere Client, you do not see information in the VM Overrides tab

In rare cases, some of the following **VM Override** settings might be unset and not handled properly: `failureInterval`, `maxFailureWindow`, `maxFailures`, `minUpTime`. As a result, when you select a cluster and open **VM Overrides**

under **Configuration**, you see no information even if such was previously present and cannot manage existing or new overrides. In the dataservice log on vCenter, you see logs such as:

```
[2025-01-08T13:17:40.030Z] [WARN ] data-service-pool-780 70063039 101588 200147
TaskExecutor [queryId: qb-85656:h5ui-
getData:com.vmware.vsphere.client.clusterui.model.configuration.VmOverridesData:1931023000]
Data Service execution for task: Property 'vmOverrides' was
fetched by 'RiseToVisePropertyProviderAdapter' has failed
java.lang.IllegalStateException: Property provider method 'getVmOverrides'
of 'com.vmware.vsphere.client.clusterui.impl.VmOverridePropertyProvider'
threw exceptionCaused by: java.lang.NullPointerException: nullat
com.vmware.vsphere.client.clusterui.util.VmOverridesUtil.getVmMonitoringSensitivityValue(VmOverri
com.vmware.vsphere.client.clusterui.util.VmOverridesUtil.instantiateVmOverridesListItem(VmOverri
com.vmware.vsphere.client.clusterui.util.VmOverridesUtil.getVmOverrideItem(VmOverridesUtil.java:
```

This issue is resolved in this release.

PR 3426784: The vpxd service fails during statistics queries from vCenter when an unmounted datastore is not properly cleaned from the inventory

In rare cases, the vpxd service might not properly clean from the inventory a datastore that is unmounted from all ESXi hosts and such datastores remain with a null `MountInfo` path. As a result, when statistics queries from vCenter run on such datastores by using the `QueryPerf` API, the vpxd service fails due to a `NullPointerException`. You can identify the issue in the last line in the vpxd logs after failure:

```
info vpxd[360244] [Originator@6876 sub=vpxLro opID=628f05c2] [VpxLRO] -- BEGIN
lro-384 -- PerfMgr -- vim.PerformanceManager.queryStats -- 52c6eba3-68a1-7a84-
f665-65987ede6d38(52ddc34d-5336-7802-b973-2e368048a410)warning vpxd[360244]
[Originator@6876 sub=StatsAggregator opID=628f05c2] counterId = 281 and instance = not
found ..info vpxd[360244] [Originator@6876 sub=StatsAggregator opID=628f05c2] There is
nothing to aggregate
```

This issue is resolved in this release.

PR 3483510: An overload of the vCenter sidecar proxy service might cause failing vCenter workloads or login to vCenter

If the memory utilization of the vCenter sidecar proxy service, Envoy-sidecar, reaches 98% of the 1GB limit, Envoy starts sending overload responses that might cause failures in the vCenter internal workloads or logins to vCenter.

This issue is partially resolved in this release. vCenter 8.0 Update 3e reduces Envoy memory consumption where possible to reduce the likelihood the issue to occur.

PR 3430738: The Customize hardware page in the Deploy From Template wizard becomes unresponsive after you change the network adapter to a different distributed port group

In the vSphere Client, when you right-click a template from a local content library and select **New VM from This Template** in the **Deploy From Template** wizard, you might not be able to complete the workflow if you change the network adapter to a different distributed port group, because the **Customize hardware** page becomes unresponsive.

This issue is resolved in this release.

PR 3442875: After patching vCenter to 8.0 Update 3, the vSphere Client Log In button is dimmed and RSA SecureID option activates automatically

Due to a vCenter session that might remain open after patching vCenter to 8.0 Update 3, the **Log In** button of the vSphere Client might not be active even after you provide login credentials. Instead, the RSA SecureID option activates automatically.

This issue is resolved in this release.

PR 3467472: If the API version of an ESXi host is empty in the vCenter database, you might see spikes in CPU utilization

If for some reason, the API version of any ESX host is set to EMPTY in the vCenter DB, the vpxd service might loop and cause spikes in CPU utilization, based on vpxd tasks.

This issue is resolved in this release.

PR 3483567: An overload of the vCenter reverse proxy service might cause failing vCenter workloads or login to vCenter

If the memory utilization of the vCenter reverse proxy service, Envoy, reaches 98% of its 1GB limit, it starts sending overload responses that might cause failures in the vCenter internal workloads or logins to vCenter.

This issue is resolved in this release.

PR 3432060: You get an invalid file when exporting system logs for virtual machines from the vSphere Client

Due to an escape character issue, when you export system logs on virtual machine from the vSphere Client, you get a 1KB invalid file.

This issue is resolved in this release.

PR 3488098: vCenter alarms do not disappear after you click Acknowledge or Reset to Green

In the **Summary** tab of the vSphere Client, vCenter alarms in the **Issues and Alarms** panel might not disappear after you click **Acknowledge** or **Reset to Green**.

This issue is resolved in this release.

PR 3470361: In the vSphere Client, you still see an NSX overlay segment after deleting it from the NSX Manager

When adding NSX port groups on a vCenter system, vCenter uses an API call to push the port groups to ESXi hosts, but in some cases the call does not complete on all hosts, if a host is temporarily disconnected for example. As a result, from vCenter perspective, all port groups exist, but from ESXi hosts perspective, some might not exist. Consequently, when you try to delete an NSX overlay segment from the NSX Manager, it still shows in the vSphere Client, because not all port groups can be removed.

This issue is resolved in this release.

PR 3465133: The vpxd service might fail to start due to duplicate Distributed Virtual Port Groups (DVPG)

An issue in NSX might cause the creation of an NSX DVPG with a name exceeding 80 characters when the security feature is active on the vSphere Distributed Switch (VDS). This results in an incomplete record in the vCenter DB, which triggers NSX to recreate the DVPG and ultimately leads to duplicate DVPGs. Consequently, the vpxd service might fail to start and the vCenter instance becomes unresponsive.

This issue is resolved in this release.

Migration Issues**PR 3474651: Distributed Resource Scheduler (DRS) fails to migrate an encrypted VM with an error for missing Cryptographer. Migrate privileges**

In some cases, when DRS schedules to migrate an encrypted VM, it might use the context of a user without `Cryptographer.Migrate` privilege and migration fails with an error in the vSphere Client such as:

```
Permission to perform this operation was denied. You do not hold privileges "virtual machine vm-xx : [Cryptographic operations > Migrate] .
```

In the vpxd.log, you see error messages such as:

```
YYYY-MM-DDThh:mm:ss.472Z info vpxd[06061] [Originator@6876 sub=drsExec opID=<op_id>]
Executing DRM recommended vMotion in [vim.ClusterComputeResource:domain-c<id>,esxi_name],
migrating [vim.VirtualMachine:vm-<vm_id>,vm_name] from vim.HostSystem:host-<src_id> to
vim.HostSystem:host-<dstn_id>.....
```

```
YYYY-MM-DDThh:mm:ss.494Z info vpxd[06061] [Originator@6876 sub=pbm opID=<op_id>-01] pre
migrate callback is skipped - for cached from pre migrate check reason
```

```
YYYY-MM-DDThh:mm:ss.494Z warning vpxd[06061] [Originator@6876 sub=CryptoManager
opID=<op_id>-01] The session xx-yy-zz of user DOMAIN\user_name does not have privilege
Cryptographer.Migrate on entity [vim.VirtualMachine:vm-<vm_id>,vm_name].
```

This issue is resolved in this release.

Networking Issues

PR 3452081: The vpxd service fails after you remove an ESXi host from a vSphere Distributed Switch (VDS)

When you remove an ESXi host from a VDS and make no changes on the physical NIC of the host, the parameter tracking if the host is a member of that VDS might be null. As a result, the vpxd service might fail while removing such hosts from a DVS, because it considers the task as illegal memory accessing due to the null parameter.

This issue is resolved in this release.

PR 3442763: You see duplicate routes in the Netplan configuration file after guest OS customization on Ubuntu VMs

For Ubuntu releases that support the Netplan utility for configuring networking on a Linux system, when you deploy an Ubuntu VM and customize it with both IPv4 and IPv6 gateways, the Netplan configuration file `/etc/netplan/99-netcfg-vmware.yaml` file of the guest OS contains two route sections under a single network interface.

This issue is resolved in this release. With the fix, customizing an Ubuntu VM with both IPv4 and IPv6 gateways creates a single route section for each NIC in the Netplan configuration file, and adds the corresponding IPv4 and IPv6 gateways of the NIC to its route section.

PR 3467808: Virtual machines connected to an NSX portgroup lose connectivity after a vSphere Distributed Switch (VDS) upgrade

During an upgrade of the VDS from version 6.6 to 7.0, if any ESXi host member of the DVS fails to upgrade, for example the host is disconnected from vCenter, the upgrade stays in **In progress** state. In such cases, vCenter incorrectly tries to repush the NSX logical switch port to the ESXi host with a dummy teaming policy, which leads to VMs losing connectivity.

This issue is resolved in this release.

Auto Deploy Issues

PR 3465275: Virtual machine operations such as power on and migration fail on ESXi hosts deployed by using Auto Deploy with an error for expired license key

The expiration of the evaluation period on an ESXi host deployed by using Auto Deploy might remove the license on the host. As a result, you continue to see the hosts as properly licensed in vCenter, but virtual machine operations such as power on and migration start failing with an error in the vSphere Client such as `License Key has expired`. Please install a valid license to use this product.

This issue is resolved in this release. For more information, see KB [379119](#).

PR 3469515: You cannot apply Auto Deploy stateful host configuration with the auto-partitioning option

If the auto partitioning option `VMkernel.Boot.autoPartition = true` is active on a host you use for stateful installs, provisioning ESXi hosts fails with an error that the partition is busy.

This issue is resolved in this release.

PR 3452906: Stateful ESXi installations by using Auto Deploy fail at boot and do not create a local datastore

Due to an issue with requesting a new license during the boot of stateful ESXi installations by using Auto Deploy, when you use a host profile to provision ESXi hosts on the local hard drive, the task fails and does not create a local datastore.

This issue is resolved in this release.

VM Migration and Snapshots Issues

PR 3493822: Scheduled tasks to delete VM snapshots might not adhere to the VM retention policy

In the vSphere Client, when you create a scheduled task to delete all snapshots of virtual machines on a given ESXi host, the task might delete all snapshots regardless of any retention policy on the VMs. The issue also occurs if you to delete snapshots by using API. It affects only ESXi hosts of version earlier than 8.0 Update 3.

This issue is resolved in this release.

vSphere Configuration Profiles Issues

PR 3470090: In the vSphere Client, you still see vSphere Configuration Profiles (VCP) host override after deleting it

With vSphere 8.x, you can use VCP as a desired configuration framework that allows you to manage ESXi host settings in the vSphere Client at the cluster-level by using a JSON configuration document. Part of the configuration customization is the ability to override the host settings for specific hosts. In some cases, after you decide to remove such an override and confirm the **Delete Override** action, you still see the override listed in the configuration document.

This issue is resolved in this release.

Cloud Native Storage Issues

PR 3491095: After upgrading to vCenter 8.x, you see frequent errors related to Cloud Native Storage (CNS) that a parameter is not correct

After an upgrade to vCenter 8.x, CNS runs a full sync on datastores to clean up stale block volumes. When a datastore has both file and block volumes, the CNS periodic sync function might incorrectly indicate file volumes as stale block volumes. In such cases, CNS calls block volume API `retrieveVStorageObject` on the file volumes, which fails with `A specified parameter was not correct` error. CNS adds the file volumes to the CNS retry table `vpv_storage_volume_update` and continues retrying the `retrieveVStorageObject` call. As a result, in the vSphere Client you see frequent `A specified parameter was not correct` messages.

This issue is resolved in this release.

Fault Tolerance and High Availability Issues

PR 3482950: If you use the thumbprint mode for managing ESXi certificates from vCenter, High Availability agents on hosts become unresponsive after an update to vCenter 8.0 Update 3

If you use the thumbprint mode for managing ESXi certificates from vCenter, High Availability agents on hosts might become unresponsive after an update to vCenter 8.0 Update 3 and later.

This issue is resolved in this release.

PR 3456483: vCenter High Availability cluster state degrades or the vpxd service might fail when root user password expires on a passive or witness node

When the root user password expires on either the passive or witness nodes, file replication in vCenter High Availability starts to fail. As a result, vCenter High Availability health reports degrade and replication is out of sync, and the vpxd service might fail.

This issue is resolved in this release.

PR 3486630: The vpxd service in a vSphere High Availability (HA) cluster might fail when one or more virtual machines are configured with Fault Tolerance

In a corner case, when on a VM configured with Fault Tolerance multiple state transitions run in parallel, the vpxd service in the vSphere HA cluster might encounter a null pointer dereference and fail.

This issue is resolved in this release.

Upgrade Issues

PR 3472936: When you change a desired ESXi configuration, a host override might remove all settings set to false for hosts with host-specific settings

When you use the `host-override` option to set ESXi host-specific values for a setting, properties set to `false` might be missing from the desired configuration. For example, if you have the following setting: `"ipv6": { "dhcp": false, "dhcp_dns": false, "auto_configuration_enabled": true }` the output might be just: `"ipv6": { "auto_configuration_enabled": true }` and in the vSphere Client you see errors such as: Required property 'xxx' missing.

This issue is resolved in this release.

PR 3459747: After upgrading from vCenter 7.x to 8.0 Update 3, PODs do not start, and Persistent Volumes fail to attach to the nodes

After an upgrade from vCenter 7.x to 8.0 Update 3, while modifying database tables to align with the vCenter 8.x database schema, the `CNS.VOLUME_INFO` table schema in the `CNS.VPX_TABLE` might be incorrect. Inserting data into the `CNS.VOLUME_INFO` table fails with the following error in `/var/log/vmware/vsan-health/vsanvcmgmtd.log` on vCenter:

```
DB Error while inserting volume ***: Column ALLOCATED_SPACE does not exist in Table
Definition of table volume_info" or "Cannot process vStorage objectxxxxxxxx-xxxx-xxxx-
xxxx-xxxxxxxxxxxx on ds:///vmfs/volumes/****.
```

As a result, creating PODs is blocked in `ContainerCreating` state and Persistent Volumes fail to attach to the nodes.

This issue is resolved in this release.

vSAN Issues

PR 3425316: After upgrading to vCenter 8.0 Update 3, you see a vSAN cluster configuration consistency health alert

If you have active vSAN encryption for a vSAN cluster in a vCenter of version earlier than 8.0, and you have done shadow rekey for the vSAN cluster, after an upgrade to vCenter to 8.0 Update 3, you might see a vSAN health alert `Host(Cluster) Has Finished Encryption Configuration But Cluster(Host) Has Not under the vSAN Cluster Configuration Consistency` health test. This alert is a false negative and you can ignore it.

This issue is resolved in this release.

PR 3488117: The disk format of virtual machines doesn't change from thick to thin after migration to a vSAN 8.x cluster

Even when you select vSAN Default Storage Policy, which defaults to thin disk format, in the **Select storage** step of the VM migration workflow, after migration to a vSAN 8.x cluster, the disk format might remain thick.

The issue is resolved in this release.

PR 3454123: You see no vSAN capacity overuse alarm after updating to vSphere 8.0 Update 3e

With vSphere 8.0 Update 3e, if you hold a VMware vSphere Foundation license, you no longer see an alarm for vSAN capacity overuse as per your previous vSphere Foundation license, as the amount of vSAN capacity included in VMware vSphere Foundation has increased to 0.25 TiB per core.

This issue is resolved in this release.

PR 3478625: vSAN File Service virtual machines fail to deploy with an error for EAM Trust issue

When you activate the vSAN File Service on a vCenter with multichain certificate configured, vSAN File Service virtual machines might fail to deploy with an error for an EAM trust issue.

This issue is resolved in this release.

vSphere Lifecycle Manager Issues

PR 3467612: vSphere Lifecycle Manager memory leaks during download from online depots might cause vCenter memory allocation errors

During daily online depot syncs, every time the vSphere Lifecycle Manager downloads a file, it creates an HTTPS session with the online depot. When the HTTPS session ends, vSphere Lifecycle Manager might not clean up the session properly, specifically the SSL context. Since syncs are daily and include downloads of multiple files, such leaks add up and might cause memory allocation errors in vCenter.

This issue is resolved in this release.

Backup and Restore Issues

PR 3450499: HTTP(S) backup fails with the error "Failed to create backup directory on backup server."

For HTTP(S) backups, if the root directory on the backup server does not have read/write access, backup fails with error message `Failed to create backup directory on backup server..`

This issue is resolved in this release.

PR 3490881: vCenter restore from backup fails with "ReferenceError: oid is not defined" error

In rare cases, depending on the key exchange algorithms that some backup servers use to communicate with the vCenter GUI Installer, you might see the error `ReferenceError: oid is not defined` in the **Enter backup details** step of the vCenter restore flow in Stage 1.

This issue is resolved in this release.

PR 3439353: SSH File Transfer Protocol (SFTP) backup fails to connect to the backup server with an error "command ssh-keyscan returned non-zero exit status 1"

Description: SFTP backups might fail to establish a connection with the backup server. In the `applmgmt.log/backup.log`, you see an error such as: `"Command ['ssh-keyscan', '-p', 'port', '-H', 'hostname'] returned non-zero exit status 1."`

This issue is resolved in this release.

vSphere Supervisor Issues

PR 3473785: An issue with VM classes in vSphere systems with multiple Supervisor instances might cause core dumps and high log intensity

In some cases, in vSphere environments with many Supervisors instances, a VM class in one vSphere namespace in a Supervisor instance might be propagated to other Supervisor instances. As a result, you see core dumps and high log

intensity with errors such as: `CreateOrUpdate VM class failed. vmclass: <NAMESPACE-NAME1>/<VMCLASS-NAME>, Operation result: unchanged, and err:namespaces "<NAMESPACE-NAME1>" not found.`

This issue is resolved in this release.

Known Issues

vCenter and vSphere Client Issues

You see increased network traffic from vSphere Client to VMware Cloud

If for some reason the browser that you use to connect to VMware Cloud by the vSphere Client cannot access `vcsa.vmware.com` by HTTPS or port 443, you might see significantly increased network traffic. The issue occurs because when the Customer Experience Improvement Program (CEIP) is active and vSphere Client cannot access `vcsa.vmware.com`, CEIP unsuccessfully keeps trying to send data.

Workaround: Check the network configuration and make sure that requests from the vSphere Client to `vcsa.vmware.com` are successful or deactivate CEIP.

Security Issues

Update Manager Download Service (UMDS) stores proxy passwords in plain text

When UMDS accesses a VMware online depot by using a proxy, if you have configured a proxy password, it is stored in the configuration as plain text, which in some environments might be a security concern.

Workaround: Instead of configuring a password, you can pass the password every time you run a `Download (-D)` and `Re-Download (-R)` command. For example, `./vmware-umds -D --proxy-username "<your_username>" --proxy-password "<your_password>" --proxy-port 80 --proxy-ip 1.1.1.1.`

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLs>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](https://kb.vmware.com/s/article/93526) and [94934](https://kb.vmware.com/s/article/94934).

If you do not open port 9087 in your firewall between ESXi hosts and vCenter, compliance checks and vSphere HA might fail after vCenter update to 8.0 Update 3

Starting from 8.0 Update 3, the vSphere Lifecycle Manager downloads updates for ESXi hosts by a HTTPS connection to the vCenter instance on port 9087. If you do not open port 9087 in your firewall between ESXi hosts and vCenter, you might see compliance check errors. For example, in the `lifecycle.log` you see messages such as:

```
<Timestamp> In(14) lifecycle[2112988]: Downloader:373 Opening https://<VC-
FQDN>:9087/vum/repository/hostupdate/__micro-depot__vendor-DEL__DEL-ESXi-8.0-Addon-
cumulative_metadata__index__.xml for download

<Timestamp> Wa(12) lifecycle[2112988]: Downloader:210 Download failed: <urlopen error
timed out>, 9 retry left...
```

In addition, vSphere HA might fail to start.

Workaround: Open port 9087 in your firewall between ESXi hosts and vCenter by using the TCP protocol.

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system

Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as `The host returns esxupdate error codes: -1`.

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:
Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails

In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

You see a "vc.health.error.dbjob3" warning after a vCenter upgrade

In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

Workaround: See [Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x](#) how to clear stats data, if not relevant anymore.

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"

Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists.`

Workaround: Delete all files under `/storage/log/vmware/analytics/stage`, `/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Verify following URLs and their respective statuses and follow KB 93526.

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```

root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback

```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings > <specific broadcom NIC> > Device Configuration Menu > Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 1:

Default VM Settings

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

You cannot upload logs packages to a VMware service request by using the vSphere Client

You cannot upload logs packages to a VMware service request by navigating in the vSphere Client to **Administration > Support > Upload File to Service Request**. Instead, you see a link to a KB article that provides steps to upload log bundles in the Broadcom Support Portal. This is expected.

Workaround: Follow the steps in [Uploading files to cases on the Broadcom Support Portal](#).

When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client

If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back

to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

You see the same license product name displayed multiple times in the vSphere Client

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

Workaround: None

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

An external gateway firewall might block vSphere Pod traffic to cluster IPs

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to Cluster IPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see [vSphere Pod Traffic to ClusterIP Time-outs](#).

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API : `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag : `KeepAfterDeleteVM`.

The vpxd agent might fail after expanding the disk size of a virtual machine with a NVMe controller in a vCenter 8.0 Update 3 mixed environment

If you have ESXi hosts of version 8.0 Update 1 and earlier in your vCenter 8.0 Update 3 system, expanding the disk of a VM with a NVMe controller might cause the vpxd service to fail with `core.vpxd-worker coredump`.

Workaround: Upgrade your ESXi hosts of version 8.0 Update 1 and earlier to 8.0 Update 3. Alternatively, use the VMware Host Client to extend the disk size from **Edit settings** of the VM.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image

When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components** and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 3d Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 3d | 21 OCT 2024 | ISO Build 24322831
Check for additions and updates to these release notes.

What's New

- This release contains additional fixes that fully address CVE-2024-38812. For more information on this vulnerability and its impact on VMware by Broadcom products, see [VMSA-2024-0019.2](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 3c Release Notes](#)
- [VMware vCenter Server 8.0 Update 3b Release Notes](#)
- [VMware vCenter Server 8.0 Update 3a Release Notes](#)
- [VMware vCenter Server 8.0 Update 3 Release Notes](#)
- [VMware vCenter Server 8.0 Update 2d Release Notes](#)
- [VMware vCenter Server 8.0 Update 1e Release Notes](#)
- [VMware vCenter Server 8.0 Update 2c Release Notes](#)
- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 3d

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00400-24322831-patch-FP.iso
Build	24322831
Download Size	8556.0 MB
sha256checksum	799d65446086c77d1ae6be51ec6431283db48c03235dd518fc13da4a820b4f76
PRs fixed	N/A
CVEs	CVE-2024-38812

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Issues

This release contains additional fixes that fully address CVE-2024-38812. For more information on this vulnerability and its impact on VMware by Broadcom products, see [VMSA-2024-0019.2](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLS>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](#) and [94934](#).

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system

Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as The host returns `esxupdate` error codes: -1.

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:
Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails

In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

You see a "vc.health.error.dbjob3" warning after a vCenter upgrade

In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

Workaround: See [Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x](#) how to clear stats data, if not relevant anymore.

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"

Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists.`

Workaround: Delete all files under `/storage/log/vmware/analytics/stage`, `/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Verify following URLs and their respective statuses and follow KB 93526.

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
```

```
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
```

```
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
```

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
```

```

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.

```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```

root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback

```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 2:**Default VM Settings**

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

You cannot upload logs packages to a VMware service request by using the vSphere Client

You cannot upload logs packages to a VMware service request by navigating in the vSphere Client to **Administration > Support > Upload File to Service Request**. Instead, you see a link to a KB article that provides steps to upload log bundles in the Broadcom Support Portal. This is expected.

Workaround: Follow the steps in [Uploading files to cases on the Broadcom Support Portal](#).

When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client

If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

You see the same license product name displayed multiple times in the vSphere Client

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

Workaround: None

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

An external gateway firewall might block vSphere Pod traffic to clusterIPs

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to ClusterIPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see [vSphere Pod Traffic to ClusterIP Time-outs](#).

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag: `KeepAfterDeleteVM`.

The vpxd agent might fail after expanding the disk size of a virtual machine with a NVMe controller in a vCenter 8.0 Update 3 mixed environment

If you have ESXi hosts of version 8.0 Update 1 and earlier in your vCenter 8.0 Update 3 system, expanding the disk of a VM with a NVMe controller might cause the vpxd service to fail with `core.vpxd-worker coredump`.

Workaround: Upgrade your ESXi hosts of version 8.0 Update 1 and earlier to 8.0 Update 3. Alternatively, use the VMware Host Client to extend the disk size from **Edit settings** of the VM.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image

When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components** and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects

vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key /`

default KMS cluster not found. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 3c Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 3c | 09 OCT 2024 | ISO Build 24305161
Check for additions and updates to these release notes.

What's New

- This release resolves an issue with the vSphere Client becoming unresponsive when a session stays idle for more than 50 minutes. For more information, see [Resolved Issues](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 3b Release Notes](#)
- [VMware vCenter Server 8.0 Update 3a Release Notes](#)
- [VMware vCenter Server 8.0 Update 3 Release Notes](#)
- [VMware vCenter Server 8.0 Update 2d Release Notes](#)
- [VMware vCenter Server 8.0 Update 1e Release Notes](#)
- [VMware vCenter Server 8.0 Update 2c Release Notes](#)
- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 3c

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00300-24305161-patch-FP.iso
Build	24305161
Download Size	8556.0 MB
sha256checksum	fc44d794f02756296b7a9f1b0829c7459810c62beff8ab72cb11ca1f0311198a
PRs fixed	3439359
CVEs	N/A

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

vCenter Server and vSphere Client Issues

PR 3439359: If a vSphere Client session stays idle for more than 50 minutes, the vSphere Client becomes unresponsive

If you start a vSphere Client session and leave it idle for more than 50 minutes, the session becomes invalid and you cannot log in or log out, or use the vSphere Client from the same browser until you clear all cookies. The issue occurs due to an upgrade to Apache Tomcat 9.0.91 in vSphere 8.0 Update 3b, which changes the default value of the property `org.apache.catalina.connector.RECYCLE_FACADES` to `TRUE` from `FALSE`.

This issue is resolved in this release.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLS>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](https://kb.vmware.com/s/article/93526) and [94934](https://kb.vmware.com/s/article/94934).

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system

Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as The host returns `esxupdate` error codes: -1.

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:
Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails

In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

You see a "vc.health.error.dbjob3" warning after a vCenter upgrade

In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

Workaround: See [Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x](#) how to clear stats data, if not relevant anymore.

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"

Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists.`

Workaround: Delete all files under `/storage/log/vmware/analytics/stage`, `/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Verify following URLs and their respective statuses and follow KB 93526.

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```
root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback
```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API

version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.

2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 3:**Default VM Settings**

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

You cannot upload logs packages to a VMware service request by using the vSphere Client

You cannot upload logs packages to a VMware service request by navigating in the vSphere Client to **Administration > Support > Upload File to Service Request**. Instead, you see a link to a KB article that provides steps to upload log bundles in the Broadcom Support Portal. This is expected.

Workaround: Follow the steps in [Uploading files to cases on the Broadcom Support Portal](#).

When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client

If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

You see the same license product name displayed multiple times in the vSphere Client

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

Workaround: None

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

An external gateway firewall might block vSphere Pod traffic to clusterIPs

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to ClusterIPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see [vSphere Pod Traffic to ClusterIP Time-outs](#).

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag: `KeepAfterDeleteVM`.

The vpxd agent might fail after expanding the disk size of a virtual machine with a NVMe controller in a vCenter 8.0 Update 3 mixed environment

If you have ESXi hosts of version 8.0 Update 1 and earlier in your vCenter 8.0 Update 3 system, expanding the disk of a VM with a NVMe controller might cause the vpxd service to fail with `core.vpxd-worker coredump`.

Workaround: Upgrade your ESXi hosts of version 8.0 Update 1 and earlier to 8.0 Update 3. Alternatively, use the VMware Host Client to extend the disk size from **Edit settings** of the VM.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image

When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components** and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects

vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key /`

default KMS cluster not found. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 3b Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Product Support Notices](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 3b 17 SEP 2024 ISO Build 24262322 Check for additions and updates to these release notes.
--

What's New

- This release resolves CVE-2024-38812 and CVE-2024-38813. For more information on these vulnerabilities and their impact on VMware by Broadcom products, see [VMSA-2024-0019](#).
- VMware vCenter Server 8.0 Update 3b delivers bug fixes documented in the [Resolved Issues](#) section.
- For VMware vSphere with Tanzu updates, see [VMware vSphere IaaS Control Plane 8.0 Release Notes](#).
- For Photon OS updates, see [VMware vCenter Server Appliance Photon OS Security Patches](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 3a Release Notes](#)
- [VMware vCenter Server 8.0 Update 3 Release Notes](#)
- [VMware vCenter Server 8.0 Update 2d Release Notes](#)
- [VMware vCenter Server 8.0 Update 1e Release Notes](#)
- [VMware vCenter Server 8.0 Update 2c Release Notes](#)
- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)

- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Product Support Notices

Deprecation of support on legacy NVIDIA GPUs for 3D hardware acceleration with Virtual Shared Graphics Acceleration (vSGA): Starting with vSphere 8.0 Update 3b, support for 3D hardware acceleration by using vSGA on legacy NVIDIA GPUs is deprecated and will be removed in a future vSphere release. You can still use vSGA mode without 3D hardware acceleration or use hardware acceleration with legacy GPUs in vGPU mode. Alternatively, you can migrate to Intel GPUs for 3D hardware acceleration with vSGA support.

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 3b

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00200-24262322-patch-FP.iso
Build	24262322
Download Size	8556.0 MB
sha256checksum	952085ba6391500b6bfa9ea459d6964da5a25c3a36414ed8982a15bee5d08a77
PRs fixed	3417930, 3404486, 3404485, 3308624, 3210868, 2993755, 3405670, 3020259, 3413604, 3080596, 3416226, 3354058, 3377801, 3423070, 3413600, 3417512, 3415467, 3421685, 3412916, 3419046
CVEs	CVE-2024-38812, CVE-2024-38813

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

vSphere Client and vCenter Issues

PR 3417930: A fresh install of vCenter by using the CLI installer might fail due to a rare race condition

When using the CLI installer, a rare race condition because of the sequence in which vCenter services start and register might fail some tasks, such as the Datacenter API call, and you cannot create a datacenter. As a result, you might see a vCenter instance getting deployed and immediately deleted.

This issue is resolved in this release. The fix increases the number of retries and adds logging in case the vpxd service is currently unavailable to make sure that vCenter deployments complete.

PR 3404486: A date rounding issue leads to inaccurate aggregate values during the monthly rollup of performance stats data on vCenter

Due to a date rounding issue, if the partition index starts at a random hour instead of 00:00:00, the `stats_rollup2_proc_PostgreSQL.sql` stored procedure can cause incorrect aggregate values in the rollup processing of performance stats data on vCenter.

This issue is resolved in this release.

PR 3404485: vCenter performance statistics from yearly partitions are missing the aggregate functions MIN, MAX, and SUM

The vCenter stats rollup data in the yearly partition tables, such as `VPX_HIST_STAT4_*`, are missing the aggregate values for MIN, MAX, and SUM, and the stats rollup procedure might be incomplete, affecting the accuracy and completeness of the vCenter performance statistics.

This issue is resolved in this release.

PR 2993755: You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.x and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter by using the vSphere Client of an 8.x vCenter.

This issue is resolved in this release.

Miscellaneous Issues

PR 3405670: In the GUI installer you set Mac Address parameters that you don't see in the vSphere Client after deployment completes

When you use the GUI installer to perform an interactive deployment of a vCenter instance, you can set Mac Address parameters, but changes to these fields do not take effect in the vSphere Client after deployment completes. For example, if you set a value in the field **Mac Address Prefix**, when you navigate to **Advanced Settings** in the vSphere Client after the deployment, you see no entry for the `config.vpxd.macAllocScheme.prefixScheme.prefix` parameter and a value.

This issue is resolved in this release.

vSphere Lifecycle Manager Issues

PR 3423360: You cannot transition from vSphere Lifecycle Manager baselines to vSphere Lifecycle Manager images due to an NSX certificate error

When you transition the management of a cluster from vSphere Lifecycle Manager baselines to vSphere Lifecycle Manager images, the SSL connection of the cluster to NSX might fail. As a result, in the vSphere Client you see an error such as `Certificate too long` and the vSphere Lifecycle Manager workflow fails.

This issue is resolved in this release.

PR 3413604: While creating a desired state configuration for a cluster, you see the error "Task failed with Unknown Error: Certificate too long (max 2103) for SSL connection"

While setting a desired host configuration at a cluster level by using vSphere Configuration Profiles, if even a single stateless Auto Deploy ESXi host is active in the vCenter inventory, the task fails. In the vSphere Client, you see the error **Task failed with Unknown Error: Certificate too long (max 2103) for SSL connection**, followed by the vCenter machine SSL certificate. The issue occurs because booted stateless ESXi hosts in the vCenter inventory do not support long certificates.

This issue is resolved in this release.

High Availability and Fault Tolerance Issues

PR 3416226: The Fault Domain Manager (FDM) agent might fail to connect to a vCenter instance while in a thumbprint mode and you see no error

When you change the vCenter certificate mode to thumbprint for some reason, the FDM agent that vCenter deploys on ESX hosts might fail to perform some CA-based verifications, but you see no warning or more details on the error.

This issue is resolved in this release. The fix adds a vSphere HA event, `HostAgentErrorEvent`, which triggers when a vCenter certificate mode is set to thumbprint.

Installation, Upgrade, and Migration Issues

PR 3423070: vSphere vMotion tasks fail with an error for null network

In rare cases, the vSphere vMotion workflow that you run from the vSphere Client might fail to load data on the **Select networks** step due to a null pointer exception. In the **Select networks** page, you see an error such as `An error occurred while trying to load data and y is null` in the **Error Details** window. This could be a temporary outage as the issue depends on the network capacity of your environment. You can retry the operation by clicking the **Retry** button. The issue might also occur if a VM has multiple network adapters and one of them is connected to a distributed portgroup that is not available to the selected destination compute resource.

This issue is resolved in this release.

PR 3413600: Update to vCenter 8.0 Update 3 might fail at 95% with the error "Exception occurred in postInstallHook for wcp:Patch"

If the roles `CNS-Datastore` or `CNS-SEARCH-AND-SPBM` already exist in your vCenter system, during the `PostInstallHook` phase of an update to vCenter 8.0 Update 3, the process might try to recreate the roles, which causes a conflict and stops the update.

In the `Patchrunner.log`, you see an error such as:

```
Failed to apply patch roles_groups_users! Error: Role CNS-Datastore (id: 1090) not found in VC and Exception occurred in postInstallHook for wcp:Patch.
```

This issue is resolved in this release. The fix adds the descriptor `SUPERVISOR` to the roles when created by the `wcpsvc` service to differentiate from roles when created by a vSphere admin.

Internationalization Issues

PR 3417512: You cannot read localized csv files exported from the vSphere Client

When you export a csv file from the vSphere Client, for example from the Virtual Machines tab, in a locale such as Japanese, you might not see characters in the file properly displayed due to a missing byte-order mark (BOM).

This issue is resolved in this release.

PR 3415467: You cannot see snapshot details of virtual machines when using Traditional Chinese in the vSphere Client

In the vSphere Client, when you select Traditional Chinese and navigate to **VM > Snapshot**, you might see no data in the details pane. The issue occurs due to a missing date locale format.

This issue is resolved in this release.

vSphere vMotion Issues

PR 3421685: The cross-cluster migration of a virtual machine in the vSphere Client stops at the Select networks step

During a cross-cluster migration, when you select a cluster as the destination compute resource, a compatibility check at the **Select networks** step might block the migration workflow. The issue occurs if even a single ESXi host in the entire cluster fails the compatibility check.

This issue is resolved in this release.

Networking Issues

PR 3412916: You don't see newly created NSX Distributed Virtual port groups on some ESXi hosts

When multiple ESXi hosts join a new transport zone, a race condition in the vCenter host sync handler might cause some ESXi hosts to fail saving the new transport zones configuration. As a result, you don't see newly created NSX Distributed Virtual port groups in the new transport zone on such hosts.

This issue is resolved in this release.

vSAN Issues

PR 3419046: vSAN tasks fail with NotSupported exception

In vSAN HCI with datastore sharing environment (formerly HCI Mesh), a vSAN datastore that belongs to one vSAN cluster can be mounted on hosts of other clusters. Such hosts provide only compute and not storage, and might not even have vSAN active. In the Site Recovery Manager reprotect workflow, if a certain API call is made to a host with no active vSAN, the call ends with a `NotSupported` exception. If such a call is made to a vSAN-enabled compute-only host, then you see no error, but the call does not return any results.

This issue is resolved in this release. The fix makes sure API calls for the SRM reprotect workflow to pick up only hosts with active vSAN.

Known Issues

vCenter Server and vSphere Client Issues

You cannot upload logs packages to a VMware service request by using the vSphere Client

You cannot upload logs packages to a VMware service request by navigating in the vSphere Client to **Administration > Support > Upload File to Service Request**. Instead, you see a link to a KB article that provides steps to upload log bundles in the Broadcom Support Portal. This is expected.

Workaround: Follow the steps in [Uploading files to cases on the Broadcom Support Portal](#).

Virtual Machines Management Issues

The vpxd agent might fail after expanding the disk size of a virtual machine with a NVMe controller in a vCenter 8.0 Update 3 mixed environment

If you have ESXi hosts of version 8.0 Update 1 and earlier in your vCenter 8.0 Update 3 system, expanding the disk of a VM with a NVMe controller might cause the vpxd service to fail with `core.vpxd-worker coredump`.

Workaround: Upgrade your ESXi hosts of version 8.0 Update 1 and earlier to 8.0 Update 3. Alternatively, use the VMware Host Client to extend the disk size from **Edit settings** of the VM.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLs>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](https://kb.vmware.com/s/article/93526) and [94934](https://kb.vmware.com/s/article/94934).

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system

Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as The host returns `esxupdate` error codes: -1.

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:
Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails

In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

You see a "vc.health.error.dbjob3" warning after a vCenter upgrade

In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

Workaround: See [Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x](#) how to clear stats data, if not relevant anymore.

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"

Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists.`

Workaround: Delete all files under `/storage/log/vmware/analytics/stage`, `/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Verify following URLs and their respective statuses and follow KB 93526.
```

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
```

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
```

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
```

```
2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```
root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback
```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 4:**Default VM Settings**

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client

If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

You see the same license product name displayed multiple times in the vSphere Client

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

Workaround: None

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

An external gateway firewall might block vSphere Pod traffic to clusterIPs

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to ClusterIPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see [vSphere Pod Traffic to ClusterIP Time-outs](#).

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag: `KeepAfterDeleteVM`.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image

When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components** and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 3a Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 3a | 18 JUL 2024 | ISO Build 24091160
Check for additions and updates to these release notes.

What's New

This release delivers critical bug fixes. For more information, see the [Resolved Issues](#) section.

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 3 Release Notes](#)
- [VMware vCenter Server 8.0 Update 2d Release Notes](#)
- [VMware vCenter Server 8.0 Update 1e Release Notes](#)
- [VMware vCenter Server 8.0 Update 2c Release Notes](#)
- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 3a

Product Patch for vCenter Server containing VMware software fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00100-24091160-patch-FP.iso
Build	24091160
Download Size	8507.3 MB
sha256checksum	2b65ce11d3d56a1acba21e52d89202ed51b33f6a5bc24381c425f92e8445a85f

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

1. Attach the `VMware-vCenter-Server-Appliance-8.0.3.00100-24091160-patch-FP.iso` file to the vCenter Server CD or DVD drive.
2. Log in to the appliance shell as a user with super administrative privileges (for example, `root`) and run the following commands:
 - To stage the ISO:


```
software-packages stage --iso
```
 - To see the staged content:


```
software-packages list --staged
```
 - To install the staged rpms:


```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Installation, Upgrade, and Migration Issues

PR 3406386: Updates to vCenter 8.0 Update 3 might fail due to unnecessary certificates in the VMware Endpoint Certificate Store (VECS)

Due to a rare issue with VECS certificates that are no longer necessary, updates to vCenter 8.0 Update 3 might fail with an error such as `Pre-install failed for vmidentity:Expand` in the vSphere Client. In the `/var/log/vmware/applmgmt/Patchrunner.log`, you see messages such as:

```
2024-07-08T11:06:39.692Z vmidentity:Expand ERROR vmidentity Something went wrong while
reading certs from TRUSTED_ROOTS or deleting ssoserver cert:
```

This issue is resolved in this release.

PR 3408378: If you do not open port 9087 in your firewall between ESXi hosts and vCenter, compliance checks and vSphere HA might fail after vCenter update to 8.0 Update 3

Starting from 8.0 Update 3, the vSphere Lifecycle Manager downloads updates for ESXi hosts by a HTTPS connection to the vCenter instance on port 9087. If you do not open port 9087 in your firewall between ESXi hosts and vCenter, you might see compliance check errors. For example, in the `lifecycle.log` you see messages such as:

```
<Timestamp> In(14) lifecycle[2112988]: Downloader:373 Opening https://<VC-
FQDN>:9087/vum/repository/hostupdate/__micro-depot__vendor-DEL__DEL-ESXi-8.0-Addon-
cumulative_metadata__index__.xml for download
```

```
<Timestamp> Wa(12) lifecycle[2112988]: Downloader:210 Download failed: <urlopen error
timed out>, 9 retry left...
```

In addition, after updating vCenter to 8.0 Update 3, vSphere HA might fail to start on all ESXi hosts with messages such as:

```
An error occurred when vCenter Server attempted to initialize the vSphere HA Agent running
on the host.
```

HA Agent Unreachable - The vSphere HA Agent on the host cannot be reached.

Cannot complete the configuration of the vSphere HA agent on the host. Applying HA VIBs on the cluster encountered failure.

A general system error occurred: Installing HA components failed on the cluster: domain-<ID>.

Cannot find vSphere HA master agent.

This issue is resolved in this release. The fix restores functionality on port 9084.

Known Issues

Installation, Upgrade, and Migration Issues

Update to vCenter 8.0 Update 3 might fail at 95% with the error "Exception occurred in postInstallHook for wcp:Patch"

If the roles `CNS-Datastore` or `CNS-SEARCH-AND-SPBM` already exist in your vCenter system, during the `PostInstallHook` phase of an update to vCenter 8.0 Update 3, the process might try to recreate the roles, which causes a conflict and stops the update.

In the `Patchrunner.log`, you see an error such as:

```
Failed to apply patch roles_groups_users! Error: Role CNS-Datastore (id: 1090) not found in VC and Exception occurred in postInstallHook for wcp:Patch.
```

Workaround: See KB [371524](#).

Upgrades from ESXi 7.0.x and 8.0.x might fail because the upgrade prechecks report that a vFAT partition was not unmounted

Starting with vSphere 8.0 Update 3, the upgrade precheck includes a vFAT partition validation step, `dosfsck`, because corrupted vFAT partitions might lead to unexpected behavior. This issue is specific to ESXi hosts, where another OS than ESXi has mounted the vFAT partitions and adds a bit that ESXi does not use. As a result of the check, in the vSphere Client you see an error such as `A problem with one or more vFAT bootbank partitions was detected`. Please refer to KB 91136 and run `dosfsck` on bootbank partitions.

Workaround: To remove the unnecessary bit, follow the steps in KB [345227](#).

If the vCenter certificate mode is set to thumbprint, vSphere HA clusters might fail to configure after update to vCenter 8.0 Update 3

If for some reason the advanced setting `vpzd.certmgmt.mode` in a vCenter instance is set to `thumbprint`, which is recommended only as a fallback option, the primary vSphere HA host might remain active but other vSphere HA clusters fail to configure after update to vCenter 8.0 Update 3.

In the `fdm.log`, you see messages such as:

```
Failed to SSL handshake; SSL(<io_obj p:0x0000002e86fda5f0, h:30, <TCP '10.10.10.10 : 8111'>, <TCP '10.10.10.11 : 11911'>>), e: 167773208(tlsv1 alert unknown ca (SSL routines)), duration: 5msec or Er(163) Fdm[1151234]:

--> The remote host certificate has these problems: Er(163) Fdm[1151234]: --> Er(163) Fdm[1151234]: --> * unable to get local issuer certificate)
```

Workaround: Change the `vpzd` certificate mode to `vmca` or `custom`, depending on your environment. For more information, see [Change the ESXi Certificate Mode](#) and KB [372329](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLs>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](#) and [94934](#).

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system

Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as The host returns `esxupdate` error codes: -1.

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:
Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails

In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

You see a "vc.health.error.dbjob3" warning after a vCenter upgrade

In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

Workaround: See [Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x](#) how to clear stats data, if not relevant anymore.

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"

Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists.`

Workaround: Delete all files under `/storage/log/vmware/analytics/stage`, `/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Verify following URLs and their respective statuses and follow KB 93526.
```

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
```

2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed. Please collect the VC support bundle.

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```
root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback
```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service. . This is expected behavior.`

Workaround: None

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 5:

Default VM Settings

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client

If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

You see the same license product name displayed multiple times in the vSphere Client

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

Workaround: None

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

An external gateway firewall might block vSphere Pod traffic to clusterIPs

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to ClusterIPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see [vSphere Pod Traffic to ClusterIP Time-outs](#).

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag: `KeepAfterDeleteVM`.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image

When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components** and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

vSphere Cluster Service Issues

In the vSphere Client, you see the vSphere HA status of Embedded vSphere Cluster Service VMs as Unprotected

vSphere 8.0 Update 3 introduces a redesign of vCLS to Embedded vCLS, which utilizes vSphere Pod technology. Deployment and lifecycle of such VMs are now managed within ESXi and are no longer managed by the vSphere ESX Agent Manager (EAM). On the **Summary** tab in the vSphere Client, you see the vSphere HA status of Embedded vSphere Cluster Service VMs as Unprotected, but this is expected, because vSphere HA does not manage Embedded vSphere Cluster Service VMs.

Workaround: Ignore the information in the vSphere HA card on the **Summary** tab in the vSphere Client. For more information, see [vSphere Cluster Services](#).

VMware vCenter Server 8.0 Update 3 Release Notes

This document contains the following sections

- [Introduction](#)
- [General Availability](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Product Support Notices](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 3 | 25 JUN 2024 | GA ISO Build 24022515
Check for additions and updates to these release notes.

General Availability

This vCenter Server 8.0 Update 3 release is a General Availability (GA) designation. For more information on the vSphere 8.0 IA/GA Release Model of vSphere Update releases, see [The vSphere 8 Release Model Evolves](#).

What's New

This release resolves CVE-2024-37087. For more information on this vulnerability and its impact on VMware products, see [VMSA-2024-0013](#).

For overview of the new features in vSphere 8.0 Update 3, see [What's New in vSphere 8 Update 3](#) and the following release notes:

- [vSphere IaaS Control Plane \(former VMware vSphere With Tanzu\)](#)
- [Tanzu Kubernetes releases](#)
- [VMware vSAN](#)
- [VMware Host Client](#)
- [VMware OVF Tool](#)

The list of new features and enhancements that follows adds some of the key highlights for vSphere 8.0 Update 3:

DPU/SmartNIC

- **High Availability with VMware vSphere Distributed Services Engine:** Starting with ESXi 8.0 Update 3, vSphere Distributed Services Engine adds support for 2 data processing units (DPUs) to provide high availability or increase offload capacity per ESXi host. Dual-DPU systems can use NVIDIA or Pensando devices. In ESXi 8.0 Update 3, dual-DPU systems are supported by Lenovo server designs. For more information, see [High Availability with VMware vSphere Distributed Services Engine](#).

vSphere IaaS control plane

- **Support for running vSphere IaaS control plane on vSAN stretched clusters:** vSphere 8.0 Update 3 adds support for vSphere IaaS control plane, formerly known as vSphere with Tanzu, on vSAN stretched clusters. This capability is available only for greenfield installations of the IaaS control plane on stretched vSAN clusters. If you already have IaaS control plane or vSphere with Tanzu plus Content Library running, you need to redeploy it before you use the vSAN Stretched Cluster Storage Policies.
Before configuring this capability, read [Running vSphere IaaS Control Plane on vSAN Stretched Cluster](#) and follow content on <https://core.vmware.com>.
- **IaaS Supervisor Control Planes backup and restore from the vCenter Server Management Interface:** Starting with vCenter 8.0 Update 3, you can record the state of the IaaS Supervisors Control Planes in a vCenter system as part of the vCenter file-based backup. While configuring a backup schedule in the vCenter Server Management Interface, you can select the **Supervisors Control Plane** checkbox to activate the backup and restore of a Supervisor in the case of any disaster scenario. For more information, see [Backup and Restore the Supervisor Control Plane](#).

Virtual Machine Management

- **New Virtual Machine Compute Policy for Best Effort Virtual Machine Evacuation:** vSphere 8.0 Update 3 adds a compute policy for best effort evacuation of virtual machines on ESXi hosts that are entering maintenance mode. When the host enters maintenance mode, all VMs are shut down. If shutdown fails, the VMs are powered off. If power off fails, you need to evacuate the VMs. With the new capability, when the VMs are in a powered-off state, vCenter attempts to power them on every few minutes on the best available ESXi host at the time. This policy overrides any DRS overrides set at the VM level, and the hosts on which the VMs are powered on might be different from the original host.

vSphere Cluster Service

- **Introducing Embedded vSphere Cluster Service (vCLS):** vSphere 8.0 Update 3 introduces a redesign of vCLS to Embedded vCLS, which utilizes vSphere Pod technology. Deployment and lifecycle of these VMs are managed within ESXi and are no longer managed by the vSphere ESX Agent Manager (EAM). Earlier versions of vCLS are termed External vCLS. Issues previously encountered with External vCLS are resolved in this release of Embedded vCLS. While existing API compatibility is preserved, minor modifications to customer scripts or automation tools might be necessary. Other products and solutions that have defined business logic around External vCLS might not work with Embedded vCLS. See individual product documentation to understand the interoperability support and impact. For more information, see [vSphere Cluster Services](#) and content available from vSphere Technical Marketing on the [Broadcom Support Portal](#).

vSAN

- **vSAN add-on licenses based on capacity per tebibyte (TiB):** With vSphere 8.0 Update 3, you can license your use of vSAN storage based on TiB capacity. The new capacity-based license replaces the core and CPU-based licenses. For more information on capacity reporting and licensing in vSAN, see [License Requirements](#) and [Counting Cores for VMware Cloud Foundation and vSphere Foundation and TiBs for vSAN](#).

Security

- **TLS 1.3 and 1.2 support by using TLS profiles:** Starting with vSphere 8.0 Update 3, you can use TLS profiles to simplify the configuration of TLS parameters and improve supportability in your vSphere system. With vSphere 8.0 Update 3, you get a default TLS profile on ESXi and vCenter Server hosts, COMPATIBLE, which supports TLS 1.3 and some TLS 1.2 connections. For more information, see [vSphere TLS Configuration](#).

- **PingFederate Identity Provider for vSphere:** With vSphere 8.0 Update 3, you can configure PingFederate as an external identity provider in your vSphere system. For more information, see [Configuring vCenter Server Identity Provider for PingFederate](#).

Storage/Memory

- **Memory Tiering:** vSphere 8.0 Update 3 launches in tech preview the Memory Tiering capability, which allows you to use NVMe devices that you can add locally to an ESXi host as tiered memory. Memory tiering over NVMe optimizes memory utilization by directing VM memory allocations to either NVMe devices or faster dynamic random access memory (DRAM) in the host. This allows you to increase your memory footprint and workload capacity, while reducing the total cost of ownership (TCO). For more details on the tech preview, see KB [95944](#).
- **Fabric Notification support for SAN clusters:** ESXi 8.0 Update 3 introduces support for Fabric Performance Impact Notifications Link Integrity (FPIN-LI). With FPIN-LI, the vSphere infrastructure layer can manage notifications from SAN switches or targets, identifying degraded SAN links and ensuring only healthy paths are used for storage devices. FPIN can also notify ESXi hosts for storage link congestion and errors.
- **Support for space reclamation requests from guest operating systems on NVMe-backed vSphere Virtual Volumes datastores and Config-vVol:** ESXi 8.0 Update 3 adds support for automatic space reclamation requests from guest operating systems on NVMe-backed vSphere Virtual Volumes datastores. ESXi 8.0 Update 3 also adds support for both command line-based and automatic unmap for vSphere Virtual Volumes objects of type Config-vVol, formatted with VMFS-6. For more information, see [Reclaim Space on the vSphere Virtual Volumes Datastores](#).
- **Manage the UNMAP load from ESXi hosts at a VMFS datastore level:** Starting with ESXi 8.0 Update 3, you can control the unmap load at datastore level to avoid time delays from space reclamation and reduce overall unmap load on the arrays in your environment. For more information, see [Space Reclamation on vSphere VMFS Datastores](#).
- **Windows Server Failover Clustering (WSFC) enhancements on vSphere Virtual Volumes:** vSphere 8.0 Update 3 adds support for a WSFC solution for NVMe-backed disks on vSphere Virtual Volumes. This capability allows NVMe reservations support in NVMe/TCP environments apart from Fibre Channel support for WSFC on vSphere Virtual Volumes. Virtual NVMe (vNVME) controllers are supported as the frontend for WSFC with NVMe-backed disks, not with SCSI-backed disks. For more information, see [VMware vSphere® Virtual Volumes Support for WSFC](#).
- **Support for active-active vSphere Metro Storage Cluster (vMSC) with vSphere Virtual Volumes:** vSphere 8.0 Update 3 introduces a new version of the VMware vSphere Storage APIs for Array Integration (VASA) to add support to active-active stretched storage clusters with vSphere Virtual Volumes, with active-active deployment topologies for SCSI block access between two sites. VASA version 6 includes new architecture and design for VASA Provider High Availability support for both stretched and non-stretched storage clusters. For more information, see [Using Stretched Storage Clustering with Virtual Volumes](#).
- **VMkernel port binding for NFS v4.1 datastores:** With vSphere 8.0 Update 3, you can bind an NFS 4.1 connection to a specific VM kernel adapter. If you use multipathing, you can provide multiple vmknics for each connection to ensure path isolation and security by directing NFS traffic across a specified subnet/VLAN, so that the NFS traffic does not use other vmknics. For more information, see [Configure VMkernel Binding for NFS 4.1 Datastores](#).
- **Support for nConnect for NFS v4.1 datastores:** ESXi 8.0 Update 3 adds support to multiple TCP connections for a NFS v4.1 volume, also referred to as nConnect. For NFS v4.1, multiple TCP connections are created for a single session that many datastores can share in parallel. It can be configured by using either the vSphere API or ESXCLI directly on an ESXi host. For more information, see [Configure Multiple TCP Connections for NFS](#).
- **Reduced time to inflate VMFS disks:** With vSphere 8.0 Update 3, a new VMFS API allows you to inflate a thin-provisioned disk to eagerzeroedthick (EZT) while the disk is in use and up to 10 times faster than previous alternative methods. During the inflation, all blocks are fully allocated and zeroed upfront to allow faster run-time performance. For more information, see [Virtual Disk Options](#).
- **Improved resiliency against memory corruption on RAM-heavy ESXi hosts:** vSphere 8.0 Update 3 adds proactive measures to prevent memory errors in systems with VMs of more than 1TB that might bring down an entire ESXi host and increase application downtime.
- **Advanced setting to block deletion and removal of disks for VMs with snapshots:** ESX 8.0 Update 3 adds a per-host advanced option `blockDiskRemoveIfSnapshot` to prevent the removal of disks from a VM that has snapshots, even if you choose to delete files, which might lead to orphaned disks. For more information, see VMware knowledge base article [94545](#).

- **Hardware accelerated move (clone operation) support on NVMe devices:** vSphere 8.0 Update 3 supports NVMe copy command for hardware accelerated move (also called clone blocks or copy offload) across or within NVMe namespaces that belong to the same NVMe subsystem.
- **Granular monitoring for VASA provider accessibility and certification authentication status on ESXi host level:** With vSphere 8.0 Update 3, you can monitor connection status and ESXi authentication with VASA storage providers from the vSphere Client and re-authenticate hosts as necessary. For more information, see [Reauthenticate VASA Client in ESXi](#).

GuestOS

- **Guest customization supports RHEL NetworkManager keyfile format:** In vSphere 8.0 Update 3, guest customization adds support for RHEL NetworkManager keyfile format and you can store network configuration in both keyfile and ifcfg format.

Drivers/Network

- **Support for Fibre Channel Extended Link Services (FC-ELS):** With vSphere 8.0 Update 3, you can use the command `esxcli storage fpin info set -e= <true/false>` to activate or deactivate the Fabric Performance Impact Notification (FPIN). The command saves the FPIN activation to both ConfigStore and the VMkernel System Interface Shell and persists across ESXi reboots. This is enabled by both Broadcom's `lpfc` and Marvell's `qlnativefc` drivers.
- **Unified Enhanced Networking Stack (UENS) driver for the Elastic Network Adapter (ENA):** vSphere 8.0 Update 3 adds a UENS driver for ENA, which provides connectivity to the AWS underlay Virtual Private Cloud (VPC).
- **Overlay Filters Supporting Tunnel End Point (TEP):** vSphere 8.0 Update 3 enhances the `i40en`, `qedentv`, and `sfvmk` NIC drivers by adding capabilities that expose overlay filters, supporting the TEP functionality.
- **Broadcom Driver Updates:**
 - **Broadcom `bnxtnet` driver:** Adds support for NetQ RSS to facilitate future Unified Enhanced Networking Stack (UENS) support.
- **Intel Driver Updates:**
 - **Intel `i40en` driver:** The number of queues per RSS engine is up from 4 to 8, in ENS mode it supports up to 16 queues.
 - **Intel `icen` driver:** Adds RoCE support on 1000GbE NIC.
 - **Intel `irdman` driver:** Adds RoCE support on 1000GbE NIC.
- **Marvell Driver Updates:**
 - **Marvell `qedentv` driver:** The max number of queues for the Default RSS engine is up from 4 to 16 in ENS mode.
- **Mellanox Driver Updates:**
 - **Mellanox `nmlx5` driver:** Adds support for hardware Large Receive Offload (LRO) in Enhanced Data Path mode to increase inbound throughput of high-bandwidth network connections by reducing CPU overhead.
 - Adds support for dual-DPU servers.
- **Pensando Driver Updates:**
 - **Pensando `ionic_en` driver:** Adds support for dual-DPU servers.
- **Routine Inbox Driver Updates and Bug Fixes**
 - Broadcom `bcm_mpi3`
 - Broadcom `lpfc`
 - Cisco `nfnc`
 - Marvell `qlnativefc`
 - Microchip `smartpqi`

CPU

- **PCIe hot plug is updated for server platforms utilizing newer generation AMD Genoa and Intel Sapphire Rapids CPUs:** Starting with vSphere 8.0 Update 3, kernel hot plug is supported for newer generation CPUs such as AMD Genoa and Intel Sapphire Rapids.

- **Support for Intel Xeon Max Series processors with integrated High Bandwidth Memory (HBM):** vSphere 8.0 Update 3 adds support for Intel Xeon Max Series processors (formerly with code name Sapphire Rapids HBM) with 64 GB of integrated HBM, aimed to enhance performance for workloads like high performance computing apps, artificial intelligence (AI), and machine learning (ML).
- **CPU C-State Power virtualization:** With vSphere 8.0 Update 3, for use cases such as Virtualized Radio Access Network (vRAN) workloads, you can configure and control the C-State power of the physical CPUs dedicated to vRAN VMs from the vSphere Client.
- **Cluster-wide option to retain virtual NUMA topology:** vSphere 8.0 Update 3 adds a setting to retain preconfigured vNUMA topology even if the VM moves, allowing for better NUMA topology tuning for VMs across all the hosts in the cluster. This is the advanced vCenter Server setting `VPXD_PersistVnuma` to keep the virtual NUMA topology on a cluster level under **Configure > Advanced Settings** in the vSphere Client.

Analytics and Metrics

- **vSphere green metrics with Running Average Power Limit (RAPL) technology:** Starting with vSphere 8.0 Update 3, in the **Advanced Performance Charts** in the vSphere Client you can see RAPL data on ESXi hosts that normally do not report individual subsystem power consumption, such as CPU and Memory, but only general host-level power consumption. With individual subsystem power consumption reports, you can plan on a more granular level to match your power and cooling budget.
- **Set VM log levels without powering off the VMs:** With vSphere 8.0 Update 3, you can set log levels between `VMW_LOG_TRIVIA` and `VMW_LOG_DEBUG_3` to avoid log spew in a healthy running VM without a power cycle by using the `SetLogLevel` service in the vAPI infrastructure.

GPU

- **Support for switching between Time Sliced and Multi-Instance GPU (MIG) modes for NVIDIA virtual GPUs:** Starting with vSphere 8.0 Update 3, you do not need to reboot an ESXi host to switch between time sliced and MIG modes for NVIDIA virtual GPUs. vGPU VMs can automatically set the correct device mode according to their vGPU type.
- **Zero-copy support for vGPUs to enhance vSphere vMotion and vSphere DRS tasks:** vSphere 8.0 Update 3 adds zero-copy support for vGPUs to enhance vSphere vMotion and vSphere DRS tasks by utilizing throughput of up to 100 Gbps.
- **Support for heterogeneous vGPU profiles on physical GPUs:** Starting with vSphere 8.0 Update 3, you can set vGPU profiles with different types or sizes on a single physical GPU to achieve greater flexibility with vGPU workloads and better utilization of GPU devices. For more information, see [Configuring vGPU Size](#).

vSphere Lifecycle Management

- **Support for parallel hardware and firmware upgrade with vSphere Lifecycle Manager:** With vCenter Server 8.0 Update 3, you can run parallel hardware and firmware remediation by using an integration between the vSphere Lifecycle Manager and the Hardware Support Manager.
- **Non-disruptive certificates for ESXi:** With vCenter Server 8.0 Update 3, you can renew or replace ESXi certificates non-disruptively without the need for rebooting ESXi or setting up the maintenance window.
- **VMware Photon™ 5.0 support for Update Manager Download Service (UMDS):** vSphere 8.0 Update 3 adds VMware Photon™ 5.0 to the supported Linux-based operating systems for installing UMDS. For more information, see [Installing UMDS](#).
- **Additional lifecycle management capabilities for standalone ESXi hosts:** Starting with vSphere 8.0 Update 3, you can add a standalone host to a data center or folder by importing an image from another ESXi host in the vCenter Server inventory or by using the current image on the host. For more information, see [Managing Standalone ESXi Hosts with vSphere Lifecycle Manager Images](#). When you move a host out of a cluster that you manage with a vSphere Lifecycle Manager image to a data center or a folder, the host becomes standalone and can retain the image from the cluster. For more information, see [Specifics of the Transitioning Workflow](#).
- **Lifecycle management of standalone ESXi hosts with VMware NSX:** Starting with vSphere 8.0 Update 3, NSX Manager and vSphere Lifecycle Manager work together to coordinate remediation of standalone ESXi hosts with

VMware NSX. For more information, see [Using vSphere Lifecycle Manager Images for Standalone Hosts with NSX 4.2 and later](#).

- **Convert baseline-managed clusters to clusters managed by a single vSphere Lifecycle Manager image:** Starting with vSphere 8.0 Update 3, to start managing a cluster with a single vSphere Lifecycle Manager image, you can use the image installed on one of the ESXi hosts inside the cluster managed with baselines. For more information, see [Use an Image from a Host in the Cluster](#).
- **Patch VMX-related security vulnerabilities without any disruption to your workloads:** Starting with vSphere 8.0 Update 3, you can use the **Live Patch** capability to apply VMX-related security patches and bug fixes to ESXi hosts in a cluster managed with a vSphere Lifecycle Manager image. A pre-check prior to remediation lists all suitable hosts in a cluster. After you activate **LivePatch** in the remediation settings, qualified hosts in a cluster do not require maintenance mode, or reboot, or VM migration during the update procedure.
- **Customize vSphere Lifecycle Manager desired state images:** Starting with vSphere 8.0 Update 3, you can remove the Host Client and VMware Tools components from the base image, remove unnecessary drivers from vendor add-ons and components, and override existing drivers in a desired image. For more information, see [Edit a vSphere Lifecycle Manager Image](#).
- **Support for dual DPUs with vSphere Lifecycle Manager:** Starting with vSphere 8.0 Update 3, you can install dual DPUs on ESXi hosts and use the vSphere Lifecycle Manager workflow to upgrade the dual DPU system.
- **Extended support for vSphere Configuration Profiles (VCP):** With vSphere 8.0 Update 3, you can use VCP with the following new capabilities:
 - **Support to baseline-managed clusters (formerly referred to as VUM clusters):** Having an image-managed cluster is no longer a prerequisite for using VCP. You can use VCP to configure either baseline-managed clusters or image-managed clusters.
 - **Support for vSphere Distributed Switch (VDS):** VCP is fully integrated with VDS and supports drift detection and remediation of VDS configurations at a cluster level.
 - **Firewall ruleset management:** You can manage custom firewall rules at a cluster level by using VCP.
 - **ESXi Lockdown Mode:** vSphere admins can use the VCP desired configuration document to enforce Lockdown Mode on all hosts in a cluster.
 - **Support for SNMP and PCI device configurations:** You can manage SNMP and PCI devices at a cluster level by using VCP.

For more information, see [Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level](#).

vSphere Client and vCenter

- **Warning on the maximum number of remote https connections for vCenter:** Starting with vSphere 8.0 Update 3, to prevent a vCenter system to become unresponsive due to exceeding the limit of 2048 https connections, you will see HTTP error code **503 Service Unavailable** and **x-envoy-local-overloaded: true** in the response headers.
- **Merging the vSAN Management SDK with the Python SDK for the VMware vSphere API:** Starting with vSphere 8.0 Update 3, the [vSAN Management SDK for Python](#) is integrated into the Python SDK for the VMware vSphere API ([pyVmomi](#)). From the Python Package Index ([PyPI](#)), you can download a single package to manage vSAN, vCenter, and ESXi. This integration streamlines the discovery and installation process and enables automated pipelines instead of the series of manual steps previously.
- **vCenter Universally Unique Identifier (VC_UUID) field in the vSphere Client:** With vSphere 8.0 Update 3, the property table under **VMs > Virtual Machines** in the vSphere Client includes a new column that contains the VC_UUID. This identifier is automatically assigned to every virtual machine within a vCenter instance. The VC_UUID field helps clarify the correlation between the VM's UUID displayed on the switch's fabric and the VM name in vCenter.
- **Default HTTP response compression:** vSphere 8.0 Update 3 adds support for HTTP response compression by default on the edge proxies of vCenter and ESXi hosts, management traffic between vCenter and ESXi hosts, and for outgoing HTTP requests on sidecar proxies. HTTP response compression reduces the HTTP traffic in a vCenter Server system, shrinks page load time and speeds up API operations. You can deactivate HTTP response compression if required, but the capability does not require any changes in your environment and is backward compatible.

- **Remove restrictions on virtual machine operations from the vSphere Client:** Starting with vCenter Server 8.0 Update 3, in the vSphere Client under **Virtual Machine > Configure > Disabled Methods** you can remove restrictions on virtual machine operations such as migration. For more information, see VMware knowledge base article [2044369](#).
- **Unified Management and Automation API Sessions:** Starting with vCenter 8.0 Update 3, you can combine vSphere Management API (VIM) and vSphere Automation API (vAPI) sessions, effectively unifying authentication across SOAP and REST vCenter endpoints. For more information see [Unified Management and Automation Session](#).
- **Migration of SPBM, SMS, EAM, and VLSM APIs to HTTP/JSON-based wire protocol:** Starting with vCenter Server 8.0 Update 3, the following 4 SOAP/XML service interfaces migrate to a HTTP/JSON-based wire protocol, providing OpenAPI 3.0 specifications for each, and integration with the [REST API](#) documentation:
 - The Storage Policy (SPBM) API, which simplifies the task of matching available storage to virtual machines.
 - VMware vCenter Storage Monitoring Service (SMS), which facilitates access to all vCenter storage information associated with VMware vCenter servers.
 - vSphere ESX Agent Manager API (EAM), which gives access to the objects that manage, monitor, and control lifecycle operations in vSphere.
 - Virtual Storage Lifecycle Management (VSLM) API that you use to manage first class disks (FCD).
 For more information, see [What's New in vSphere Automation: vSphere APIs, JSON, OpenAPI](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 2d Release Notes](#)
- [VMware vCenter Server 8.0 Update 1e Release Notes](#)
- [VMware vCenter Server 8.0 Update 2c Release Notes](#)
- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Product Support Notices

Deprecation of VMware vShield Endpoint: Starting with vSphere 8.0 Update 3, VMware vShield Endpoint is deprecated and will be removed in a future major release. vShield Endpoint is supported until the end of General Support for vSphere 8.x. Work with your vendors for agent-based solutions.

Deprecation of APIs for vCenter for Windows to Linux migration:

Starting with vSphere 8.0 Update 3, APIs from vSphere 6.7.x for migrating vCenter for Windows to Linux during upgrade are deprecated and will be removed in a future vSphere release.

Deprecation of VMkernel API (vmkapi) version 2.5:

Starting with vSphere 8.0 Update 3, vmkapi version 2.5 is deprecated and will be removed in a future major release. This requires an update of any third-party component released for vSphere 6.7.

Removal of Integrated Windows Authentication (IWA)

vSphere 8.0 Update 3 is the final release to support Integrated Windows Authentication. IWA was deprecated in vSphere 7.0 and will be removed in the next major release. To ensure continued secure access, migrate from IWA to Active Directory over LDAPS or to Identity Federation with Multi-Factor Authentication. For more information, see [vSphere Authentication with vCenter Single Sign-On](#) and [Deprecation of Integrated Windows Authentication](#).

Removal of VMware Enhanced Authentication Plug-in (EAP):

Starting with vSphere 8.0 Update 3, you cannot use EAP to log in to a vCenter system by using the vSphere Client. For more information, see [Removing the deprecated VMware Enhanced Authentication Plug-in \(EAP\) to address CVE-2024-22245 and CVE-2024-22250](#).

Deprecation of the vCenter Service Lifecycle Management API:

vSphere 8.0 Update 3 is deprecating the use of vCenter Service Lifecycle Management (vmonapi service) API and the service is not active by default, you must manually activate it. The service will be removed in a future release. For more information, see VMware knowledge base article [80775](#).

Removal of the internal runtime option execInstalledOnly:

Starting with ESXi 8.0 Update 3, the internal runtime option that deactivates the security option execInstalledOnly is deprecated and will be removed in the next major release. The boot option execInstalledOnly, which helps protect hosts against ransomware attacks, will be activated on ESXi hosts by default in the next major release.

Deprecation of Storage DRS Load Balancer and Storage I/O Control (SIOC):

The Storage DRS (SDRS) I/O Load Balancer, SDRS I/O Reservations-based load balancer, and vSphere Storage I/O Control Components will be deprecated in a future vSphere release. Existing 8.x and 7.x releases will continue to support this functionality.

The deprecation affects I/O latency-based load balancing and I/O reservations-based load balancing among datastores within a Storage DRS datastore cluster. In addition, enabling of SIOC on a datastore and setting of Reservations and Shares by using SPBM Storage policies are also being deprecated.

Storage DRS Initial placement and load balancing based on space constraints and SPBM Storage Policy settings for limits are not affected by the deprecation.

Deprecation of vSphere Trust Authority (vTA):

Starting with vSphere 8.0 Update 3, vSphere Trust Authority is deprecated. vSphere continues to offer advanced workload attestation in its baseline functionality.

Removal of Patch Manager APIs:

Patch Manager APIs are supported in vSphere 8.x, but support will discontinue in a future release of vSphere. Instead of Patch Manager APIs, you can use the latest vSphere APIs, documented in the [vSphere API automation reference guide](#).

Deprecation of locales:

Beginning with the next major release, we will be reducing the number of supported localization languages. The three supported languages will be:

- Japanese

- Spanish
- French

The following languages will no longer be supported:

- Italian, German, Brazilian Portuguese, Traditional Chinese, Korean, Simplified Chinese

Impact:

- Users who have been using the deprecated languages will no longer receive updates or support in these languages.
- All user interfaces, help documentation, and customer support will be available only in English or in the three supported languages mentioned above.

Removal of vSphere Lifecycle Manager baselines:

Support for managing clusters with vSphere Lifecycle Manager baselines and baseline groups (legacy vSphere Update Manager workflows) will drop in a future vSphere release. Instead of baselines and baseline groups, you can use vSphere Lifecycle Manager images to perform tasks on a standalone host or on a cluster level such as install a desired ESXi version on all hosts in a cluster, install and update third-party software, update, and upgrade ESXi or firmware, generate recommendations, and use a recommended image for your cluster.

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 3

Product Patch for vCenter Server containing VMware software fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00000-24022515-patch-FP.iso
Build	24022515
Download Size	8507.3 MB
sha256checksum	f611bba1fca57bfc81a021b0de2433a1df284b5283e0750f49eb2272fdd908ed

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions, see [Download Broadcom products and software](#).

1. Attach the VMware-vCenter-Server-Appliance-8.0.3.00000-24022515-patch-FP.iso file to the vCenter Server CD or DVD drive.
2. Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:
 - To stage the ISO:


```
software-packages stage --iso
```
 - To see the staged content:


```
software-packages list --staged
```
 - To install the staged rpms:


```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

vSphere Client and vCenter Issues

PR 3308624: ESXi hosts fail to boot with an error such as Could not boot image: Invalid argument (http://ipxe.org/err/1c0de8)

After a change in the root certificate authority in a vCenter instance, when you use vSphere Auto Deploy to boot ESXi hosts, the TLS handshake with the vCenter might not succeed. As a result, booting of the hosts fails with a message such as `Could not boot image: Invalid argument (http://ipxe.org/err/1c0de8)`.

This issue is resolved in this release. For more information, see KB [96827](#).

PR 3210868: vCenter support bundles fill up the /storage/log partition

Due to a change in the vCenter permissions model, the vpxd service might not be able to delete the vCenter support bundles it temporarily stores in the `/storage/log` partition. As a result, you might see vCenter support bundles fill up the `/storage/log` partition.

This issue is resolved in this release. If you already face the issue, you can safely delete such files manually.

PR 2993755: You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.x and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.x vCenter Server.

This issue is resolved in this release.

Miscellaneous Issues

The PCI slot of a virtual device might change after upgrade from hardware version 19

When you upgrade a virtual machine with hardware version 19 and configured with EFI firmware to a newer virtual hardware version, the PCI device addresses of the virtual devices in the VM might change. As a result, the guest operating system of the VM might not apply configurations on a PCI device at a particular address.

This issue is resolved in this release.

PR 3280955: You see incorrect Maximum Size value for thin-provisioned virtual disks

In the vSphere Client, when you look at the settings of a VM with a thin-provisioned virtual disk, you might see the **Maximum Size** value for the disk larger than the capacity of the datastore where the disk is located. For example, if a datastore capacity is 100GB of which 90GB are available, and a thin-provisioned virtual disk has 50GB capacity, of which only 10GB are utilized, you might see a **Maximum Size** value of 140 GB, adding the available datastore capacity to the overall disk capacity, not its actual utilization.

This issue is resolved in this release.

PR 3020259: You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

This issue is resolved in this release.

vSphere Lifecycle Manager Issues

PR 3080596: During the generation of a new image recommendation by the vSphere Lifecycle Manager, the vpxd service fails due to memory exhaustion

If the desired state image on a vSphere Lifecycle Manager cluster has several additional Components, and especially if some of them are not up to date, during the generation of a new image recommendation, the memory required to calculate all possible combinations of Components that have new versions might grow to 85% of the overall memory of the vpxd service. As a result, the service might fail with `coreDump.#####` files. In the `/var/log/vmware/vmware-updatemgr/vum-server/vmware-vum-server.log`, you see errors such as:

```
2022-12-08T11:00:12.208 info vmware-vum-server[50005] [Originator@6876
sub=RecommendationEngine::ReUtil] [reUtil 108] GenerateImageUnitsCombination: Generating
ImageUnit combinations
```

```
2022-12-08T11:05:34.821 error vmware-vum-server[50005] [Originator@6876 sub=Default]
Unable to allocate memory
```

This issue is resolved in this release.

High Availability and Fault Tolerance Issues

PR 3354058: vSphere HA advanced options `das.respectVmVmAntiAffinityRules` and `das.respectVmVmAffinityRules` do not work as expected

In some cases, the vSphere HA advanced options `das.respectVmVmAntiAffinityRules` and `das.respectVmVmAffinityRules` might not work as expected. For example, when you set `das.respectVmVmAntiAffinityRules` to `FALSE`, expecting that the vm-vm anti-affinity rules to be ignored during a vSphere HA failover, some VMs fail to failover with a `RuleViolation` fault. In other cases, you might expect a VM failover to respect the vm-vm anti-affinity rule as by default `das.respectVmVmAntiAffinityRules` is `TRUE`, but some VMs fail over and the rule is violated.

This issue is resolved in this release.

Known Issues

vCenter Server and vSphere Client Issues

When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client

If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

You see the same license product name displayed multiple times in the vSphere Client

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

Workaround: None

An external gateway firewall might block vSphere Pod traffic to clusterIPs

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to ClusterIPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see [vSphere Pod Traffic to ClusterIP Time-outs](#).

vSphere Lifecycle Manager Issues

In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image

When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components** and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

Installation, Upgrade, and Migration Issues

If you do not open port 9087 in your firewall between ESXi hosts and vCenter, compliance checks and vSphere HA might fail after vCenter update to 8.0 Update 3

Starting from 8.0 Update 3, the vSphere Lifecycle Manager downloads updates for ESXi hosts by a HTTPS connection to the vCenter instance on port 9087. If you do not open port 9087 in your firewall between ESXi hosts and vCenter, you might see compliance check errors. For example, in the `lifecycle.log` you see messages such as:

```
<Timestamp> In(14) lifecycle[2112988]: Downloader:373 Opening https://<VC-
FQDN>:9087/vum/repository/hostupdate/__micro-depot__vendor-DEL__DEL-ESXi-8.0-Addon-
cumulative_metadata__index__.xml for download

<Timestamp> Wa(12) lifecycle[2112988]: Downloader:210 Download failed: <urlopen error
timed out>, 9 retry left...
```

In addition, vSphere HA might fail to start.

Workaround: Open port 9087 in your firewall between ESXi hosts and vCenter.

vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system

Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as The host returns `esxupdate` error codes: -1.

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:
Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails

In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

You see a "vc.health.error.dbjob3" warning after a vCenter upgrade

In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

Workaround: See [Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x](#) how to clear stats data, if not relevant anymore.

Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"

Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists`.

Workaround: Delete all files under `/storage/log/vmware/analytics/stage`, `/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

vSphere Cluster Service Issues

In the vSphere Client, you see the vSphere HA status of Embedded vSphere Cluster Service VMs as Unprotected

vSphere 8.0 Update 3 introduces a redesign of vCLS to Embedded vCLS, which utilizes vSphere Pod technology. Deployment and lifecycle of such VMs are now managed within ESXi and are no longer managed by the vSphere ESX Agent Manager (EAM). On the **Summary** tab in the vSphere Client, you see the vSphere HA status of Embedded vSphere Cluster Service VMs as Unprotected, but this is expected, because vSphere HA does not manage Embedded vSphere Cluster Service VMs.

Workaround: Ignore the information in the vSphere HA card on the **Summary** tab in the vSphere Client. For more information, see [vSphere Cluster Services](#).

Security Features Issues

Adding a Standard Key Provider to your vCenter system from the vSphere Client fails with `Vim.fault.DatabaseError`

Starting with vCenter 8.0 Update 1c, you must use a RSA certificate when adding a Standard Key Provider to vCenter otherwise in the vSphere Client you see a `Vim.fault.DatabaseError` error. In the `vmafdd.log` file, you see the following line:

```
<Timestamp> [vmafdd][ERROR] Certificate uses an unsupported signature algorithm
(NID=ecdsa-with-SHA256). Only SHA-2 RSA algorithms are supported on the vCenter Server.
```

Workaround: Use a RSA certificate with a SHA-2 digital signature algorithm when trusting a KMS.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to https://kb.vmware.com/s/article/93526 to trust the URLs: <LIST_OF_URLS>
```

Or

```
Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against
trusted certificates by the System! Verify that the ESX Agent Manager extension is running
properly on the source vCenter Server instance and https://VC\_IP/eam/mob presents correct
data. If log in to the MOB is not successful, try resolving the issue with https://kb.vmware.com/s/article/94934.
```

Workaround: For more information, see VMware knowledge base articles [93526](https://kb.vmware.com/s/article/93526) and [94934](https://kb.vmware.com/s/article/94934).

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to

detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: Update 8.0.2.00000 for component vlcm is not found.

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Verify following URLs and their respective statuses and follow KB 93526.
```

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot /storage/core unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as No such file or directory: '/storage/core/software-update/install_operation'. In the software-packages.logs, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
```

```
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
```

```

2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf

2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.

```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```

root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback

```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed

machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service..` This is expected behavior.

Workaround: None

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VM DirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the `ntg3` driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

vCenter Server and vSphere Client Issues

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 6:

Default VM Settings

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag: `KeepAfterDeleteVM`.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 2e Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)

Introduction

VMware vCenter Server 8.0 Update 2e | 21 OCT 2024 | ISO Build 24321653
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2024-38812 and CVE-2024-38813. For more information on these vulnerabilities and their impact on VMware by Broadcom products, see [VMSA-2024-0019.2](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 2d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2c Release Notes](#)
- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 2e

Product Patch for vCenter Server containing software fixes.

Download Filename	VMware-vCenter-Server-Appliance-8.0.2.00500-24321653-patch-FP.iso
Build	24321653
Download Size	7874.2 MB
sha256checksum	aeca2582a4ec72c22cc72409546b2776b01f176882df4ec330b6b34ce2040c1a
PRs fixed	N/A
CVEs	CVE-2024-38812, CVE-2024-38813

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Issues

This release resolves CVE-2024-38812 and CVE-2024-38813. For more information on these vulnerabilities and their impact on VMware by Broadcom products, see [VMSA-2024-0019.2](#).

VMware vCenter Server 8.0 Update 2d Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware vCenter Server 8.0 Update 2d | 17 JUN 2024 | ISO Build 23929136
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2024-37079, CVE-2024-37080, and CVE-2024-37081. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0012](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 2c Release Notes](#)
- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 2d

Product Patch for vCenter Server containing VMware software fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.2.00400-23929136-patch-FP.iso
Build	23929136
Download Size	7874.2 MB
sha256checksum	566d3a9b866ce7af27c0655ae2a7115fc76ca226f1917aa26d4890175243b2e0
PRs fixed	NA
CVEs	CVE-2024-37079, CVE-2024-37080, CVE-2024-37081

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

1. Attach the VMware-vCenter-Server-Appliance-8.0.2.00400-23929136-patch-FP.iso file to the vCenter Server CD or DVD drive.
2. Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:
 - To stage the ISO:

```
software-packages stage --iso
```
 - To see the staged content:

```
software-packages list --staged
```
 - To install the staged rpms:

```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Issues

This release resolves CVE-2024-37079, CVE-2024-37080, and CVE-2024-37081. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0012](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLS>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](#) and [94934](#).

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Verify following URLs and their respective statuses and follow KB 93526.

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```
root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback
```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

You see incorrect Maximum Size value for thin-provisioned virtual disks

In the vSphere Client, when you look at the settings of a VM with a thin-provisioned virtual disk, you might see the **Maximum Size** value for the disk larger than the capacity of the datastore where the disk is located. For example, if a datastore capacity is 100GB of which 90GB are available, and a thin-provisioned virtual disk has 50GB capacity, of which only 10GB are utilized, you might see a **Maximum Size** value of 140 GB, adding the available datastore capacity to the overall disk capacity, not its actual utilization.

Workaround: None

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service. . This is expected behavior.`

Workaround: None

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues**Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail**

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 7:

Default VM Settings

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API : `setVStorageObjectControlFlags`. You can invoke the FCD API at https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags and pass the control flag : `KeepAfterDeleteVM`.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 2c Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware vCenter Server 8.0 Update 2c | 26 MAR 2024 | ISO Build 23504390
Check for additions and updates to these release notes.

What's New

- VMware vCenter Server 8.0 Update 2c serves as a vehicle for [VMware vSphere with Tanzu](#) and does not deliver vCenter Server fixes.

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 2b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 2c

Product Patch for vCenter Server containing VMware software fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.2.00300-23504390-patch-FP.iso
Build	23504390
Download Size	7873.2 MB
sha256checksum	050ed8d0fdee1e174f4e70f51c31382429ae7634f5c9211aaa9efb498cfaa790

Download and Installation

Download this patch from [Broadcom Support Portal](#).

1. Attach the VMware-vCenter-Server-Appliance-8.0.2.00300-23504390-patch-FP.iso file to the vCenter Server CD or DVD drive.
2. Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:
 - To stage the ISO:

```
software-packages stage --iso
```
 - To see the staged content:

```
software-packages list --staged
```
 - To install the staged rpms:

```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to https://kb.vmware.com/s/article/93526 to trust the URLs: <LIST_OF_URLS>
```

Or

```
Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against
trusted certificates by the System! Verify that the ESX Agent Manager extension is running
properly on the source vCenter Server instance and https://VC\_IP/eam/mob presents correct
data. If log in to the MOB is not successful, try resolving the issue with https://kb.vmware.com/s/article/94934.
```

Workaround: For more information, see VMware knowledge base articles [93526](#) and [94934](#).

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: Update 8.0.2.00000 for component vlcm is not found.

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Verify following URLs and their respective statuses and follow KB 93526.

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```
root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
```

```
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback
```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

You see incorrect Maximum Size value for thin-provisioned virtual disks

In the vSphere Client, when you look at the settings of a VM with a thin-provisioned virtual disk, you might see the **Maximum Size** value for the disk larger than the capacity of the datastore where the disk is located. For example, if a datastore capacity is 100GB of which 90GB are available, and a thin-provisioned virtual disk has 50GB capacity, of which only 10GB are utilized, you might see a **Maximum Size** value of 140 GB, adding the available datastore capacity to the overall disk capacity, not its actual utilization.

Workaround: None

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service. . This is expected behavior.`

Workaround: None

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 8:

Default VM Settings

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?`

`moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag : `KeepAfterDeleteVM`.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a

vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 2b Release Notes

This document contains the following sections

- [Introduction](#)

- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 2b | 29 FEB 2024 | GA ISO Build 23319993

Check for additions and updates to these release notes.

What's New

Solution License for VMware Cloud Foundation (VCF) and VMware vSphere Foundation (VVF):

Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VVF and VCF.

For more information on applying the Solution License to VVF components, see [VMware vSphere Foundation \(VVF\) Licensing](#).

For more information on VCF Solution License, see [VMware Cloud Foundation 5.1.1 Release Notes](#).

100 GiB of included storage capacity per licensed core:

Starting with vSphere 8.0 Update 2b, as part of the VMware vSphere Foundation Solution License, you can use up to 100 gibibytes (GiB) of included vSAN storage per host licensed core. For capacity larger than 100 GiB per core, you must purchase vSAN capacity per terabyte (TiB) and apply a vSAN license key that reflects the total raw storage capacity of the vSAN cluster. For more information on capacity reporting and licensing in vSAN, see [Demystifying Capacity Reporting in vSAN](#) and [Counting Cores for VMware Cloud Foundation and vSphere Foundation and TiBs for vSAN](#).

Aggregation of GPU statistics on host and cluster level:

With vCenter Server 8.0 Update 2b, you can track host and cluster-level aggregation of GPU metrics, which benefit generative AI workloads, traditional AI workloads, and non-AI workloads accelerated by GPUs. In the vSphere Client, under **Monitor > Performance**, you see:

- on a host level: performance overview charts and advanced performance charts for the GPU compute usage, memory allocation, and temperature on hosts.
- on a cluster level: advanced performance charts with aggregated statistics from the hosts for the GPU memory and the GPU utilization.

For more information, see [Working with Advanced and Custom Charts](#) and [Overview Performance Charts](#).

Other updates:

- VMware vCenter Server 8.0 Update 2b delivers bug and security fixes documented in the [Resolved Issues](#) section.
- For VMware vSphere with Tanzu updates, see [VMware vSphere with Tanzu Release Notes](#).
- For Photon OS updates, see [VMware vCenter Server Appliance Photon OS Security Patches](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 2a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)

- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 2b

Product Patch for vCenter Server containing VMware software fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.2.00200-23319993-patch-FP.iso
Build	23319993
Download Size	7862.5 MB
sha256checksum	48532b79c8329add780637671d90d631f34f8f9090c8f54a9b48b7892e5011ef
PRs fixed	3288037, 3298859, 3299658, 3273124, 3301131, 3347798, 3289235, 3294936, 3290053, 3290248, 3267090, 3293934, 3321231, 3277663, 3296876, 3281906, 3291336, 3298209, 3319926, 3291909, 3292182, 3297001, 3273930, 3239235, 3297450

Download and Installation

Download this patch from [Broadcom Support Portal](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

1. Attach the VMware-vCenter-Server-Appliance-8.0.2.00200-23319993-patch-FP.iso file to the vCenter Server CD or DVD drive.
2. Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:
 - To stage the ISO:


```
software-packages stage --iso
```
 - To see the staged content:


```
software-packages list --staged
```
 - To install the staged rpms:


```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Installation, Upgrade, and Migration Issues

PR 3296582: If your vCenter system has limited or no Internet access, vCenter upgrade to 8.0 Update 2 might stay at 93% for a few hours

Upgrades to vCenter 8.0 Update 2 include an upgrade from Photon 3 to Photon 4, which requires access to the Photon online repository. If your vCenter system cannot to access the repository due to limited or no Internet access, the system makes multiple retries and the upgrade progress stays at 93% for a few hours. The upgrade eventually completes. In the `/var/log/vmware/applmgmt/update_microservice.log`, you see errors such as:

```
2023-10-18 15:14:23,120 - 34391 - update_script:: collectRequirements: 495 - DEBUG -
Source VCSA version = 8.0.1.00200

2023-10-18 15:14:23,120 - 34391 - update_script:: collectRequirements: 500 - INFO - Target
VCSA version = 8.0.2.00000.

2023-10-18 15:20:37,336 - 45585 - functions_target:: createPhotonRepo: 270 - INFO -
appliance-photon repo created successfully

2023-10-18 15:20:37,336 - 45585 - functions_target:: runCommand: 79 - DEBUG - Invoking
command [/bin/bash --login -c /usr/bin/photon-upgrade.sh --skip-update --assume-yes --
upgrade-os --repos=appliance-photon --install-all..

2023-10-18 16:51:04,634 - 54435 - update_script::_executePatchRunnerAndSupplyUserData: 132
- DEBUG - Patch command patch completed successfully

2023-10-18 16:51:11,472 - 45585 - update_b2b_target:: install:2820 - DEBUG - Installation
completed successfully!
```

This issue is resolved in this release. For more information, see KB [326203](#).

PR 3288037: The deployment of a new vCenter instance of version 8.0 Update 2 fails at around 80% during stage 1 with an error for RPM installation

The deployment of a new vCenter instance of version 8.0 Update 2 might fail at around 80% during stage 1 due to a rare kernel panic. In the GUI installer, you see an error such as `Waiting for RPM installation to start`. In the console logs, you see errors such as:

```
ecdh: Party A: compute shared secret test failed. err -14

Kernel panic - not syncing: alg: self-tests for ecdh-generic (ecdh) failed in fips mode!
```

This issue is resolved in this release.

PR 3298859: Patching vCenter to 8.0 Update 2 or later fails with an error such as Installation of Photon3 RPM failed

Updates to vCenter 8.0 Update 2 or later might fail due to a rare issue with the `vmware-vmon.service.bak` file in `/etc/systemd/system/`. In the `update_microservice.log`, you see errors such as:

```
stderr [ package createrepo_c-0.11.1-5.ph3.x86_64 is already installed]

update_b2b_target:: executeInstallRpms:1920 - ERROR - RPM installation failed.

update_b2b_target:: installRpms:2358 - ERROR - Installation of one or more RPMs failed.

update_b2b_target:: installRpms:2359 - ERROR - Installation of Photon3 RPM failed.
```

This issue is resolved in this release.

PR 3299658: Migration of virtual machines with Virtual Trusted Platform Module (vTPM) across vSphere DRS-enabled clusters fails

In the vSphere Client, when you try to migrate VMs with vTPM across DRS-enabled clusters without selecting any hosts, the final step of the wizard fails with an error such as `A DRS fault has occurred while retrieving VM placement recommendations: xvmotionPlacement`. If you use the PlaceVM API to live relocate a VM with a vTPM from one DRS cluster to another DRS cluster, the operation also fails.

This issue is resolved in this release. With the fix, you can safely migrate VMs with vTPM across DRS-enabled clusters.

PR 3273124: You see significant delay in downloading ESXi patches from the vSphere Update Manager Download Service when using a proxy server

Due to URL accessibility checks, you might see a delay of up to 24 hours in downloading ESXi patches from the Update Manager Download Service when using a proxy server.

This issue is partially resolved in this release. The time for downloading ESXi patches from the Update Manager Download Service when using a proxy server is reduced to few hours.

PR 3301131: The vSphere DRS migration threshold slider does not display correctly in the vSphere Client

On vCenter Server 8.0 Update 2, when you navigate to **vSphere DRS > Edit > Automation > Migration Threshold**, the slider appears much shorter than in previous releases. This is a cosmetic issue that does not affect the functionality of the slider.

This issue is resolved in this release.

Virtual Machine Management Issues

PR 3347798: The VPXD service might fail during a consolidation task due to file locked exception

During a consolidation task in a vCenter environment, such as **Consolidate virtual machine disk files**, multiple processes lock more than one file and in some environments it is possible to see multiple `File Locked` exceptions. As a result, the VPXD service might fail while fetching the lock details for the affected files.

This issue is resolved in this release.

PR 3289235: You cannot manually set a static MAC address on a virtual machine by using the vSphere Client

In the vSphere Client, when you right-click on a virtual machine, select **Edit Settings > Virtual Hardware > Network adapter > MAC Address** and select **Manual** from the drop-down menu, then enter a MAC Address and click **OK**, the **Reconfigure virtual machine** task appears to complete successfully. But when you check the configured MAC address, it is still set to **Automatic** and the previous MAC Address is unchanged.

This issue is resolved in this release.

PR 3294936: Encrypted virtual machines in a vSphere DRS cluster might power on slowly

Power-on of encrypted VMs in a DRS-enabled cluster might take more than 90 seconds because of the time it takes to synchronize the VM encryption keys to each ESXi host within the cluster.

This issue is resolved in this release. The fix removes redundant encryption key synchronizations.

PR 3290053: You cannot edit the settings of powered on VMs of hardware version 9 or earlier

In the vSphere Client, when you open **Edit Settings** on a powered on VM of hardware version 9 or earlier, you see a red exclamation mark beside the **Virtual Hardware** and **VM Options** tabs and no content in the dialog box. The **Edit Settings** dialog works for powered off VMs.

This issue is resolved in this release. The fix makes sure that the edit settings of VMs of hardware version 9 or earlier display properly in the vSphere Client.

PR 3290248: In the vSphere Client, you see only 15 objects per page when migrating a VM to a new compute resource

In the vSphere Client, when you migrate a virtual machine and opt to change only the compute resource, on the **Select compute resource** page you see up to 15 objects as compared to the default of 35 for inventory lists.

This issue is resolved in this release. The fix makes sure you see a list of up to 35 objects on the **Select compute resource** page.

Storage Issues

PR 3267090: You see unusually high CPU usage when you open the Storage DRS recommendations page under the vSphere Client datastore view

In the **Monitor** tab in the vSphere Client datastore view, when you look at the **Storage DRS** recommendations, you might see unusually high CPU usage, up to 100%. The issue occurs due to a fault search algorithm in the DRS recommendation service.

This issue is resolved in this release.

Server Configuration Issues

PR 3293934: You see an error for invalid certificate bytes when configuring new identity source for Active Directory over LDAP

If a certificate .pem file contains the Windows-specific line endings `\r\n`, or CRLF, you might not be able to configure a new identity source for Active Directory over LDAP. The configuration workflow fails with an error **Cannot configure identity source due to invalid certificate bytes**.

This issue is resolved in this release.

Miscellaneous Issues

PR 3321231: Virtual machine tasks such as power off and clone take unusually long

In rare cases, when a service account password for the vService Manager expires, running operations on virtual machines such as power on, clone, and delete might take unusually long to complete. You can see the operations in the task list, but you see no progress for more than 10 minutes.

This issue is resolved in this release. If you are already facing the issue, log into the vCenter VM and run the command `service-control --restart vsm`.

PR 3277663: You see many errors for HTTP failure due to the waitForUpdates method in the vpx service logs

Canceling the `waitForUpdates()` method is part of the normal workflow when the vpx service waits for a task running on the vpxa or hostd service to complete. However, the canceled call results in a failed HTTP request and subsequently in a log message at the `info` level. In some environments, such messages, which are not requesting any specific action, might significantly load the vpx service logs.

This issue is resolved in this release. The fix lowers the log level to `verbose` in case the fault for the `waitForUpdates()` method is `RequestCanceled` to avoid logging of unnecessary messages in the vpx service logs.

PR 3296876: You cannot edit the vCenter alarm Datastore usage on disk based on available capacity

When you select the **Datastore usage on disk based on available capacity** option in the alarm definition wizard from the **Configure** tab in the vSphere Client, the Edit task displays indefinitely in progress and you cannot make any changes.

This issue is resolved in this release.

Networking Issues

PR 3281906: You see a black screen after activating the VMware Remote Console (VMRC) proxy service

After you activate the VMware Remote Console, you see a black screen with an error such as `The webmks ticket has expired.` instead of configuration parameters.

This issue is resolved in this release.

PR 3291336: Port ID selections in the Add Port Mirroring Session wizard might not work

When you create a port mirroring session by using the vSphere Client to mirror vSphere Distributed Switch traffic to ports, uplinks, and remote IP addresses, Port ID selections in the **Select sources** and **Select destinations** steps might not work. For example, in the **Select sources** page, when you select a Port ID check box and filter to select a host, after selecting a host, the selection on the Port ID disappears. As a result, you cannot filter an ESXi host.

This issue is resolved in this release.

PR 3298209: In the vSphere Client, you do not see all ESXi hosts available to add to a vSphere Distributed Switch

In the vSphere Client, when you associate ESXi hosts in a cluster with a vSphere Distributed Switch, not all hosts in the cluster are visible and available for selection under the **Select hosts** page of the **Add and Manage Hosts** wizard. For example, you might see 8 hosts when the total number of available hosts in the cluster is 20 and you cannot scroll down to see the full list. When you click **Select All**, you see 20 selected hosts but they remain not visible.

This issue is resolved in this release.

vCenter Server and vSphere Client Issues

PR 3319926: Login and vSphere Client sessions might be affected from connectivity issues between VMware Cloud Gateway and VMware Cloud

With vSphere 8.0 Update 2b, if you use vSphere 8.0 Update 2 or later, you no longer need to honor the 45-minute window of vSphere Client sessions to maintain connectivity between VMware Cloud Gateway and VMware Cloud. If your vSphere system is of version earlier than vSphere 8.0 Update 2, after upgrading to 8.0 Update 2b, you can log in to the vSphere Client at any time and see no error for loss of connectivity between VMware Cloud Gateway and VMware Cloud.

This issue is resolved in this release.

PR 3291909: In air-gapped environments, you see untranslated or partial messages in Japanese locale in the vSphere Client for vSAN troubleshooting options

Only in air-gapped environments, when you navigate in the vSphere Client to **vSAN > Skyline Health > Overview > vSAN optimal datastore default policy configuration > Troubleshoot**, you might see untranslated or partial messages in the Japanese locale.

This issue is resolved in this release.

PR 3292182: Log in to the vSphere Client intermittently fails

In large environments where the vsphere-ui service handles multiple concurrent tasks and a single task request might exceed 20 MB, the service might run out of memory and fail. As a result, log in to the vSphere Client intermittently fails.

This issue is resolved in this release.

PR 3297001: After upgrading to vCenter 8.0 Update 2, you see many warnings in the vSphere Client that a privilege check fails

In some environments, during the prestart stage of the VMware Service Manager (VSM), a missing privilege for the VsmSvcRole might cause vCenter to issue multiple warnings after an upgrade to version 8.0 Update 2. In the vSphere Client, under **Monitor > Events**, you might see hundreds of warnings per day.

This issue is resolved in this release. The fix makes sure the VsmSvcRole gets all required privileges.

PR 3273930: You cannot open the vSphere Client Summary tab from a VMware Aria Operations object

When you navigate to a VMware Aria Operations object, for example a VM under **Environments > vSphere > vSphere Host and Cluster > vSphere World > vCenter Server > Datacenter**, and select **Open Virtual Machine in vSphere Client** from the **Actions** drop down menu, you see an error. The issue occurs due to accidental encoding of some characters in the navigation segment. As a result, the vSphere Client cannot open the **Summary** tab of the virtual machine as the current object and you see the error `Missing Route This route is not supported:<link to object>`.

This issue is resolved in this release.

Security Issues

PR 3297450: vCenter Server 8.0 Update 2b provides the following security updates:

- The Guava suite is updated to version 32.1.0-jre.
- Apache Tomcat is updated to versions 8.5.95/9.0.82.
- Eclipse Jetty is updated to version 9.4.53.v20231009.
- Apache Log4j is updated to version 2.21.1.

Known Issues

Installation, Upgrade, and Migration Issues

In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode

In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**

Workaround: Power-off the incompatible virtual machine.

Storage Issues

Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state

In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article [2147750](#).

vCenter Server and vSphere Client Issues

You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box

In the vSphere Client, when you select an ESXi host and click **Configure > Virtual Machines > VM Startup/Shutdown > Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

- System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
- Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
- Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
- Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

Workaround: None. See the screenshot to figure out the sequence of labels:

Figure 9:

Default VM Settings

System influence	☐ Automatically start and stop the virtual machines with the system
Startup delay	120
Shutdown delay	120
	☐ Continue if VMware Tools is started
Shutdown action	Power off ▼

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLs>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](https://kb.vmware.com/s/article/93526) and [94934](https://kb.vmware.com/s/article/94934).

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](#).

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Verify following URLs and their respective statuses and follow KB 93526.
```

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result,

the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf

2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation

2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf

2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```
root@localhost []# cat /etc/hosts

# Begin /etc/hosts (network card version)

127.0.0.1 localhost.localdomain

127.0.0.1 localhost

::1 localhost ipv6-localhost ipv6-loopback
```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss` : Use the DRSS parameter.
- `drss` : Use the DRSS parameter.
- `rss` : Use the RSS parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

You see incorrect Maximum Size value for thin-provisioned virtual disks

In the vSphere Client, when you look at the settings of a VM with a thin-provisioned virtual disk, you might see the **Maximum Size** value for the disk larger than the capacity of the datastore where the disk is located. For example, if a datastore capacity is 100GB of which 90GB are available, and a thin-provisioned virtual disk has 50GB capacity, of which only 10GB are utilized, you might see a **Maximum Size** value of 140 GB, adding the available datastore capacity to the overall disk capacity, not its actual utilization.

Workaround: None

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service..` This is expected behavior.

Workaround: None

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API : `setVStorageObjectControlFlags` . You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag : `KeepAfterDeleteVM` .

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations` , which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted` . The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 2a Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 2a | 26 OCT 2023 | GA ISO Build 22617221
Check for additions and updates to these release notes.

What's New

- For fixes and known issues in vCenter Server 8.0 Update 2a, see the sections [Resolved Issues](#) and [Known Issues](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 2 Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 2a

Product Patch for vCenter Server containing VMware software fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.2.00100-22617221-patch-FP.iso
--------------------------	---

Build	22617221
Download Size	7856.2 MB
sha256checksum	e41a8407d8cec0f523f779608fb64f60d75b132d4749176a360a098ce9c648a1

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

vSAN Issues

Cannot apply vSAN standard license to a vSAN two-node cluster after upgrade

After you upgrade a vCenter system to version 8.0 Update 2, the vSAN evaluation license is applied to vSAN clusters. If you attempt to apply a vSAN standard license to a vSAN two-node cluster, you might see the following error: **The license does not support some of the features that assets currently use**. This issue occurs because the stretched cluster feature is reported as in-use state for vSAN two-node clusters.

This issue is resolved in this release.

Install, Upgrade and Migration Issues

A set of VMDKs might get associated to incorrect mount points in a vCenter instance

In very rare cases, the Linux kernel might associate an incorrect size of the mount point for some VMDKs in a vCenter instance. You see the issue when you run the `lsblk` command and compare the result with the default disk layout for vCenter. For example, the default size for a `/storage/seat` mount point for a vCenter of medium deployment size is expected to be 50GB and for `/storage/netdump` is 10GB, but after running the `lsblk` command, you might see these sizes reversed, or 10GB for the `/storage/seat` mount point and 50GB for `/storage/netdump`. The issue might affect any number or sequence of VMDKs. The issue occurs during a vCenter deployment and might result in failure of backup and restore operations.

This issue is resolved in this release.

You see an error "Failed to perform Cleanup" during patching to vCenter Server 8.0 Update 2

The Logical Volume Manager (LVM) might take longer than the 4 min time out for cleanup operations during an update to vCenter 8.0 Update 2. As a result, in the vSphere Client you see the error **Failed to perform Cleanup**. The update operation completes successfully, but some post-update tasks fail.

This issue is resolved in this release. The fix increases the timeout to 30 min.

Upgrade or fresh installation of vCenter 8.0 Update 2 might fail during the start of the vc-ws1a-broker service

During an upgrade or installation of vCenter 8.0 Update 2, if a DNS server in your network is offline or not configured correctly, during the start of the vc-ws1a-broker service, which enables new identity providers such as Microsoft EntraID and Okta, the PostgreSQL database connection process might time out. As a result, the service fails to start, and the upgrade or installation job is not successful.

This issue is resolved in this release.

Server Configuration Issues

After update or upgrade to vCenter Server 8.0 Update 2, ESXi hosts disconnect from vCenter

In rare cases, when a vCenter holds a legacy non-CA certificate in its TRUSTED_ROOTS store, after an upgrade or update to 8.0 Update 2, the hostd service filters out such certificates and certain ESXi hosts cannot reconnect to this vCenter. Attempts to manually add such hosts to the vCenter system also fail.

This issue is resolved in this release.

Known Issues

Installation, Upgrade, and Migration Issues

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Please refer to <https://kb.vmware.com/s/article/93526> to trust the URLs: <LIST_OF_URLs>

Or

Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check against trusted certificates by the System! Verify that the ESX Agent Manager extension is running properly on the source vCenter Server instance and https://VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try resolving the issue with <https://kb.vmware.com/s/article/94934>.

Workaround: For more information, see VMware knowledge base articles [93526](https://kb.vmware.com/s/article/93526) and [94934](https://kb.vmware.com/s/article/94934).

vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"

If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article [94840](https://kb.vmware.com/s/article/94840).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: Update 8.0.2.00000 for component vlcm is not found.

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

```
Source ESX Agent Manager Configuration contains URLs that are not trusted by the System!
Verify following URLs and their respective statuses and follow KB 93526.
```

Workaround: For more information, see VMware knowledge base article [93526](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot /storage/core unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as No such file or directory: '/storage/core/software-update/install_operation'. In the software-packages.logs, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf
```

```
2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation
```

```
2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'
```



```

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf
2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.

```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```

root@localhost []# cat /etc/hosts
# Begin /etc/hosts (network card version)
127.0.0.1 localhost.localdomain
127.0.0.1 localhost
::1 localhost ipv6-localhost ipv6-loopback

```

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is not active, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliance$` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

You see incorrect Maximum Size value for thin-provisioned virtual disks

In the vSphere Client, when you look at the settings of a VM with a thin-provisioned virtual disk, you might see the **Maximum Size** value for the disk larger than the capacity of the datastore where the disk is located. For example, if a datastore capacity is 100GB of which 90GB are available, and a thin-provisioned virtual disk has 50GB capacity, of which only 10GB are utilized, you might see a **Maximum Size** value of 140 GB, adding the available datastore capacity to the overall disk capacity, not its actual utilization.

Workaround: None

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service..` This is expected behavior.

Workaround: None

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VM DirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag: `KeepAfterDeleteVM`.

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which

expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers.

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

VMware vCenter Server 8.0 Update 2 Release Notes

This document contains the following sections

- [Introduction](#)
- [General Availability](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Product Support Notices](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 2 | 21 SEP 2023 | GA ISO Build 22385739
Check for additions and updates to these release notes.

General Availability

This vCenter Server 8.0 Update 2 release is a General Availability (GA) designation. For more information on the vSphere 8.0 IA/GA Release Model of vSphere Update releases, see [The vSphere 8 Release Model Evolves](#).

What's New

NEW: This release resolves CVE-2023-34048 and CVE-2023-34056. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2023-0023](#).

For overview of the new features in vSphere 8.0 Update 2, see [What's New in vSphere 8 Update 2](#) and the following release notes:

- [VMware vSphere With Tanzu](#)
- [Tanzu Kubernetes releases](#)
- [VMware vSAN](#)
- [VMware Host Client](#)
- [VMware vSphere+ Release Notes](#)

The list of new features and enhancements that follows adds some of the key highlights for vCenter systems in vSphere 8.0 Update 2:

vCenter and vSphere Client

- **Reduce upgrade downtime for a self-managed vCenter:** vSphere 8.0 Update 2 introduces a migration-based approach for upgrading vCenter between major versions that reduces downtime to several minutes. Reduced downtime upgrade is only supported for single self-managed vCenter instances that have no vCenter HA and do not participate in Enhanced Linked Mode (ELM). For more information, see [Reduced Downtime Upgrade](#) and VMware knowledge base article [92659](#).
- **Automatic pre-update Logical Volume Manager (LVM) snapshots for quick and reliable rollback:** Starting with vSphere 8.0 Update 2, the vCenter patching orchestrator automatically creates a LVM snapshot of the vCenter before performing a patch installation to allow safe roll back in case the update fails. You can use the rollback capability only for updates from vCenter 8.0 Update 2 to a later version. For more information, see [Patching vCenter Server Using the vCenter Server Management Interface](#) and [Resilient vCenter Patching](#).
- **Detailed information about File lock owners and processes in vCenter:** vSphere 8.0 Update 2 adds detailed information to file lock errors such as `failed to lock the file` to enable root cause analysis and troubleshooting. The new messages include details on which file is locked, which host owns the file lock and which process on that host is responsible for the file lock.
- **API for vSphere plug-ins to classify vCenter events:** Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecated support for local plug-ins and some vCenter alarms were only available in the local plug-ins. With vSphere 8.0 Update 2, you can view and manage alarms related to vCenter in the vSphere plug-in and no longer switch to the legacy local plug-ins.
- **Keep network configuration changes after a vCenter backup:** With vSphere 8.0 Update 2, when you make changes to the network configuration or the vSphere Distributed Switch after a vCenter backup, such as adding or removing an ESXi host to VDS or an NSX port group, all changes persist when a vCenter is restored from backup.
- **Non-disruptive SSL certificate renewal and replacement for vCenter:** With vSphere 8.0 Update 2, you can renew or replace the vCenter machine and trusted-root certificates with no restart of the vCenter services. For more information, see [Non-disruptive Certificate Management](#).

- **vCenter support for identity federation with Microsoft Entra Active Directory (former Azure AD):** vSphere 8 Update 2 adds Microsoft Entra Active Directory as an external identity provider to reduce the attack surface for vCenter. For more information, see [Configure vCenter Server Identity Provider Federation for Azure AD](#).
- **Activate retreat mode for vSphere Cluster Services (vCLS) from the vSphere Client:** With vSphere 8.0 Update 2, you can activate retreat mode on a vCLS cluster by using the **EDIT VCLS MODE** button at **Configure > vSphere Cluster Services > General**.

GPU

- **Reduced fragmentation of vSphere DRS placement for GPUs:** With vSphere 8.0 Update 2, DRS placement of VMs with vGPUs happens on ESXi hosts that are already utilized and have available capacity instead of spreading them over other hosts. This applies for full vGPU and not for partial vGPU profile sizes. For more information, see [Improving VM Placement to Servers to Optimize Your GPU Usage in VMware vSphere 8 Update 2](#) and [DRS Placement of vGPUs](#).
- **vMotion stun time limit for vGPU virtual machines:** Starting with vSphere 8.0 Update 2, you can set a stun time limit per virtual machine when you migrate NVIDIA vGPU-powered virtual machines with vSphere vMotion. Setting a stun time limit can prevent vCenter from powering on the virtual machine or migrating it to a host and network whose estimated maximum stun time exceeds that limit. For more information, see [How to Set a Stun Time Limit for Your vGPU Virtual Machines](#) and [Migrating VMs with vGPUs](#).

Storage

- **Migration of persistent container volumes between datastores in the vSphere Cloud Native Storage:** Starting with vSphere 8.0 Update 2, you can migrate persistent container volumes between datastores in the vSphere Cloud Native Storage by using the vSphere Client. For more information, see [Migrating Container Volumes in vSphere](#).
- **Storage Policy Based Management (SPBM) Support for Asynchronous CNS Queries:** vSphere 8.0 Update 2 adds asynchronous API support for `QueryAssociatedProfiles`, `QueryAssociatedEntity`, `QueryAssociatedEntities`, and `FetchEntityHealthStatusExt` to improve concurrency.
- **Container Storage Interface (CSI) snapshot support for Tanzu Kubernetes Grid Service:** With vSphere 8 Update 2, you can take Persistent Volume snapshot directly with CSI in a standardized way, no need for third-party tools, to increase workload protection. For more information, see [Creating Snapshots in a TKG Cluster in vSphere with Tanzu](#).
- **Offline consolidation improvement with VMFS SE Sparse:** With vSphere 8 Update 2, you see faster offline consolidation, which leads to lower Recovery Point Objective (RPO) and Recovery Time Objective (RTO) with VMFS snapshot offline consolidation operations, without resulting in VMFS version change.
- **Extend online shared disks in clustered applications with vSphere Virtual Volumes:** With vSphere 8.0 Update 2, you can increase the size of shared disks, either physical bus sharing or multi-writer, without taking the application cluster offline and independent of raw device mapping (RDM). For Windows Server Failover Clusters (WSFC), only SCSI is supported. For Oracle RAC, both SCSI and NVMeoF (FC, TCP) are supported. For more information, see [VMware vSphere Virtual Volumes Support for WSFC](#) and [You Can Hot Extend a Shared vVol Disk](#).
- **Support for NVMe in-band migration:** With vSphere 8.0 Update 2, you can manage in-band migration of vSphere Virtual Volumes namespaces between Asymmetric Namespace Access (ANA) groups to allow better balancing of I/O loads across protocol endpoints.
- **Support for third-party Multipathing Plug-In (MPP) for vSphere Virtual Volumes:** vSphere 8.0 Update 2 adds support to third-party MPPs for vSphere Virtual Volumes for NVMe over Fibre Channel and NVMe over TCP datastores.
- **Directory Name Lookup Cache (DNLC) for NFS 4.1:** With vSphere 8.0 Update 2, you can use DNLC for NFS 4.1 to facilitate searches, power on, or listing VMs on large datastores.
- For more information on storage-related features, see [What's New with vSphere 8 Core Storage](#).

vSphere Lifecycle Manager and Configuration Manager

- **Support for vSAN witness node desired state image independent of a vSAN cluster:** With vSphere 8.0 Update 2, you can upgrade a vSAN witness node with a desired state image independent of the vSAN cluster. You no longer

need a dedicated witness per vSAN cluster as shared witness deployments is supported. For more information, see [vSphere Lifecycle Manager and the vSAN Witness Hosts](#).

- **vSphere Configuration Profiles Drafts:** Starting with vSphere 8.0 Update 2, you can create a draft configuration and edit the configuration settings directly in the vSphere Client. For more information, see [Create a Draft Configuration in vSphere Client](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Product Support Notices

Removal of the Disable acceleration option from the Virtual Machine Edit Settings dialog: Starting with vCenter Server 8.0 Update 2, the option to **Disable acceleration** under the **VM Options** tab of the **Virtual Machine Edit Settings** dialog is removed and not supported.

Deprecation of ESX Agent Manager (EAM) VIB lifecycle API: vCenter Server 8.0 Update 2 supports EAM for managing the VIB lifecycle on clusters, but support will drop in a future vSphere release. Instead of creating EAM agents with VIB URLs, you can use vSphere Lifecycle Manager solutions to perform tasks on a cluster level such as installing, updating, and uninstalling a desired solution version on all hosts in a cluster. ESX Agent Manager agent Virtual Machines lifecycle API is not affected by this deprecation.

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 2

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.2.00000-22385739-patch-FP.iso
Build	22385739
Download Size	7856.1 MB
sha256checksum	b8ad59bd3e269011eb0e2f627608cb52fcd14eb8e36c760846a618fc44ac73e2

Download and Installation

Download this patch from [Broadcom Support Portal](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

1. Attach the `VMware-vCenter-Server-Appliance-8.0.2.000000-22385739-patch-FP.iso` file to the vCenter Server CD or DVD drive.
2. Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:
 - To stage the ISO:


```
software-packages stage --iso
```
 - To see the staged content:


```
software-packages list --staged
```
 - To install the staged rpms:


```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Issues

This release resolves CVE-2023-34048 and CVE-2023-34056. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2023-0023](#).

Miscellaneous Issues

You cannot change SMTP credentials in vCenter

As a side effect of vCenter security enhancements, the `vpzd` service might fail to modify a configuration file required to set SMTP credentials in a vCenter instance. When you populate the username field by using either the vSphere Client or the Virtualization Management Object Management Infrastructure (VMOMI), vCenter fails to edit the configuration values. As a result, you cannot change any of the following settings:

- `mail.smtp.server`
- `mail.smtp.password`
- `mail.smtp.username`
- `mail.smtp.port`

This issue is resolved in this release.

vSphere Client and vCenter Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

This issue is resolved in this release.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

This issue is resolved in this release.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

This issue is resolved in this release.

Installation, Upgrade and Migration Issues

You cannot specify IP properties on virtual machines deployed by using a customized OVA on vCenter 8.0 Update 1

While deploying virtual machines by using a customized OVA file on a vCenter 8.0 Update 1 system, the field to provide IP properties of the virtual machine might not appear and you must manually add IP properties after the deployment.

This issue is resolved in this release.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

This issue is resolved in this release.

Known Issues

Installation, Upgrade, and Migration Issues

A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner

During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found.`

Workaround: For more information, see VMware knowledge base articles [94779](#) and [92659](#).

You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade

During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

Source ESX Agent Manager Configuration contains URLs that are not trusted by the System! Verify following URLs and their respective statuses and follow KB 93526.

Workaround: For more information, see VMware knowledge base article [93526](#).

You see vCenter update status as failed although it completes successfully

A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In State._get
using state file /etc/applmgmt/appliance/software_update_state.conf

2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found operation
in progress /storage/core/software-update/install_operation

2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
read JSON file /storage/core/software-update/install_operation [Errno 2] No such file or
directory: '/storage/core/software-update/install_operation'

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is orphaned

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation in
progress is finished

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing to state
file from State._get

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/software_update_state.conf

2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation failed.
Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname

When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/Post install hook failed`.

Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

See this example:

```
root@localhost []# cat /etc/hosts

# Begin /etc/hosts (network card version)

127.0.0.1 localhost.localdomain

127.0.0.1 localhost

::1 localhost ipv6-localhost ipv6-loopback
```

During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)

With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

Workaround: Make sure your file servers are FIPS compliant.

Miscellaneous Issues

You see incorrect Maximum Size value for thin-provisioned virtual disks

In the vSphere Client, when you look at the settings of a VM with a thin-provisioned virtual disk, you might see the **Maximum Size** value for the disk larger than the capacity of the datastore where the disk is located. For example, if a datastore capacity is 100GB of which 90GB are available, and a thin-provisioned virtual disk has 50GB capacity, of which only 10GB are utilized, you might see a **Maximum Size** value of 140 GB, adding the available datastore capacity to the overall disk capacity, not its actual utilization.

Workaround: None

You cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service..` This is expected behavior.

Workaround: None

vSphere Client and vCenter Issues

VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides

the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article [2109887](#).

You do not see the option to push root certificates to vCenter hosts

In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see [Non-disruptive Certificate Management](#).

A scheduled task fails and doesn't schedule further runs

With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see [Scheduling vSphere Tasks](#).

vSphere Lifecycle Manager Issues

You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers

When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

Workaround: Use numeric values in fields that expect numbers.

Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations

Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances

After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

Virtual Machine Management

In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted

In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API: `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag: `KeepAfterDeleteVM`.

Security Issues

You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1

When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary` and `A general runtime error occurred. Key / default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.

2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is not active, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This

issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0 Update 1e Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 1e | 17 JUN 2024 | ISO Build 24005165

Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2024-37079 and CVE-2024-37080. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0012](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 1d Release Notes](#)
- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)
- [VMware vSphere 8.0 Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 1e

Product Patch for vCenter Server containing a VMware security fix.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.1.00500-24005165-patch-FP.iso
Build	24005165
Download Size	7814.7 MB
sha256checksum	3c5151c593af50b113d9b2c78887d98f16c710bcc23fa3caaa3d48403b433291

PRs fixed	NA
CVEs	CVE-2024-37079, CVE-2024-37080

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

1. Attach the `VMware-vCenter-Server-Appliance-8.0.1.00500-24005165-patch-FP.iso` file to the vCenter Server CD or DVD drive.
2. Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:
 - To stage the ISO:


```
software-packages stage --iso
```
 - To see the staged content:


```
software-packages list --staged
```
 - To install the staged rpms:


```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Issues

This release resolves CVE-2024-37079 and CVE-2024-37080. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2024-0012](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as `Failed to get ceip status`. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example vmnic6 and vmnic7) and one management controller (for example vmnic8):

```
:~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
:~] esxcfg-nics -ls
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers vmnic6 and vmnic7 register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller vmnic8 registers first with the vSphere Distributed Switch, the high-speed ethernet controllers vmnic6 or vmnic7 uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
```

```
vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:~] `esxcfg-module -u ionic_en`:~] `esxcfg-module ionic_en`:~] `localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind`.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a

VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.

2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0 Update 1d Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 1d | OCT 24 2023 | ISO Build 22368047
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2023-34048. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2023-0023](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 1c Release Notes](#)
- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)
- [VMware vSphere 8.0 Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 1d

Product Patch for vCenter Server containing a VMware security fix.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.1.00400-22368047-patch-FP.iso
Build	22368047
Download Size	7813.7 MB
sha256checksum	e687c65e6ddaac7318d779463b2befa2939ea7f0f8eab1dbdda0cd c57ca43d94

Download and Installation

Download this patch from [Broadcom Support Portal](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

- Attach the VMware-vCenter-Server-Appliance-8.0.1.00400-22368047-patch-FP.iso file to the vCenter Server CD or DVD drive.

2. Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:

- To stage the ISO:
`software-packages stage --iso`
- To see the staged content:
`software-packages list --staged`
- To install the staged rpms:
`software-packages install --staged`

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Issues

This release resolves CVE-2023-34048. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2023-0023](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as `Failed to get ceip status`. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some nmlx5_core driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.

- `drss` : Use the DRSS parameter.
- `rss` : Use the RSS parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example `vmnic6` and `vmnic7`) and one management controller (for example `vmnic8`):

```
:~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
:~] esxcfg-nics -ls
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers `vmnic6` and `vmnic7` register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller `vmnic8` registers first with the vSphere Distributed Switch, the high-speed ethernet controllers `vmnic6` or `vmnic7` uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:

```
:~] esxcfg-module -u ionic_en
:~] esxcfg-module ionic_en
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind.
```

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the ntg3 driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on ntg3 drivers of versions earlier than 4.1.3 or the tg3 driver. This

issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such

hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0 Update 1c Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 1c | JUL 27 2023 | ISO Build 22088981
Check for additions and updates to these release notes.

What's New

- VMware vCenter Server 8.0 Update 1c delivers bug and security fixes documented in the [Resolved Issues](#) section.
- For VMware vSphere with Tanzu updates, see [VMware vSphere with Tanzu Release Notes](#).
- For Photon OS updates, see [VMware vCenter Server Appliance Photon OS Security Patches](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 1b Release Notes](#)
- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)
- [VMware vSphere 8.0 Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 1c

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.1.00300-22088981-patch-FP.iso
Build	22088981
Download Size	7813.7 MB
sha256checksum	233df6bb919c0b0863e5a3c9f95ff4d61ceae8abdaecfa53ba358d2c2b83c4a4

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Miscellaneous Issues

API command `com.vmware.appliance.version1.resources.*.stats.*` in vCenter Server fails with an error

When you run the API command `com.vmware.appliance.version1.resources.*.stats.*` to get CPU statistics in vCenter Server, you see an error such as:

```
Error in method: Operation Failed. (code com.vmware.applmgmt.err_operation_failed)
```

The issue occurs due to a XML-RPC integer limit.

This issue is resolved in this release.

You cannot migrate only the configuration and inventory of a vCenter by using the vCenter GUI installer

When you perform a vCenter migration by using the GUI installer, if you select only the option **Configuration and Inventory**, you get data for the full list of options, including **Tasks and Events**, and **Performance Metrics**. The issue occurs, because regardless of the selection in the GUI installer, the selection in the backend is for all data.

This issue is resolved in this release.

The timestamp property of the `DatastoreInfo` object continuously changes and might even take old timestamps

In case of shared datastores, timestamps in the `DatastoreInfo` object can vary from host to host. Timestamps indicate the time at which an ESXi host pulls data such as freespace and capacity from the storage stack and different hosts can trigger a refresh at different time. During host sync to update the `datastoreInfo` object, vCenter might get an older timestamp than the one that another host has just sent and you see a fluctuation in the timestamps.

This issue is resolved in this release.

Core files of the SNMP service fill up the `/storage/core` directory

The internal buffer that holds the CPU information of a virtual file system in the file `/proc/cpuinfo` might not be correctly calculated and cause many core files of the SNMP service to generate and fill up the `/storage/core` directory.

This issue is resolved in this release. The fix makes sure that the buffer can hold the entire `cpuinfo` file properly.

After a vCenter upgrade, SNMP might return the base version instead of the upgraded vCenter version information

After a vCenter upgrade, when you run a `snmpwalk` command to get the vCenter version information, the command might return the base version instead of the upgraded version information.

This issue is resolved in this release.

You cannot download `.vmx` files from a vCenter system when host encryption is enabled

If an ESXi host has encryption mode enabled, you cannot download `.vmx` files from this host even when individual VMs are not encrypted, unless you have `Cryptographer.Access` privilege on the vCenter system.

This issue is resolved in this release. For more information, see VMware knowledge base article [92919](#).

Multiple ESXi hosts intermittently lose connectivity but quickly reconnect

If you have an invalid Key Management Server (KMS) configuration in your vCenter system, such as a stale testing instance, the `vpzd` service periodically tries to connect to that KMS to check its status, which leads to wrong file descriptors. As a result, ESXi hosts lose connectivity. Hosts disconnect one at a time but reconnect automatically after 1 second.

This issue is resolved in this release.

Performance of certain nested virtual machines on AMD CPUs might degrade

Nested virtual machines on AMD CPUs with operational systems such as Windows with virtualization-based security (VBS) might experience performance degradation, timeouts, or unresponsiveness due to an issue with the virtualization of AMD's Rapid Virtualization Indexing (RVI), also known as Nested Page Tables (NPT).

This issue is resolved in this release.

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

This issue is resolved in this release.

Storage Issues

The parent resource pool or vApp might not be preserved when performing scheduled VM migration within the same cluster

When you schedule a task to migrate VMs and the source and destination hosts are in the same cluster, the VMs might move out of the parent resource pool or vApp after the migration completes.

This issue is resolved in this release.

Some virtual disks in a cluster with enabled Storage DRS and I/O reservations might get multiple placement recommendations

In some cases, in clusters with enabled Storage DRS and I/O reservations, virtual disks that are pending for placement recommendation might get more than one recommendation. As a result, you see multiple Storage vMotion operations between datastores and possibly datastore reservation violations.

This issue is resolved in this release.

vCenter and vSphere Client Issues

vCenter backup might fail due to a database error

During vCenter backup, config files are copied onto a backup server. A write operation on any of the vCenter database files during the copy operation causes it to fail. As a result, the backup operation also fails.

This issue is resolved in this release. With the fix, instead of directly copying files, vCenter makes a separate DB dump first and then completes the backup.

When you clone a VM, some customization properties might not be as expected

In the vSphere Client, when you clone a VM and try to apply a VM customization specification to the new VM, you might be prompted to provide `subnet mask` and `gateway` values, even if these properties are already configured in the VM customization specification.

This issue is resolved in this release.

In the vSphere Client, you cannot set CPU affinity on a VM with a scheduling affinity range exceeding the number of CPUs on the VM

In the vSphere Client, when you configure CPU affinity on a virtual machine and specify a Scheduling Affinity range that exceeds the number of CPUs on the VM, you might see an error such as `The affinity specification string is not formatted properly`. The error prevents you from applying the specified Scheduling Affinity on the VM, unless you modify the setting to correspond to the actual number of CPUs. If you use the ESXi Host Client, you see no error and the task completes successfully.

This issue is resolved in this release.

You cannot create a scheduled task in the vSphere Client to clone virtual machines on a weekly frequency

In the vSphere Client, when you create a new scheduled task to clone a virtual machine, when you select **Weekly** frequency, you see an error such as `Operation failed!` and the task scheduling does not complete. The workflow works for other schedule frequencies.

This issue is resolved in this release.

If you use a screen reader, you see "Filter Button Filter" instead of "Filter Button" in the Tasks and Events tab in the vSphere Client

If you use a screen reader, when you navigate to **Monitor > Tasks and Events**, you see the filter button in the Tasks and Events tab with a label "Filter Button Filter" instead of "Filter Button".

This issue is resolved in this release.

vSphere vMotion Issues

Migration of ESXi 6.x VMs to 7.0 Update 2 and later fails with an error 'Module 'MonitorLoop' power on failed.'

Migrating a 6.x virtual machine to an ESXi host of version 7.0 Update 2 or later by using vSphere vMotion might fail because the destination host requires larger VM overhead memory than the source host. In the backtrace, you see an error such as `Module 'MonitorLoop' power on failed..`

This issue is resolved in this release. The fix enhances the vSphere DRS advanced option `NewDeltaAlloc`.

vSAN Issues

Transient vSAN health check warning: Network configuration is out of sync

vSAN Skyline health might randomly report that network configuration is out of sync. This transient issue occurs when the vSAN health service uses an outdated vCenter configuration to perform unicast check.

This issue has been resolved in this release.

Networking Issues

If you do not change the default target, VMKNIC port binding for new NFS3 datastores on individual ESXi hosts is not configured

In the vSphere Client, when you create a new NFS3 datastore on an individual ESXi host and select the **Bind to vmknics** check box, you see a default bind target such as `vmk0`. If you do not change the default, the task completes without an error, but port binding is not configured. The issue affects only the workflow for creating a new datastore in the vSphere Client for individual ESXi hosts, not for clusters.

This issue is resolved in this release.

Security Issues

vCenter Server 8.0 Update 1c provides the following security updates:

- Apache Tomcat is updated to version 8.5.89.
- The Spring Framework is updated to version 5.3.27.
- The Jackson package is updated to version 2.15.2.
- The Commons Lang3 software library is updated to version 3.12.0.
- The Commons-Collections4 software library is updated to version 4.4.
- The Commons-Fileupload software library is updated to version 1.5.
- Velocity-Engine-Core is updated to version 2.3.

- The sqlite-jdbc is updated to version 3.42.0.0.
- Protocol Buffers (Protobuf) is updated to version 3.22.3.
- Jetty is updated to version 9.4.51.v20230217.
- See the [PhotonOS](#) release notes for open source changes.

When you do not select a key size while generating a Certificate Signing Request (CSR) in the vSphere Client, the key size might be different from the default 2048 bits

In the vSphere Client, when you generate a CSR for a machine SSL certificate, if you do not change the key size value, you might see a different value from the default 2048 bits after the task completes, for example 3072 bits.

This issue is resolved in this release. If you already face the issue, change the key size to another value than 2048 and then revert to 2048.

Virtual Machine Management Issues

Migration of multiple virtual machines across datacenters might fail with an error Operation Failed!

In rare cases, when you migrate many VMs across datacenters, an issue with a property computation might cause the migration operation to fail with an error `Operation Failed!`.

This issue is resolved in this release.

Upgrade Issues

Patching a vCenter to 8.0 Update 1 by using command line might fail with the error 'Invalid Type, expected String, instead got NoneType'

Patching a vCenter to 8.0 Update 1 by using command line with the patch iso disk mounted from a local system with the VMware Remote Console (VMRC) might fail with an error such as `Invalid Type, expected String, instead got NoneType`. The issue occurs when the vCenter VM is managed by the vCenter on which the patching is attempted and VMRC is accessed by the same vSphere Client. As a result, the VMRC console loses connection to the iso mounted from the local system as vCenter services stop during the patching.

This issue is resolved in this release.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as `Failed to get ceip status`. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example vmnic6 and vmnic7) and one management controller (for example vmnic8):

```
:~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
:~] esxcfg-nics -ls
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers vmnic6 and vmnic7 register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller vmnic8 registers first with the vSphere Distributed Switch, the high-speed ethernet controllers vmnic6 or vmnic7 uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
```

```
vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:~] `esxcfg-module -u ionic_en`:~] `esxcfg-module ionic_en`:~] `localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind`.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a

VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.

2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0 Update 1b Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 1b | JUNE 22 2023 | Build 21860503
Check for additions and updates to these release notes.

What's New

- This release resolves CVE-2023-20892, CVE-2023-20893, CVE-2023-20894, CVE-2023-20895, and CVE-2023-20896. For more information on these vulnerabilities and their impact on VMware products, see [VMSA-2023-0014](#).
- VMware vCenter Server 8.0 Update 1b delivers bug fixes documented in the [Resolved Issues](#) section.

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 1a Release Notes](#)
- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)
- [VMware vSphere 8.0 Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 1b

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.1.00200-21860503-patch-FP.iso
Build	21860503
Download Size	7965.6 MB
sha256checksum	2c7df46fc563212f55db509331be8179e3dc8774f9ccacb8806add55d948f08

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Upgrade Issues

vCenter 8.x installation might fail during firstboot with an error for a missing install parameter

While installing vCenter 8.x, you might see the error `cis.exceptions.InstallParameterException` during the firstboot step due to a missing `upgrade.import.directory` install parameter. The error occurs if vCenter does not find the parameter due to VM restarts, or restores to a VM Snapshot without memory before running the firstboot configuration.

This issue is resolved in this release. The fix makes sure that the `upgrade.import.directory` install parameter is correctly passed during vCenter installation.

Installation and upgrades to vCenter 8.0 Update 1 fail when a vCenter Single Sign-On domain contains uppercase letters

If your vCenter Single Sign-On domain contains uppercase letters such as `administrator@vSphere.Local` or `administrator@VSPHERE.LOCAL`, installation, patching and upgrades to vCenter 8.0 Update 1 might fail. During installation or upgrade from an earlier version, in the install wizard, you see an error such as `Failed to start the Workload Control Plane Service`. During patching from 8.x, in the Virtual Appliance Management Interface, you see the error `Exception occurred in postInstallHook`. Messages in the `/var/log/vmware/wcp/wcpsvc.log` file shows a failure such as:

```
2023-04-17T13:04:54.017Z error wcp [gcauth/gc_auth.go:94] Failed to parse VC JWKS: invalid character '<' looking for beginning of value
```

```
2023-04-17T13:04:54.017Z error wcp [gcauth/gc_auth.go:59] Unable to get VC public key configuration : invalid character '<' looking for beginning of value
```

```
2023-04-17T13:04:54.017Z fatal wcp [gcauth/gc_auth.go:47] Unable to get VC public key configuration : invalid character '<' looking for beginning of value
```

You can see the case of an SSO domain by using the following command: `/usr/lib/vmware-vmafd/bin/vmafd-cli get-domain-name --server-name localhost`. Example output of the above command might be `VSPHERE.LOCAL`.

This issue is resolved in this release.

Miscellaneous Issues

If no vSphere vMotion bandwidth for migration of VMs with vGPUs is available, the vpxd service might fail at regular intervals

If vpxd restarts during a migration of VMs with vGPUs and for some reason vSphere vMotion has no bandwidth available, vpxd might start failing on regular intervals.

This issue is resolved in this release. If you already face the issue, restart the vpxd service, re-add ESXi hosts using the Quickstart workflow, and reconfigure the vSphere vMotion network to make sure bandwidth is available.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as `Failed to get ceip status`. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.

2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example vmnic6 and vmnic7) and one management controller (for example vmnic8):

```
:~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
:~] esxcfg-nics -ls
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers `vmnic6` and `vmnic7` register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -
n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller `vmnic8` registers first with the vSphere Distributed Switch, the high-speed ethernet controllers `vmnic6` or `vmnic7` uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:

```
:~] esxcfg-module -u ionic_en
:~] esxcfg-module ionic_en
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind.
```

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

```
Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server
version cannot be validated. They may not function properly after vCenter Server upgrade.
```

```
Resolution: Please contact the plug-in vendor and make sure the package is compatible with
the new vCenter Server version.
```

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

Networking Issues

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed in the backtrace`. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') . The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.`

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_XXX. Extraction of image from host XXX.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0 Update 1a Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 1a | JUNE 01 2023 | Build 21815093
Check for additions and updates to these release notes.

What's New

- VMware vCenter Server 8.0 Update 1a delivers a fix for a vCenter upgrade issue due to a rare race condition. This patch is the base vCenter Server Bill of Materials (BOM) version for VMware Cloud Foundation (VCF) 5.0. For vCenter systems without a VCF environment, this patch is optional. In the rare case of encountering the issue, see VMware knowledge base article [92462](#) or the [Resolved Issues](#) section.

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0 Update 1 Release Notes](#)
- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)
- [VMware vSphere 8.0 Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 1a

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.1.00100-21815093-patch-FP.iso
Build	21815093
Download Size	7965.5 MB
sha256checksum	0ad65465fc57d7c164b22748abaa3b258c1b541fa7225cdab6a4b0be986b902b

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Upgrade Issues

vCenter upgrade might fail due to a race condition with the statsmonitor service

vCenter upgrade might fail while starting the statsmonitor service due to a race condition with another service that creates entries in the system metrics DB of the vCenter system.

This issue is resolved in this release.

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as `Failed to get ceip status`. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example vmnic6 and vmnic7) and one management controller (for example vmnic8):

```
:~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
:~] esxcfg-nics -ls
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers vmnic6 and vmnic7 register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -
n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller vmnic8 registers first with the vSphere Distributed Switch, the high-speed ethernet controllers vmnic6 or vmnic7 uplink might end up operating with RSS set to 1 receive queue.:

```
localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
```

```
vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:~] `esxcfg-module -u ionic_en`:~] `esxcfg-module ionic_en`:~] `localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind`.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as *An error occurred while saving the snapshot: A digest operation has failed* in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0 Update 1 Release Notes

This document contains the following sections

- [Introduction](#)
- [General Availability](#)
- [What's New](#)
- [Earlier Releases of vCenter Server 8.0](#)
- [Installation and Upgrades for This Release](#)
- [Compatibility](#)
- [Product Support Notices](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0 Update 1 | 18 APR 2023 | GA ISO Build 21560480
Check for additions and updates to these release notes.

General Availability

This vCenter Server 8.0 Update 1 release is a General Availability (GA) designation. For more information on the vSphere 8.0 IA/GA Release Model of vSphere Update releases, see [The vSphere 8 Release Model Evolves](#).

What's New

- **vSphere Configuration Profiles:** vSphere 8.0 Update 1 officially launches vSphere Configuration Profiles, which allow you to manage ESXi cluster configurations by specifying a desired host configuration at the cluster level, automate the scanning of ESXi hosts for compliance to the specified Desired Configuration and remediate any host that is not compliant. vSphere Configuration Profiles require that you use vSphere Lifecycle Manager images to manage your cluster lifecycle, a vSphere 8.0 Update 1 environment, and Enterprise Plus or vSphere+ license. For more information, see [Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level](#).
- **With vSphere 8.0 Update 1, vSphere Distributed Services Engine adds support for:**
 - NVIDIA BlueField-2 DPUs to server designs from Lenovo (Lenovo ThinkSystem SR650 V2).
 - 100G NVIDIA BlueField-2 DPUs to server designs from Dell.
 - UPTv2 for NVIDIA BlueField-2 DPUs.
 - AMD Genoa CPU based server designs from Dell.
- **Recorder for required privileges on vCenter Server workflows:** With vSphere 8.0 Update 1, to help you identify the minimal set of privileges that are required to run a vCenter Server workflow, you can record workflows and store the required privileges. You can use the stored data to create specific roles for users that run the same workflows, similar to templates. Searching through vCenter Server workflow recordings allows you to retrieve lists of privilege checks along with the corresponding sessions, users, managed objects, and operation IDs (opIDs). For more information, see [Performing Privilege Checks Operations](#).
- **Support for heterogeneous virtual graphics processing unit (vGPU) profiles on the same GPU hardware:** vSphere 8.0 Update 1 removes the requirement that all vGPUs on a physical GPU must be of the same type and you can set different vGPU profiles, such as compute, graphics, or Virtual Desktop Infrastructure workload, on one GPU to save cost by higher GPU utilization and reduced workload fragmentation.
- **Integration of VMware Skyline™ Health Diagnostics™ with vCenter:** Starting with vSphere 8.0 Update 1, you can detect and remediate issues in your vSphere environment by using the VMware Skyline Health Diagnostics self-service diagnostics platform, which is integrated with the vSphere Client. For more information, see [VMware Skyline Health Diagnostics for vSphere Documentation](#).

- **VM-level power consumption metrics:** Starting with vSphere 8.0 Update 1, you as a vSphere admin can track power consumption at a VM level to support the environmental, social, and governance goals of your organization.
- **NVSwitch support:** vSphere 8.0 Update 1 adds support to up to 8 GPUs with NVSwitch connections between them on a single ESXi host. With the NVSwitch technology, you can run high-performance computing (HPC) and AI applications such as deep learning, scientific simulations, and big data analytics, which require multiple GPUs working together in parallel.
- **Okta Identity Federation for vCenter:** With vSphere 8.0 Update 1, you can configure vCenter Server Identity Provider Federation for Okta as an external identity provider. For more information, see [Configure vCenter Server Identity Provider Federation for Okta](#).
- **Support for Fault Tolerance of virtual machines that use a virtual TPM (vTPM) module:** With vSphere 8.0 Update 1, you can use FT for VMs with a vTPM module to ensure continuous availability and security for your mission critical VMs.
- **Quick Boot support for servers with TPM 2.0 chips:** With vSphere 8.0 Update 1, you no longer need to disable TPM 2.0 for Quick Boot, which allows you to save time and eliminate security gaps.
- **vSphere API for Storage Awareness (VASA) version 5 for vSphere Virtual Volumes:** VASA version 5 for vSphere Virtual Volumes enhances security, certification management for multi vCenter deployments, and usability support. VASA version 5 deprecates self-signed certificates but provides backward compatibility support for earlier VASA versions. Before upgrading to ESXi 8.0 Update 1, see VMware knowledge base article [91387](#) to ensure continued availability of vSphere Virtual Volumes datastores.
- **Sidecar files become regular files in Config-vVol instead of vSphere Virtual Volumes objects:** To improve scalability and performance of vSphere Virtual Volumes, starting with vSphere 8.0 Update 1, sidecar files are no longer vSphere Virtual Volumes objects but regular files in Config-vVol. When sidecar files are created as vSphere Virtual Volumes objects, for solutions such as First Class Disk (FCD) that create a large number of small sidecar files, they can cause an overhead of VASA operations such as binding and unbinding the volumes to a protocol endpoint. Creating sidecar files under the Config-vVol home directory removes such overhead from VASA and storage. After updating to ESXi 8.0 Update 1, new virtual machines, or virtual disks in the Config-vVol namespace are not supported on ESXi hosts of earlier versions. For more information and resolution, see VMware knowledge base article [90791](#).
- **Increased default capacity for vSphere Virtual Volumes objects of type Config-vVol:** Starting with vSphere 8.0 Update 1, to allow the use of folders over vSphere Virtual Volumes datastores as content repositories, vSphere Virtual Volumes objects of type Config-vVol are created by default as thin-provisioned 255 GB, up from 4 GB in earlier releases. Also, starting with ESXi 8.0 Update 1, Config-vVol objects use VMFS-6 format instead of VMFS-5.
- **NVMe over TCP support for vSphere Virtual Volumes:** vSphere 8.0 Update 1 adds NVMe over TCP support for vSphere Virtual Volumes.
- **Extended XCOPY support:** With vSphere 8.0 Update 1, you can use `EXTENDED COPY` (`XCOPY`) commands to optimize the data copy between VMFS datastores that are in different storage arrays, not only between datastores within the same array. `XCOPY` operations across storage arrays are supported only for VMFS 6 datastores and allow you to offload, migrate and clone workloads. While ESXi enables this feature, the actual data migration across the arrays must be implemented on the storage array side. For more information on how to use SCSI T10 commands, see [Managing Hardware Acceleration on Block Storage Devices](#).
- **New file type for OSDATA volumes on SSD devices:** vSphere 8.0 Update 1 adds a new file system type, VMFSOS, specifically for the ESX-OSData system partition on local SSD devices, which allows you to continue using virtual flash resources on other devices. The new file type prevents cases when you format an ESX-OSData volume on a local SSD device, and `fsType` returns a file of type Virtual Flash File System (VFFS). As a result, the disk backing of the ESX-OSData volume is listed under the Virtual Flash resources in vCenter, but such a disk belongs to the ESX-OSData volume and is not a part of the Virtual Flash resource pool.
- **NFS traffic isolation enhancements:** With vSphere 8.0 Update 1, to isolate NFS traffic, you can bind an NFS3 datastore on an NFS share over the network to a VM kernel adaptor on an ESXi host from a cluster. VMKNIC port binding is not supported in this release for vSphere Virtual Volumes datastores backed with NFS3 configurations. For more information, see [Configure VMkernel Binding for NFS 3 Datastore](#).

- **Fourfold increase in NVMe-oF namespace capacity:** Starting with vSphere 8.0 Update 1, you can use 32 paths to a NVMe-oF namespace from 8 in earlier releases.
- **Support for end-to-end NVMe stack without protocol translation:** With vSphere 8.0 Update 1, the Pluggable Storage Architecture (PSA) extends NVMe capabilities to support end-to-end NVMe stack without any protocol translation in any of the layers.
vSphere 8.0 Update 1 also extends NVMe functionality by adding support for third-party multi-pathing plug-ins to control and manage NVMe arrays.
- **Increased maximum for Windows Server Failover Clusters (WSFC):** vSphere 8.0 Update 1 increases the maximum WSFC clusters per ESXi host to 16 from 3.
- **Scale the number of NVMe over TCP adapters:** vSphere 8.0 Update 1 increases the number of NVMe over TCP storage adapters per ESXi host to 8 from 2 to allow more redundancy and better bandwidth.
- **Support hot-add and hot-remove for VMDirectPath I/O devices:** With vSphere 8.0 Update 1, by using vSphere API you can add or remove a VMDirectPath I/O device without powering off VMs. For more information, see [Hot-add Support for VMDirectPath I/O Devices](#).
- **Increased number of PCI passthrough devices per VM:** With vSphere 8.0 Update 1, you can add up to 64 PCI passthrough devices per VM, up from 32 in previous releases, to allow better support for Virtualized Radio Access Network (vRAN) use cases.
- **Local depot overrides for Remote Office and Branch Office (ROBO) standalone ESXi hosts:** Starting with vSphere 8.0 Update 1, to meet demand of certain business use cases, you can manage local depot overrides for Remote Office and Branch Office (ROBO) standalone ESXi hosts in the vSphere Client and by using the vSphere Automation API. For more information, see [Manage Depot Overrides for a Cluster or a Standalone Host](#).
- **Advanced filters in the vSphere Client:** With vSphere 8.0 Update 1, in addition to the existing quick filter option, you can set advanced filters on most inventory objects in the vSphere Client, such as VM lists, Hosts and Clusters, Datastores, and linked vCenters. You can use strings, such as Name, Status, and Guest OS, and enumerations, such as IP address. To refine your filters, you can use various logical operators and 2 conditions.
- **Increased maximum number of NFSv3 datastores that can be mounted with multiple connections:** Starting with vSphere 8.0 Update 1, you can configure up to 8 connections to NFSv3 datastores on an ESXi host, from 1 in earlier releases, and have up to 256 connections used by the NFSv3 datastores on the host, depending on the combination of the number of datastores and the number of connection configured for each datastore. For more information, see VMware knowledge base article [91481](#).
- **HTTP/JSON-based wire protocol as an alternative to SOAP/XML:** vSphere 8.0 Update 1 adds support for a new HTTP/JSON-based wire protocol as an alternative to SOAP/XML. The new protocol is described using the industry standard OpenAPI specification, version 3.0, and you can use it to access the popular VIM APIs. For more information, see the OpenAPI schema and Web Services Programming Guide, part of the [vSphere Management SDK](#).

Earlier Releases of vCenter Server 8.0

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- [VMware vCenter Server 8.0c Release Notes](#)
- [VMware vCenter Server 8.0b Release Notes](#)
- [VMware vCenter Server 8.0a Release Notes](#)
- [VMware vSphere 8.0 Release Notes](#)

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Installation and Upgrades for This Release

Before upgrading to VASA version 5, refer to VMware knowledge base article [91387](#) to ensure continued availability of your vSphere Virtual Volumes datastores.

Compatibility

Newly created virtual machines or virtual disks of version ESXi 8.0 Update 1 in the Config-vVol namespace of vSphere Virtual Volumes are not supported on ESXi hosts of earlier versions. For more information and resolution, see VMware knowledge base article [90791](#).

Product Support Notices

- Since the SHA-1 signature algorithm is no longer supported in vSphere 8.0, if you use Enhanced Linked Mode (ELM) to connect vCenter systems of versions 7.0x and 8.0x, you must not add a SHA-1 trusted root certificate to the 7.0x vCenter instance during upgrades, because the SHA-1 certificate might be propagated to the 8.0x vCenter and cause unexpected results, such as disconnecting ESXi hosts.
If you already have 7.0x vCenter instances connected with ELM, and using a SHA-1 trusted root certificate, do not add an 8.0x vCenter instance to the group, because the SHA-1 certificate might be replicated to the 8.0x instance.

Patches Contained in This Release

Patch for VMware vCenter Server 8.0 Update 1

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.1.00000-21560480-patch-FP.iso
Build	21560480
Download Size	7965.6 MB
sha256checksum	1016f37c466e46957879f4084e176e62d96803fb0d1ad16cb1b8f6fc0e852b24

Download and Installation

Download this patch from [Broadcom Support Portal](#).

For download instructions for earlier releases, see [Download Broadcom products and software](#).

- Attach the `VMware-vCenter-Server-Appliance-8.0.1.00000-21560480-patch-FP.iso` file to the vCenter Server CD or DVD drive.
- Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:
 - To stage the ISO:


```
software-packages stage --iso
```
 - To see the staged content:


```
software-packages list --staged
```
 - To install the staged rpms:


```
software-packages install --staged
```

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Features Issues

Enabling TLS 1.0 on ESXi 8.0 hosts causes connection drops

ESXi 8.0 supports OpenSSL 3.0 and the only TLS protocol enabled by default is TLS 1.2. If you try to enable TLS 1.0 in `/UserVars/ESXiVPsDisabledProtocols` by using ESXCLI commands, this leads to connection drops.

This issue is resolved in this release.

Miscellaneous Issues

You might see wrong metadata for delete operation of a vSphere Lifecycle Manager depot that you use to create an image to manage a standalone ESXi host

Starting with vSphere 8.0, you can create an image based on a vSphere Lifecycle Manager depot, online or offline, to manage the lifecycle of any standalone ESXi host that is part of your vCenter Server inventory. In rare cases, when you delete such a depot, the metadata for the delete task might not be correct. For example, you see the `hostid` and `hostIP` details populated under `clusterName` and `clusterID`, such as:

```
clusterName = 10.161.153.136,
clusterId = host-54,
entityName = <null>,
entityId = <null>,
```

The issue also occurs when you delete a depot downloaded by using VMware vSphere Update Manager Download Service (UMDS), which is part of the desired state of a standalone host. The issue has no impact on any vSphere Lifecycle Manager operations and affects only task metadata for standalone hosts.

This issue is resolved in this release.

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

This issue is resolved in this release.

Server Configuration Issues

Changing an Input/Output Operations Per Second (IOPS) limit might cause a significant drop in the I/O throughput of virtual machines

When you change an IOPS limit based on the Storage I/O Control (SIOC) by using a Storage Policy Based Management (SPBM) policy, you might see significantly slower VM performance. Normally, when you set an SPBM policy, IOPS limits are handled by an I/O filter, while `mClock`, the default I/O scheduler, handles reservations and shares. Due to a logic fault, when you change an existing IOPS limit, I/Os might throttle at the `mClock` scheduler instead at the I/O filter. As a result, I/Os get with a significant delay to the I/O filter which causes a drop in the I/O throughput of virtual machines.

This issue is resolved in this release. The fix makes sure that IOPS limits are handled by the I/O filter, while mClock handles the reservations and shares. For more information, see VMware knowledge base article [89951](#).

Known Issues

Installation, Upgrade, and Migration Issues

You cannot specify IP properties on virtual machines deployed by using a customized OVA on vCenter 8.0 Update 1

While deploying virtual machines by using a customized OVA file on a vCenter 8.0 Update 1 system, the field to provide IP properties of the virtual machine might not appear and you must manually add IP properties after the deployment.

Workaround: If the VM deployment succeeds, you can manually set the IP properties from the **vApp Options** under the **Configure** tab of the newly deployed VM. If the VM deployment fails, you must first add the property `ovf:required="false"` to the properties with attribute `vmw:qualifiers='Ip'` in the OVF file before applying the workaround. For more information, see VMware knowledge base article [93677](#).

Upgrades to vCenter Server 8.0 Update 1 fail with an error Exception occurred in postInstallHook

If a vCenter Single Sign-On domain contains a capitalization such as `administrator@vSphere.Local` or `administrator@VSPHERE.LOCAL`, upgrades to vCenter Server 8.0 Update 1 might fail with the error `Exception occurred in postInstallHook`.

In the vSphere Client, you see an error such as `wcp service failed to start`.

In the `PatchRunner.log`, you see an error such as `An error occurred while starting service 'wcp'`.

In the `wcpSvc.log`, you see the following errors: `Failed to parse VC JWKS: invalid character '<' looking for beginning of value` or `Unable to get VC public key configuration : invalid character '<' looking for beginning of value`.

After such failure, you must restore vCenter from a backup or revert from a snapshot.

Workaround: See VMware knowledge base article [92436](#).

Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host

Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

1. You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.
2. You trigger a page refresh after the image compliance reports are generated.

In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state

An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article [91260](#).

You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1

During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as `Failed to get ceip status`. This is expected and does not indicate an actual issue with the vCenter system.

Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

Miscellaneous Issues

In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

Networking Issues

Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

Workaround: Enable the hotplug functionality as described in [Hot-add and Hot-remove support for VMDirectPath I/O Devices](#).

Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail

With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

vSphere Lifecycle Manager Issues

You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task

In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic [Configure the vSphere Lifecycle Manager Automatic Download Task](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support the auto negotiate option, or the option is deactivated, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed

machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of `BMCNetworkEnable` is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example `vmnic6` and `vmnic7`) and one management controller (for example `vmnic8`):

```
:~] esxcfg-nics -l

vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller

vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller

:~] esxcfg-nics -ls

vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers `vmnic6` and `vmnic7` register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller `vmnic8` registers first with the vSphere Distributed Switch, the high-speed ethernet controllers `vmnic6` or `vmnic7` uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:

```
:~] esxcfg-module -u ionic_en
:~] esxcfg-module ionic_en
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind.
```

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS.`

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed in the backtrace`. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') . The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.`

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail

If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article [89493](#).

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_XXX. Extraction of image from host XXX.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

VMware vCenter Server 8.0c Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0c | 30 MAR 2023 | ISO Build 21457384
Check for additions and updates to these release notes.

What's New

- VMware vCenter Server 8.0c Release Notes serves as a vehicle for the [VMware vSphere with Tanzu Release Notes](#) and does not deliver vCenter Server fixes.

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0c

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.0.10300-21457384-patch-FP.iso
Build	21457384
Download Size	6471.3 MB
sha256checksum	336f03622ba8fb1e84a715bfdbcc82e5934d15a10971bf1f5e1d8724f459bc01

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

vCenter Server update fails with error for failed load of repository manifest data

If you set up a custom repository that has authentication and self-signed certificates that vCenter Server does not trust, updates and upgrades to vCenter Server 8.0 and later might fail. In the vSphere Client, you see an error such as `Failed to load the repository manifest data for the configured upgrade after the upgrade precheck completes`.

Workaround: See VMware knowledge base article [90259](#).

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail.

as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

You cannot revert configuration changes for the `vmware-rbd-watchdog` service

When you use the TLS Configuration utility to revert configuration changes, the `vmware-rbd-watchdog` service does not restore back to the expected TLS protocols.

Workaround: Use the argument `-s <service-name>` to individually call the `rbd` service, for example:

```
root@sc1-10-168-175-98 [ /usr/lib/vmware-TlsReconfigurator/VcTlsReconfigurator ]# ./
reconfigureVc restore -d /tmp/20221124T100959 -s vmware-rbd-watchdog
```

In the vSphere API Explorer, VMware Datacenter CLI (DCLI) and PowerCLI, you see an API option "contentinternal" that is not functional

You see an API option **contentinternal** in the metadata of either the vSphere API Explorer, DCLI and PowerCLI. For example, when you open <https://<your vCenter IP>/ui/app/devcenter/api-explorer>, you see the option in the **select API** drop-down menu. This option is not functional.

Workaround: Ignore the **contentinternal** API option and do not use it.

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reinstall ESXi. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.`

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS.`

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the `icen` driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You

see an error such as `An error occurred while saving the snapshot: A digest operation has failed in the backtrace. The snapshot creation task completes upon retry.`

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

VMware vCenter Server 8.0b Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Product Support Notices](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0b | 14 FEB 2023 | ISO Build 21216066
Check for additions and updates to these release notes.

What's New

- For fixes and known issues in vCenter Server 8.0b, see the sections [Resolved Issues](#) and [Known Issues](#).
- For VMware vSphere with Tanzu updates, see [VMware vSphere with Tanzu Release Notes](#).
- For Photon OS updates, see [VMware vCenter Server Appliance Photon OS Security Patches](#).

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

Patch for VMware vCenter Server 8.0b

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.0.10200-21216066-patch-FP.iso
Build	21216066
Download Size	6451.6 MB
sha256checksum	bea9e7499da68f0ea0a89594db068bc9bb272e10e8f7c4d56097950b93c2bd8f

Download and Installation

[Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Product Support Notices

Deprecation of TLS 1.0 and 1.1: With vCenter Server 8.0b, TLS 1.0 and 1.1 are deprecated and will be removed in upcoming vCenter Server release.

Resolved Issues

Security Issues

If you use Integrated Windows Authentication (IWA) for vCenter Server authentication, you might not be able to log in to the vCenter Server system because the STS service (vmware-stsd) might fail

If you use IWA for vCenter Server authentication, the STS service might fail during attempts to log in to the vCenter Server system, because the Active Directory might return a `null` LDAP hostname. Logins fail until the STS service restarts.

The `vMonCoredumper.log` file contains entries similar to the following:

```
2022-12-29T07:16:15.305Z In(05) host-12345 Notify vMon about pool-2-thread-4 dumping core.
Pid : 3315
```

```
2022-12-29T07:16:15.320Z In(05) host-12345 Successfully notified vMon.
```

```
2022-12-29T07:16:16.785Z In(05) host-12345 Successfully generated core file /var/core/
core.pool-2-thread-4.3315.
```

The `hs_err_sts_pid*.log` file contains entries such as:

```
# A fatal error has been detected by the Java Runtime Environment:
```

```
#
```

```
# SIGSEGV (0xb) at pc=0x00007f4d80203d06, pid=3315, tid=0x00007f4ce3d1f700
```

This issue is resolved in this release. Due to performance issues with IWA-based authentication, vSphere 8.0 deprecates the use of IWA. You can use AD over LDAP or ADFS.

vCenter Server 8.0b provides the following security updates:

- The Gson Java library is updated to version 2.9.0.
- The SQLite JDBC library is updated to version 3.39.2.1.
- The PostgreSQL JDBC is updated to version 42.5.0.
- Spring Security is updated to version 5.7.5.
- The Jackson package is updated to version 2.14.0.
- The Spring library is updated to version 5.3.22.
- Apache Tomcat is updated to version 9.0.68/8.5.82.
- Apache Commons Configuration is updated to version 2.8.0.
- Eclipse Jetty is updated to version 9.4.48.v20220622.
- vPostgres is updated to version 13.8.
- Apache Commons Text library is updated to version 1.10.0.

- See the [PhotonOS](#) release notes for open source changes.

The field `organizationalUnitName` is required for Secure Sockets Layer (SSL) certificates

Starting from vCenter Server 8.0.b, if you use vSphere Automation API, Certificate Manager utility, or certificate management CLIs as an interface to manage vCenter Server certificates, you can leave the `organizationalUnitName` field empty. The change does not affect existing workflows for certificate replacement or Certificate Signing Requests (CSRs). The vSphere Client and PowerCLI 12.4 still require data for the `organizationalUnitName` field.

This issue is resolved in this release.

Networking Issues

A VM clone might fail with an error and the newly created dvPort remains in the vCenter Server database

If a virtual machine clone spec has the `guestinfo.resync.mac.addr` key in the `extraConfig` elements, such as the following: `<ovf:Info>Virtual hardware requirements</ovf:Info><vmw:ExtraConfig vmw:key="guestinfo.resync.mac.addr"/>`, a cloning operation might fail, because in some cases, a value for the key might not pass and remains unset. In the vpxd logs, you see with an error such as `vmodl.fault.InvalidArgument` and error messages like `A specified parameter was not correct: config.extraConfig["guestinfo.resync.mac.addr"]`. As a result, the newly created dvPort for the failed clone remains unused in the vCenter Server database.

This issue is resolved in this release. The fix makes sure to remove unused dvPort from the vCenter Server database. If the error persists, you must remove the `guestinfo.resync.mac.addr` key in the `extraConfig` elements or set a valid MAC address as value.

A link aggregation group (LAG) might intermittently go down and up upon a vSphere Distributed Switch uplink removal

After a removal of a free VDS uplink, a LAG might go down and up in a short time. As a result, virtual machines that use that LAG as uplink might temporarily lose network connectivity.

This issue is resolved in this release.

GET VM REST API fails with an internal server error

When you use a REST API function such as `/rest/vcenter/vm/{vm-id}`, the call might fail with an error message such as `Internal server error: Error: Http error 500 while requesting '/rest/vcenter/vm/vm-xx'`. The issue occurs because the `networkBootProtocol` value might not persist in the vCenter Server database after reconfiguring the VM. As a result, when vCenter Server restarts, the `networkBootProtocol` value is not available and the function `/rest/vcenter/vm/{vm-id}` fails.

This issue is resolved in this release. The fix makes sure that the fixed now and the `networkBootProtocol` value persists in the vCenter Server database after a VM reconfigure and vCenter Server restart.

Downloading small files over NFC stops or fails intermittently

If your vSphere system is busy with multiple calls to connect to ESX hosts to download files from datastores by using NFC, many attempts to retry a delayed or failed call might accumulate unnecessary load of datastore refresh operations. As a result, NFC performance aggravates and downloading small files might stop or fail intermittently.

This issue is resolved in this release. The fix reduces unnecessary load of datastore refresh operations.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

This issue is resolved in this release.

vSphere Client Issues

In the vSphere Client, when you edit a virtual machine settings, you see duplicate entries in the SCSI Bus Sharing drop-down menu

In the vSphere Client, when you create a new virtual machine or reconfigure an existing one, under **Edit Settings > Virtual Hardware > SCSI controller 0 > SCSI Bus Sharing**, you might see duplicate entries in the drop-down menu.

This issue is resolved in this release.

Virtual Machine Management Issues

vSphere Storage DRS might stop load balancing due to an I/O reservation computation issue

In rare cases, vSphere Storage DRS might stop load balancing due to an I/O reservation computation issue. For example, if some VMs in a storage cluster have storage I/O reservations, vSphere Storage DRS might fail to balance the storage cluster and throw an error such as `The available storage IOPs capacity is not sufficient for the operation.`

This issue is resolved in this release. The manual workaround is to set the vSphere Storage DRS advanced option `EnforceIOReservations` to 0.

Storage Issues

Disabling Storage I/O Control on datastores results in high read I/O on all ESXi hosts

A rare issue with calculating the Max Queue Depth per datastore device might cause queue depth mismatch warnings from the underlying SCSI layer. As a result, when you disable Storage I/O Control on datastores, you might see high read I/O on all ESXi hosts.

This issue is resolved in this release.

Miscellaneous Issues

The vpxd service might fail and generate a core file due to bad authentication data when performing vCenter Server tagging operations

A rare issue with the authentication of the tagging API might cause the vpxd service to fail and generate a core file in the `/var/core` directory. Such failures of the vpxd service might in turn cause unexpected failovers to passive nodes in vCenter Server High Availability environments.

This issue is resolved in this release.

Rsyslog does not write log messages of the VMware Directory Service (VMDIRD)

The rsyslog daemon does not write log messages to the `/vmdird` directory and you cannot use vmdird logs for debugging.

This issue is resolved in this release.

In the vSphere Client, you cannot export lists of Inventory objects

In the vSphere Client, when you try to export a file with a list of Inventory objects, such as VMs, Hosts, and Datastores, the task fails with an error `Export Data Failure`.

This issue is resolved in this release.

Enter Maintenance Mode task fails when you try to perform maintenance on ESXi hosts where instant clones reside

In some cases, when you try to put in maintenance mode the ESXi hosts where instant clones reside, the task might fail with a `ManagedObjectNotFound` error. The issue occurs because some virtual machine objects enqueued for evacuation might be deleted while in-flight.

This issue is resolved in this release.

In very rare cases, Virtual Desktop Infrastructure (VDI) tasks might fail as result of intermittent vCenter Server unresponsiveness

Rarely, as VM events that potentially lead to a bad state accrue, this could lead to a VM misconfiguration and in very rare cases, to a missing VM layout. As a result, vCenter Server might become intermittently unresponsive when querying the database for a VM layout that does not exist. In such cases, VDI tasks might also fail.

This issue is resolved in this release.

Storage I/O Control quickly generates a large volume of logs marked as critical that might also fill up the datastore

In VMware Aria Operations for Logs, formerly vRealize Log Insight, you might see a large volume of logs generated by Storage I/O Control such as `Invalid share value: 0. Using default.` and `Skipping device naa.xxxx` either due to VSI read error or abnormal state. The volume of logs varies depending on the number of ESXi hosts in a cluster and the number of devices in switched off state. When the issue occurs, the log volume generates quickly, within 24 hours, and VMware Aria Operations for Logs might classify the messages as critical. However, such logs are harmless and do not impact the operations on other datastores that are online.

This issue is resolved in this release. The fix moves such logs from error to trivia to prevent misleading logging.

Installation and Upgrade Issues

Upgrade to vCenter Server 8.0 fails with a duplicate key error

During vSphere vMotion operations across vCenter Server instances in Enhanced Linked Mode, for VMs that have vSphere tags, it is possible that the tag association creates a record in the `CIS_KV_KEYVALUE` table in the vCenter Server DB, along with the records in the `vpvpx_entity` table. During upgrade operations, vCenter Server might try to copy the tag associations record from the `CIS_KV_KEYVALUE` table to the `vpvpx_entity` table. As a result, the upgrade fails with an error such as `psycopg2.IntegrityError: duplicate key value violates unique constraint "pk_vpvpx_entity_tag_xref"`.

This issue is resolved in this release.

Upgrades to vCenter Server 8.0 might fail during the data import phase due to an error in the handling of a partitioned table

During the data import phase of upgrades to vCenter Server 8.0, the operation might fail due to a rare issue related to the handling of a partitioned table. In the `vcddb_import.log` file, you see an error such as `relation already exists` for the `vsan_historical_health` table in the vCenter Server database.

This issue is resolved in this release. The fix enhances the handling of partitioned tables such as `vsan_historical_health` during vCenter Server database export and import operations.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

This issue is resolved in this release.

vSAN Issues

RDMA not allowed for vSAN with large cluster support

RDMA networking for vSAN cannot be configured if large cluster support is enabled.

This issue is resolved in this release.

Server Configuration Issues

You might see false vCenter Server alarms for a failed email delivery

If a task for sending emails from vCenter Server times out, you might see a false alarm with a message such as `cannot send email to` although the email is sent successfully.

This issue is resolved in this release. The fix increases the timeout for the sendmail task.

Known Issues

Installation and Upgrade Issues

vCenter Server update fails with error for failed load of repository manifest data

If you set up a custom repository that has authentication and self-signed certificates that vCenter Server does not trust, updates and upgrades to vCenter Server 8.0 and later might fail. In the vSphere Client, you see an error such as `Failed to load the repository manifest data for the configured upgrade after the upgrade precheck completes`.

Workaround: See VMware knowledge base article [90259](#).

Miscellaneous Issues

You cannot revert configuration changes for the vmware-rbd-watchdog service

When you use the TLS Configuration utility to revert configuration changes, the vmware-rbd-watchdog service does not restore back to the expected TLS protocols.

Workaround: Use the argument `-s <service-name>` to individually call the rbd service, for example:

```
root@sc1-10-168-175-98 [ /usr/lib/vmware-TlsReconfigurator/VcTlsReconfigurator ]# ./
reconfigureVc restore -d /tmp/20221124T100959 -s vmware-rbd-watchdog
```

Security Issues

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Miscellaneous Issues

Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs

If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration`. Note that this operation cannot be performed when a managed DPU is present. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

Workaround: Reinstall ESXi. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

In the vSphere API Explorer, VMware Datacenter CLI (DCLI) and PowerCLI, you see an API option "contentinternal" that is not functional

You see an API option **contentinternal** in the metadata of either the vSphere API Explorer, DCLI and PowerCLI. For example, when you open `https://<your vCenter IP>/ui/app/devcenter/api-explorer`, you see the option in the **select API** drop-down menu. This option is not functional.

Workaround: Ignore the **contentinternal** API option and do not use it.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as *An error occurred while saving the snapshot: A digest operation has failed* in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

vSphere Lifecycle Manager Issues

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

```
A general system error occurred: After host ... remediation completed, compliance check
reported host as 'non-compliant'. The image on the host does not match the image set for
the cluster. Retry the cluster remediation operation.
```

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

VMware vCenter Server 8.0a Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Patches Contained in This Release](#)
- [Resolved Issues](#)
- [Known Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

vCenter Server 8.0a | 15 DEC 2022 |ISO Build 20920323
Check for additions and updates to these release notes.

What's New

- vCenter Server 8.0a delivers new features and fixes for VMware vSphere with Tanzu and vCenter Server. For VMware vSphere with Tanzu updates, see [VMware vSphere with Tanzu Release Notes](#). For vCenter Server fixes and known issues, see the sections [Resolved Issues](#) and [Known Issues](#).
- This release resolves CVE-2021-22048. For more information on this vulnerability and its impact on VMware products, see [VMSA-2021-0025](#).

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the [VMware vSphere 8.0 Release Notes](#).

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article [67077](#).

Patches Contained in This Release

This release of vCenter Server 8.0a delivers the following patch:

- [Patch for VMware vCenter Server 8.0a](#)

Patch for VMware vCenter Server 8.0a

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.0.10100-20920323-patch-FP.iso
Build	20920323
Download Size	6402.2 MB
sha256checksum	80291814a109b1e1c538c307e19268d3e13a49f0566a44ea1249552b70f67bc4

Download and Installation

Log in to the [Broadcom Support Portal](#) to download this [patch](#).

For more information on using the vCenter Server shells, see VMware knowledge base article [2100508](#).

For more information on patching vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#).

For more information on staging patches, see [Upgrading the vCenter Server Appliance](#).

Resolved Issues

Security Issues

After upgrade to vCenter Server 8.0, synchronization of password-protected subscribed content libraries might fail

After upgrade to vCenter Server 8.0, synchronization of password-protected subscribed content libraries might fail due to missing security key details. In the logs, you see an error such as: `javax.crypto.BadPaddingException: Error finalising cipher data: pad block corrupted`.

This issue is resolved in this release.

Networking Issues

LoadBalancers and Guest Clusters are not created when two SE Groups exist on NSX-ALB

If a second SE Group is added to NSX-ALB with or without SEs or virtual services assigned to it, the creation of new supervisor or guest clusters fails, and existing supervisor clusters cannot be upgraded. The virtual service creation on NSX-ALB controller fails with the following error:

```
get() returned more than one ServiceEngineGroup - it returned 2
```

As a result, new load balancers are unusable and you cannot create new workload clusters successfully.

For more information, see the VMware Knowledge Base article [90386](#).

This issue is resolved in this release.

Known Issues

Installation and Upgrade Issues

vCenter Server update fails with error for failed load of repository manifest data

If you set up a custom repository that has authentication and self-signed certificates that vCenter Server does not trust, updates and upgrades to vCenter Server 8.0 and later might fail. In the vSphere Client, you see an error such as `Failed to load the repository manifest data for the configured upgrade after the upgrade precheck completes`.

Workaround: See VMware knowledge base article [90259](#).

Miscellaneous Issues

You cannot revert configuration changes for the vmware-rbd-watchdog service

When you use the TLS Configuration utility to revert configuration changes, the vmware-rbd-watchdog service does not restore back to the expected TLS protocols.

Workaround: Use the argument `-s <service-name>` to individually call the rbd service, for example:

```
root@sc1-10-168-175-98 [ /usr/lib/vmware-TlsReconfigurator/VcTlsReconfigurator ]# ./reconfigureVc restore -d /tmp/20221124T100959 -s vmware-rbd-watchdog
```

Known Issues from Previous Releases

Installation, Upgrade, and Migration Issues

If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error

Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions

Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

VMNICs might be down after an upgrade to ESXi 8.0

If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

Workaround: Use either of these 2 options:

1. Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings > <specific broadcom NIC> > Device Configuration Menu > Media Auto Detect**. Reboot the host.
2. Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail

In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

After upgrade to ESXi 8.0, you might lose some `nmlx5_core` driver module settings due to obsolete parameters

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

Second stage of vCenter Server restore procedure freezes at 90%

When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

In the vSphere API Explorer, VMware Datacenter CLI (DCLI) and PowerCLI, you see an API option "contentinternal" that is not functional

You see an API option `contentinternal` in the metadata of either the vSphere API Explorer, DCLI and PowerCLI. For example, when you open `https://<your vCenter IP>/ui/app/devcenter/api-explorer`, you see the option in the **select API** drop-down menu. This option is not functional.

Workaround: Ignore the `contentinternal` API option and do not use it.

If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen

If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs

Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU

In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

Workaround: None.

If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch vSwitchBMC with uplink vusb0 gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch vSwitchBMC and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on

NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode

Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example `vmnic6` and `vmnic7`) and one management controller (for example `vmnic8`):

```
:~] esxcfg-nics -l
```

```
vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
```

```
:~] esxcfg-nics -ls
```

```
vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando
Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
```

The high-speed ethernet controllers `vmnic6` and `vmnic7` register first and operate with RSS set to 16 receive queues.

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -
n vmnic6...Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

However, in rare cases, if the management controller `vmnic8` registers first with the vSphere Distributed Switch, the high-speed ethernet controllers `vmnic6` or `vmnic7` uplink might end up operating with RSS set to 1 receive queue.:

```
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n
vmnic6...Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1,
vlan_rx_strip=1, geneve_offload=1 }
```

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en` driver on ESXi by using the following commands:

```
:~] esxcfg-module -u ionic_en
:~] esxcfg-module ionic_en
:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/deviceInternal bind.
```

In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage

Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

In the Pre-Update Check Results screen, you see an error such as:

Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.

Workaround: Refer to the [VMware Compatibility Guide](#) and [VMware Product Interoperability Matrix](#) or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog [Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution](#) and VMware knowledge base article [87880](#).

You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled

Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled.

Workaround: See VMware knowledge base article [89638](#).

If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1

If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article [89639](#).

If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected

Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU

If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

Workaround: None.

You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later

When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to 1. The `noPhyStateSet` parameter defaults to 0 and is not required in most environments, except they face the issue.

VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error

An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

Workaround: Wait for 10 min and retry the NSX install or upgrade.

If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses

If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0

Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts

When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

Workaround: None.

Storage Issues

VASA API version does not automatically refresh after upgrade to vCenter Server 8.0

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager

Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed

A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed in the backtrace`. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

The Utilization view of resource pools and clusters might not automatically refresh when you change the object

When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

Workaround: Click the refresh icon.

If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server

In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

In the vSphere Client, you do not see banner notifications for historical data imports

Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see [Monitor and Manage Historical Data Migration](#).

You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment

If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster') . The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.`

Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

Virtual Machine Management Issues

When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected

When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next > Finish**.

Security Features Issues

If you use an RSA key size smaller than 2048 bits, RSA signature generation fails

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.

vSphere Lifecycle Manager Issues

If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation

With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

Workaround: None.

If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails

If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_XXX. Extraction of image from host XXX.eng.vmware.com failed.`

Workaround: Use a different host from the cluster to extract an image.

You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs

When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

`A general system error occurred: After host ... remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

VMware Host Client might display incorrect descriptions for severity event states

When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Guest OS Issues

The guest operating system of a VM might become unresponsive due to lost communication over the Virtual Machine Communication Interface (VMCI)

In very specific circumstances, when a vSphere vMotion operation on a virtual machine runs in parallel with an operation that sends VMCI datagrams, services that use VMCI datagrams might see unexpected communication or loss of communication. Under the same conditions, the issue can also happen when restoring a memory snapshot, resuming a suspended VM or using CPU Hot Add. As a result, the guest operating system that depends on services communicating over VMCI might become unresponsive. The issue might also affect services that use vSockets over VMCI. This problem does not impact VMware Tools. The issue is specific for VMs on hardware version 20 with a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE`. For more information, see VMware knowledge base article [89683](#).

Linux guest operating system cannot complete booting if Direct Memory Access (DMA) remapping is enabled

If the advanced processor setting `Enable IOMMU in this virtual machine` is enabled on a virtual machine, and the guest operating system has enabled DMA remapping, the Linux guest operating system might fail to complete the booting process. This issue affects VMs with hardware version 20 and a Linux distribution that has specific patches introduced in Linux kernel 5.18 for a VMCI feature, including but not limited to up-to-date versions of RHEL 8.7, Ubuntu 22.04 and 22.10, and SLES15 SP3, and SP4.

Workaround: Set the advanced option `vmci.dmaDatagramSupport` to `FALSE` or disable the `Enable IOMMU in this virtual machine` option. For more information, see VMware knowledge base article [89683](#).

VMware Host Client Release Notes

VMware Host Client 2.18.0 Release Notes

This document contains the following sections

- [Introduction](#)
- [What's in the Release Notes](#)
- [Product Support Notice](#)
- [Resolved Issues](#)
- [Known Issues from Previous Releases](#)

Introduction

VMware vSphere 8.0 Update 3 | 25 JUN 2024
 VMware Host Client 2.18.0 | 25 JUN 2024 | Build 23595382
 Check for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

VMware Host Client Overview

The VMware Host Client is an HTML5-based client that is used to connect to and manage single ESXi hosts. It can be used to perform administrative tasks to manage host resources such as virtual machines, networking, and storage. The VMware Host Client can also be helpful for troubleshooting individual virtual machines or hosts if vCenter Server and the vSphere Client are unavailable.

For more information, see the [vSphere 8.0 Update 2 Documentation](#).

Previous Releases

- VMware Host Client 2.14.0
- VMware Host Client 2.12.0
- VMware Host Client 2.5.0
- VMware Host Client 1.37.0
- VMware Host Client 1.34.8
- VMware Host Client 1.34.4
- VMware Host Client 1.34.0
- VMware Host Client 1.33.3
- VMware Host Client 1.30.0
- VMware Host Client 1.25.0
- VMware Host Client 1.21
- VMware Host Client 1.8
- VMware Host Client 1.4
- VMware Host Client 1.2
- VMware Host Client 1.0

Product Support Notice

Beginning with the next major release, we will be reducing the number of supported localization languages. The three supported languages will be:

- Japanese
- Spanish
- French

The following languages will no longer be supported:

- Italian, German, Brazilian, Portuguese, Traditional Chinese, Korean, Simplified Chinese

Impact:

- Users who have been using the deprecated languages will no longer receive updates or support in these languages.
- All user interfaces, help documentation, and customer support will be available only in English or in the three supported languages mentioned above.

Resolved Issues

You can select the guest operating system type that has supportedForCreate=0

Hardware version 18 now includes the following guest operating systems that were not previously available in the list of options in the **New Virtual Machine** wizard.

- Amazon Linux 3 (64-bit)

- Apple macOS 10.17
- Asianux 9 (64-bit)
- CentOS 9 (64-bit)
- Microsoft Windows Server 2021
- Oracle Linux 9 (64-bit)
- Red Hat Enterprise Linux 9 (64-bit)
- SUSE Linux Enterprise 16 (64-bit)

This issue is resolved in this release.

Known Issues from Previous Releases

The names of future Windows OS releases appear with placeholders or code names

Because the official names of the future Windows OS releases Windows 12 and Windows 2025 are still unknown when you create a virtual machine, you might see placeholders or code names on the **Select a name and guest OS** page of the **New virtual machine** wizard.

Workaround: None

The vendor's name of some local storage devices displays incorrect or incomplete

In the VMware Host Client inventory, when you navigate to **Storage > Devices**, in the **Vendor** column of the table, you might see part of the device model name instead of the vendor's name.

Workaround: None

You cannot customize the firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize the firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts.

Perform the following steps.

1. Navigate to **Networking > Firewall rules**, and select **DHCP client**.
2. Click **Edit settings**.
3. Select **Only allow connections from the following networks** and enter an IP address.

When you click **OK**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service.`

This is expected behavior.

Workaround: None

You might not change the Load balancing setting of NIC teaming

The Load balancing setting of NIC teaming might not change when you select the **Inherit from vSwitch** option.

Perform the following steps to reproduce the issue:

1. Open Google Chrome and log in to the VMware Host Client.
2. Click **Networking** in the VMware Host Client inventory and click **Management Network**.
3. Click **Edit settings**.
Result: The **Edit port group - Management Network** window appears.
4. Expand **NIC teaming**.
 - a. From the **Load balancing** drop-down menu, select **Inherit from vSwitch** and click **Save**.

5. Complete steps 1-4.
6. The **Load balancing** option remains unchanged.

Workaround:

Override the **Policy** with the same settings as vSwitch, without using **Inherit from vSwitch**.

Do not edit **configstore**.

When you create a VM through the vSphere Client and the VMware Host Client, the maximum number of cores per socket might be different for the same guest OS

When you create a virtual machine for the same guest operating system through the vSphere Client and the VMware Host Client, the CPU Cores per Socket drop-down menu might display different numbers.

Use vSphere Client because the VMware Host Client is in maintenance mode.

The **Current reading** parameter of the Fan Redundancy health sensor displays **Redundancy lost**

Although the system fan health sensor is in green status, the **Current reading** parameter displays **Redundancy lost**. When you run the ESXCLI command `esxcli hardware ipmi sdr list`, the **Computed reading** returns a **Fully Redundant** value. As both **Current reading** and **Computed reading** parameters are the same, the **Redundancy lost** value is false.

Perform the following steps to reproduce this issue:

1. Install an ESXi 8.0 guest operating system.
2. Log in to the VMware Host Client.
3. Click **Monitor** in the VMware Host Client inventory.
4. On the **Hardware tab**, click **System sensors**.
5. Search for Fan sensors.

Result: **Current reading** column for **System Board 1 Fan Redundancy 0** displays **Redundancy lost**.

Workaround: None

VMware Host Client 2.14.0 Release Notes

This document contains the following sections

- [Introduction](#)
- [What's in the Release Notes](#)
- [Known Issues](#)

Introduction

VMware vSphere 8.0 Update 2 | 21 SEP 2023
 VMware Host Client 2.14.0 | 21 SEP 2023 | Build 21993069
 Check for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

VMware Host Client Overview

The VMware Host Client is an HTML5-based client that is used to connect to and manage single ESXi hosts. It can be used to perform administrative tasks to manage host resources such as virtual machines, networking, and storage. The

VMware Host Client can also be helpful for troubleshooting individual virtual machines or hosts if vCenter Server and the vSphere Client are unavailable.

For more information, see the [vSphere 8.0 Update 2 Documentation](#).

Previous Releases

- VMware Host Client 2.12.0
- VMware Host Client 2.5.0
- VMware Host Client 1.37.0
- VMware Host Client 1.34.8
- VMware Host Client 1.34.4
- VMware Host Client 1.34.0
- VMware Host Client 1.33.3
- VMware Host Client 1.30.0
- VMware Host Client 1.25.0
- VMware Host Client 1.21
- VMware Host Client 1.8
- VMware Host Client 1.4
- VMware Host Client 1.2
- VMware Host Client 1.0

Known Issues

The names of future Windows OS releases appear with placeholders or code names

Because the official names of the future Windows OS releases Windows 12 and Windows 2025 are still unknown when you create a virtual machine, you might see placeholders or code names on the **Select a name and guest OS** page of the **New virtual machine** wizard.

Workaround: None

The vendor's name of some local storage devices displays incorrect or incomplete

In the VMware Host Client inventory, when you navigate to **Storage > Devices**, in the **Vendor** column of the table, you might see part of the device model name instead of the vendor's name.

Workaround: None

You cannot customize the firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts

Starting with vSphere 8.0 Update 2, you cannot customize the firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts.

Perform the following steps.

1. Navigate to **Networking > Firewall rules**, and select **DHCP client**.
2. Click **Edit settings**.
3. Select **Only allow connections from the following networks** and enter an IP address.

When you click **OK**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service.`

This is expected behavior.

Workaround: None

You might not change the Load balancing setting of NIC teaming

The Load balancing setting of NIC teaming might not change when you select the **Inherit from vSwitch** option.

Perform the following steps to reproduce the issue:

1. Open Google Chrome and log in to the VMware Host Client.
2. Click **Networking** in the VMware Host Client inventory and click **Management Network**.
3. Click **Edit settings**.
Result: The **Edit port group - Management Network** window appears.
4. Expand **NIC teaming**.
 - a. From the **Load balancing** drop-down menu, select **Inherit from vSwitch** and click **Save**.
5. Complete steps 1-4.
6. The **Load balancing** option remains unchanged.

Workaround:

Override the **Policy** with the same settings as vSwitch, without using **Inherit from vSwitch**.

Do not edit **configstore**.

When you create a VM through the vSphere Client and the VMware Host Client, the maximum number of cores per socket might be different for the same guest OS

When you create a virtual machine for the same guest operating system through the vSphere Client and the VMware Host Client, the CPU Cores per Socket drop-down menu might display different numbers.

Use vSphere Client because the VMware Host Client is in maintenance mode.

The **Current reading** parameter of the Fan Redundancy health sensor displays **Redundancy lost**

Although the system fan health sensor is in green status, the **Current reading** parameter displays **Redundancy lost**. When you run the ESXCLI command `esxcli hardware ipmi sdr list`, the **Computed reading** returns a **Fully Redundant** value. As both **Current reading** and **Computed reading** parameters are the same, the **Redundancy lost** value is false.

Perform the following steps to reproduce this issue:

1. Install an ESXi 8.0 guest operating system.
2. Log in to the VMware Host Client.
3. Click **Monitor** in the VMware Host Client inventory.
4. On the **Hardware tab**, click **System sensors**.
5. Search for Fan sensors.
Result: **Current reading** column for **System Board 1 Fan Redundancy 0** displays **Redundancy lost**.

Workaround: None

VMware Host Client 2.12.0 Release Notes

This document contains the following sections

- [Introduction](#)
- [What's in the Release Notes](#)
- [Known Issues](#)

Introduction

VMware vSphere 8.0 Update 1 | 18 APR 2023
VMware Host Client 2.12.0 | 18 APR 2023 | Build 21495799
Check for additions and updates to these release notes.

What's in the Release Notes

The release notes cover the following topics:

VMware Host Client Overview

The VMware Host Client is an HTML5-based client that is used to connect to and manage single ESXi hosts. It can be used to perform administrative tasks to manage host resources such as virtual machines, networking, and storage. The VMware Host Client can also be helpful for troubleshooting individual virtual machines or hosts if vCenter Server and the vSphere Client are unavailable.

For more information, see the [vSphere 8.0 Update 1 Documentation](#).

Previous Releases

- VMware Host Client 2.5.0
- VMware Host Client 1.37.0
- VMware Host Client 1.34.8
- VMware Host Client 1.34.4
- VMware Host Client 1.34.0
- VMware Host Client 1.33.3
- VMware Host Client 1.30.0
- VMware Host Client 1.25.0
- VMware Host Client 1.21
- VMware Host Client 1.8
- VMware Host Client 1.4
- VMware Host Client 1.2
- VMware Host Client 1.0

Known Issues

When you create a VM through the vSphere Client and the VMware Host Client, the maximum number of cores per socket might be different for the same guest OS

When you create a virtual machine for the same guest operating system through the vSphere Client and the VMware Host Client, the CPU Cores per Socket drop-down menu might display different numbers.

Workaround: Use vSphere Client because the VMware Host Client is in maintenance mode.

When you create a VM through the vSphere Client and the VMware Host Client, the maximum number of cores per socket might be different for the same guest OS

When you create a virtual machine for the same guest operating system through the vSphere Client and the VMware Host Client, the CPU Cores per Socket drop-down menu might display different numbers.

Use vSphere Client because the VMware Host Client is in maintenance mode.

You might not change the Load balancing setting of NIC teaming

The Load balancing setting of NIC teaming might not change when you select the **Inherit from vSwitch** option.

Perform the following steps to reproduce the issue:

1. Open Google Chrome and log in to the VMware Host Client.
2. Click **Networking** in the VMware Host Client inventory and click **Management Network**.
3. Click **Edit settings**.
Result: The **Edit port group - Management Network** window appears.
4. Expand **NIC teaming**.
 - a. From the **Load balancing** drop-down menu, select **Inherit from vSwitch** and click **Save**.
5. Complete steps 1-4.
6. The **Load balancing** option remains unchanged.

Workaround:

Override the **Policy** with the same settings as vSwitch, without using **Inherit from vSwitch**.

Do not edit **configstore**.

The **Current reading** parameter of the Fan Redundancy health sensor displays **Redundancy lost**

Although the system fan health sensor is in green status, the **Current reading** parameter displays **Redundancy lost**. When you run the ESXCLI command `esxcli hardware ipmi sdr list`, the **Computed reading** returns a **Fully Redundant** value. As both **Current reading** and **Computed reading** parameters are the same, the **Redundancy lost** value is false.

Perform the following steps to reproduce this issue:

1. Install an ESXi 8.0 guest operating system.
2. Log in to the VMware Host Client.
3. Click **Monitor** in the VMware Host Client inventory.
4. On the **Hardware tab**, click **System sensors**.
5. Search for Fan sensors.
Result: **Current reading** column for **System Board 1 Fan Redundancy 0** displays **Redundancy lost**.

Workaround: None

ESXi Installation and Setup

VMware ESXi Installation and Setup describes how to install and configure VMware ESXi™.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we have updated this guide to remove instances of non-inclusive language.

Intended Audience

VMware ESXi Installation and Setup is intended for experienced administrators who want to install and configure ESXi.

This information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations. The information about using the Image Builder and VMware vSphere® Auto Deploy™ is written for administrators who have experience with Microsoft PowerShell and VMware vSphere® PowerCLI™.

About VMware ESXi Installation and Setup

VMware ESXi Installation and Setup describes how to install and configure VMware ESXi™.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we have updated this guide to remove instances of non-inclusive language.

Intended Audience

VMware ESXi Installation and Setup is intended for experienced administrators who want to install and configure ESXi.

This information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations. The information about using the Image Builder and VMware vSphere® Auto Deploy™ is written for administrators who have experience with Microsoft PowerShell and VMware vSphere® PowerCLI™.

- [About VMware ESXi Installation and Setup](#)
- [Introduction to vSphere Installation and Setup](#)
- [What is VMware vSphere Distributed Services Engine](#)
- [Installing and Setting Up ESXi](#)
- [Troubleshooting ESXi Booting](#)
- [Decommission an ESXi Host](#)

Introduction to vSphere Installation and Setup


vSphere 8.0 provides various options for installation and setup that define a corresponding sequence of tasks.

The two core components of vSphere are ESXi and vCenter Server. ESXi is the virtualization platform on which you can create and run virtual machines and virtual appliances. vCenter Server is a service that acts as a central administrator for ESXi hosts connected in a network. vCenter Server lets you pool and manage the resources of multiple hosts.

You deploy the vCenter Server appliance, a preconfigured virtual machine optimized for running vCenter Server and the vCenter Server components. You can deploy the vCenter Server appliance on ESXi hosts or on vCenter Server instances.

How to Install and Set Up vSphere

vSphere is a sophisticated product with multiple components to install and set up. To ensure a successful vSphere deployment, understand the sequence of tasks required.

Figure 10: vSphere Installation and Setup Workflow

Installing vSphere includes the following tasks:

1. Read the vSphere release notes.
2. Install ESXi.
3. Install vCenter Server.

About ESXi Evaluation and Licensed Modes

You can use evaluation mode to explore a set of features equal to the vSphere Enterprise Plus license.

You can use evaluation mode to explore the entire set of features for ESXi hosts. The evaluation mode provides the set of features equal to a vSphere Enterprise Plus license. Before the evaluation mode expires, you must assign to your hosts a license that supports all the features in use. For example, in evaluation mode, you can use vSphere vMotion technology,

the vSphere HA feature, the vSphere DRS feature, and other features. If you want to continue using these features, you must assign a license that supports them.

The installable version of ESXi hosts is always installed in evaluation mode. ESXi Embedded is preinstalled on an internal storage device by your hardware vendor. It might be in evaluation mode or prelicensed.

The evaluation period is 60 days and begins when you turn on the ESXi host. At any time during the 60-day evaluation period, you can convert from licensed mode to evaluation mode. The time available in the evaluation period is decreased by the time already used.

For example, suppose that you use an ESXi host in evaluation mode for 20 days and then assign a vSphere Standard Edition license key to the host. If you set the host back in evaluation mode, you can explore the entire set of features for the host for the remaining evaluation period of 40 days.

For ESXi hosts, license or evaluation period expiry leads to disconnection from vCenter Server. All powered on virtual machines continue to work, but you cannot power on virtual machines after they are powered off. You cannot change the current configuration of the features that are in use. You cannot use the features that remained unused before the license expiration.

For information about managing licensing for ESXi hosts, see the *vCenter Server and Host Management* documentation.


What is VMware vSphere Distributed Services Engine®

vSphere Distributed Services Engine® is a core vSphere capability that enables customers to use DPUs with vSphere and VMware Cloud Foundation.

vSphere 8.0 enables breakthrough workload performance to meet ever-increasing throughput and latency needs. With vSphere Distributed Services Engine, infrastructure services are distributed across the different compute resources available on the ESXi host, with networking functions offloaded to the DPU. Such a capability works well for modern applications, which are developed using a microservices architecture approach that seeks to break down the application into multiple independent but cooperating services. This increased complexity places new demand for the CPU. For example, processing storage requests or shuttling network traffic for these microservices leaves fewer CPU cycles for the actual workload. In this context, purpose-built accelerators such as DPUs can take on the new compute burden and help you improve the performance and efficiency of infrastructure.

With vSphere Distributed Services Engine, DPUs can accelerate the performance of your network and increase data throughput, while placing no operational burden of managing the lifecycle of DPUs, as the existing Day-0, Day-1, and Day-2 vSphere experience does not change. vSphere Distributed Services Engine is supported by DPUs from NVIDIA and AMD, and server designs from Dell, HPE, Lenovo, and Fujitsu. vSphere Distributed Services Engine is available on servers with pre-installed DPUs.

Starting with vSphere 8.0, you can offload functionality that runs on the core CPU onto the DPU to significantly improve the network and security performance. As illustrated in the Evolving vSphere Architecture diagram, DPUs can also handle additional capabilities such as storage offload and bare metal management, but these additional capabilities are currently not supported.

Figure 11: Evolving vSphere Architecture.

vSphere Distributed Services Engine offloads and accelerates infrastructure functions on the DPU by introducing a VMware vSphere Distributed Switch on the DPU and VMware NSX Networking and Observability, which allows you to proactively monitor, identify, and mitigate network infrastructure bottlenecks without complex network taps. The DPU becomes a new control point to scale infrastructure functions and enables security controls that are agentless and decoupled from the workload domain.

With vSphere Distributed Services Engine, you can:

- Install and update ESXi images simultaneously on the x86 server and the attached supported DPU to reduce operational overhead of DPU lifecycle management with integrated vSphere workflows. For more information, see [Using vSphere Lifecycle Manager with VMware vSphere Distributed Services Engine](#).
- Set alarms for DPU hardware alerts and monitor performance metrics on core, memory, and network throughput from the familiar vCenter interfaces, without the need of new tools. For more information, see [CPU \(DPU\)](#) and [Memory \(DPU\)](#).
- Accelerate vSphere Distributed Switch on the DPU to improve network performance and utilize available CPU cycles to achieve higher workload consolidation per ESXi host. For more information, see [What is Network Offloads Compatibility](#) and [Create a vSphere Distributed Switch](#).
- Get vSphere DRS and vSphere vMotion support for VMs running on hosts with DPUs attached to get the benefits of passthrough without sacrificing on VM portability.
- Improve the security of infrastructure with zero-trust security. For more information, see [vSphere Distributed Services Engine Security Best Practices](#).

vSphere Distributed Services Engine does not require a separate ESXi license. An internal network that is isolated from other networks, connects the DPUs with ESXi hosts. ESXi 8.0 server builds are unified images, which contain both x86 and DPU content. In your vSphere system, you see DPUs as new objects during installation and upgrade, and in networking, storage, and host profile workflows.

High Availability with VMware vSphere Distributed Services Engine

With ESXi 8.0 Update 3, you can opt for a VMware vSphere Distributed Services Engine installation with 2 data processing units (DPUs) to achieve high availability.

In vSphere systems with a single DPU, the device might become the single point of failure for workloads offloaded to the DPU, such as networking functions, and impact data and productivity. With ESXi 8.0 Update 3, vSphere Distributed Services Engine is also available on servers with 2 pre-installed DPUs, which provides hardware redundancy and resiliency.

You can utilize the two DPUs in Active/Standby mode to provide high availability. Such configuration provides redundancy in the event one of the DPUs fails. In the high availability configuration, both DPUs are assigned to the same NSX-backed vSphere Distributed Switch. For example, DPU-1 is attached to vmnic0 and vmnic1 of the vSphere Distributed Switch and DPU-2 is attached to vmnic2 and vmnic3 of the same vSphere Distributed Switch.

You can also utilize the two DPUs as independent devices to increase offload capacity per ESXi host. Each DPU is attached to a separate vSphere Distributed Switch and you have no failover between DPUs in such configuration.

Dual-DPU systems can use NVIDIA or Pensando devices. In ESXi 8.0 Update 3, dual-DPU systems are supported by Lenovo server designs. The DPU devices on a dual DPU server must be identical in all aspects: same vendor, same hardware version and same firmware. For a list of current vendors and server designs for VMware vSphere Distributed Services Engine, see the [VMware Compatibility Guide](#).

Installation of VMware vSphere Distributed Services Engine with 2 DPUs

vSphere Distributed Services Engine does not require a separate ESXi license. ESXi 8.0 Update 3 server builds are unified images, which contain both x86 and DPU content, and you cannot install x86 and DPU content separately. The installation procedure on both DPUs, either interactive or scripted, also happens in parallel and you see minimal performance loss as compared to a single-DPU system.

With vSphere 8.0 Update 3, you can get a pre-installed server configuration with 2 DPUs from Dell or Lenovo, or add a second DPU to a single DPU system on the supported dual DPU servers from Dell or Lenovo.

Note: In any case, you need to run a complete fresh ESXi 8.0 Update 3 installation on your system, not only on the newly added DPUs.

For more information on the installation, see [Install ESXi Interactively](#) and [Installation and Upgrade Scripts Used for ESXi Installation](#).

Error Handling, Failover, and Rollback for VMware vSphere Distributed Services Engine

Before installing VMware vSphere Distributed Services Engine, see the error handling, failover, and rollback options.

Error Handling

An installation failure of either x86 and DPU content on an ESXi host marks the entire installation procedure as failed.

While the expectation is that the software state of DPUs remains identical at all times, in the unlikely case of an error during a lifecycle operation, such as installation or upgrade of a Component, the operation might pass on one DPU but fail on the other. Since each lifecycle operation occurs within the boundaries of each DPU, errors do not affect the state of the other DPU, but the overall result of the installation is still marked as a failure.

During interactive install, in vSphere Lifecycle Manager workflows, and when you use ESXCLI, you receive information about the DPU on which the operation failed.

After a successful installation, in case of DPU errors, the recommended action is to restart the affected ESXi host. If the DPU is still accessible from the host, the general log bundle collection is sufficient for troubleshooting. If the DPU is not accessible from the host, logging in to the DPU from a BMC, iLO, or iDRAC interface can provide troubleshooting logs.

Failover

Failover support in vSphere 8.0 Update 3 is limited to one of the DPUs becoming non-functional due to software errors within the DPU or a physical disconnect of one of the DPUs, such as cable disconnect. Failover due to Peripheral Component Interconnect (PCI) level errors is not supported.

Rollback

Rollback is a best effort mechanism to restore the system to a previous working state in case of a failure before the jumpstart phase of the ESXi boot. Rollback on both x86 servers and the attached supported DPUs is automatic in case of an error during booting. You can also opt for a manual rollback by pressing `Shift+R` before the bootloader starts, to return to a previous good state.

Any failure after the jumpstart phase starts does not result in a rollback.

Table 1: Rollback scenarios for VMware vSphere Distributed Services Engine installation

Scenario	Number of reboots required
Both DPUs boot correctly. ESXi does not boot correctly.	2
Both DPUs do not boot correctly. ESXi boots correctly.	1
One of the DPUs boots with an earlier version than the other DPU and ESXi.	2
One of the DPUs boots with an earlier version than the other DPU and ESXi does not boot correctly.	2

Installing and Setting Up ESXi

You can install and set up ESXi on your physical hardware so that it acts as a platform for virtual machines.

In the chapter, you can read about:

- [ESXi Requirements](#)
- [Preparing for Installing ESXi](#)
- [Customizing Installations with vSphere ESXi Image Builder](#)
- [Installing ESX](#)
- [Setting Up ESX](#)
- [After You Install and Set Up ESXi](#)
- [Troubleshooting ESXi Booting](#)
- [Decommission an ESXi Host](#)

ESXi Requirements

To install or upgrade ESXi, your system must meet specific hardware and software requirements.

- [ESXi System Storage Overview](#)

- [ESXi Hardware Requirements](#)
- [Using Remote Management Applications](#)
- [Recommendations for Enhanced ESXi Performance](#)
- [Incoming and Outgoing Firewall Ports for ESXi Hosts](#)
- [Required Free Space for System Logging](#)
- [VMware Host Client System Requirements](#)
- [ESXi Passwords and Account Lockout](#)

ESXi System Storage Overview

ESXi8.0 has a system storage layout that allows flexible management of partitions and third-party components, while facilitating debugging.

ESXi System Storage

The ESXi8.0 system storage layout consists of four partitions:

Table 2: ESXi system storage partitions:


Partition	Use	Type
System Boot	Stores boot loader and EFI modules.	FAT16
Boot-bank 0	System space to store ESXi boot modules.	FAT16
Boot-bank 1	System space to store ESXi boot modules.	FAT16
ESX-OSData	Acts as the unified location to store additional modules. Not used for booting and virtual machines. Consolidates the legacy <code>/scratch</code> partition, locker partition for VMware Tools, and core dump destinations. CAUTION: In case the installation media is a USB or an SD card device, best practice is to create ESX-OSData partitions on persistent storage device that is not shared between ESXi hosts.	VMFS-L

The ESX-OSData volume is divided into two high-level categories of data, persistent and non-persistent data. Persistent data contains of data written infrequently, for example, VMware Tools ISOs, configurations, and core dumps.

Non-persistent data contains of frequently written data, for example, logs, VMFS global traces, vSAN Entry Persistence Daemon (EPD) data, vSAN traces, and real-time databases.

Figure 12: Consolidated system storage in ESXi 8.0

ESXi 6.x
system storage layout



ESXi System Storage Sizes

Partition sizes, except for the system boot partition, can vary depending on the size of the boot media used. If the boot media is a high-endurance one with capacity larger than 142 GB, a VMFS datastore is created automatically to store virtual machine data.

You can review the boot media capacity and the automatic sizing as configured by the ESXi installer by using the vSphere Client and navigating to the **Partition Details** view. Alternatively, you can use ESXCLI, for example the `esxcli storage filesystem list` command.

Table 3: ESXi System Storage Sizes, Depending on the Used Boot Media and Its Capacity.

Boot Media Size	8-10 GB	10-32 GB	32-128 GB	>128 GB
System Boot	100 MB	100 MB	100 MB	100 MB
Boot-bank 0	500 MB	1 GB	4 GB	4 GB
Boot-bank 1	500 MB	1 GB	4 GB	4 GB
ESX-OSData	remaining space	remaining space	remaining space	up to 128 GB
VMFS datastore				remaining space for media size > 142 GB

You can use the ESXi installer boot option `systemMediaSize` to limit the size of system storage partitions on the boot media. If your system has a small footprint that does not require the maximum of 128 GB of system storage size, you can limit it to the minimum of 32 GB. The `systemMediaSize` parameter accepts the following values:

- min (32 GB, for single disk or embedded servers)
- small (64 GB, for servers with at least 512 GB of RAM)
- default (128 GB)
- max (consume all available space, for multi-terabyte servers)

The selected value must fit the purpose of your system. For example, a system with 1 TB of memory must use the minimum of 64 GB for system storage. To set the boot option at install time, for example `systemMediaSize=small`, refer to [Enter Boot Options to Run an Installation or Upgrade ScriptBoot Option](#). For more information, see Knowledge Base article [81166](#).

ESXi System Storage Links

The sub-systems that require access to the ESXi partitions, access these partitions by using the following symbolic links:

Table 4: ESXi system storage symbolic links.

System Storage Volume	Symbolic Link
Boot-bank 0	/bootbank
Boot-bank 1	/altbootbank
Persistent data	/productLocker /locker /var/core /usr/lib/vmware/isoimages /usr/lib/vmware/floppies
Non-persistent data	/var/run /var/log

System Storage Volume	Symbolic Link
	/var/vmware /var/tmp /scratch

Storage Behavior

When you start ESXi, the host enters an autoconfiguration phase during which system storage devices are configured with defaults.

When you reboot the ESXi host after installing the ESXi image, the host configures the system storage devices with default settings. Starting with ESXi 7.0, you can activate the option `autoPartition`, which automatically formats all available empty devices with VMFS, except for legacy SD and USB devices. The default is `autoPartition=FALSE`, which formats with VMFS only boot devices with size larger than 128 GB. For more information, see VMware knowledge base article [77009](#).

CAUTION: ESXi overwrites any disks that appear to be blank. Disks are considered to be blank if they do not have a valid partition table or partitions. If you are using software that uses such disks, in particular if you are using logical volume manager (LVM) instead of, or in addition to, conventional partitioning schemes, ESXi might cause local LVM to be reformatted. Back up your system data before you power on ESXi for the first time.

On the hard drive or USB device that the ESXi host is booting from, the disk-formatting software retains existing diagnostic partitions that the hardware vendor creates. In the remaining space, the software creates the partitions described below.

Partitions Created by ESXi on the Host Drive

For fresh installations, several new partitions are created for the system boot, boot banks, and ESX-OSData. Fresh ESXi installations use GUID Partition Tables (GPT) instead of MSDOS-based partitioning. The installer creates boot banks of varying size depending on the size of the disk. For more information on the scratch partition see [About the Scratch Partition](#).

The installer affects only the installation disk. The installer does not affect other disks of the server. When you install on a disk, the installer overwrites the entire disk. When the installer autoconfigures storage, the installer does not overwrite hardware vendor partitions.

To create the VMFS datastore, the ESXi installer requires a minimum of 128 GB of free space on the installation disk.

You might want to override this default behavior if, for example, you use shared storage devices instead of local storage. To prevent automatic disk formatting, detach the local storage devices from the host under the following circumstances:

- Before you start the host for the first time.
- Before you start the host after you reset the host to the configuration defaults.

To override the VMFS formatting if automatic disk formatting already occurred, you can remove the datastore. See the *vCenter Server and Host Management* documentation.

About the Scratch Partition

For new installations of ESXi, during the autoconfiguration phase, a scratch partition is created as part of the ESX-OSDATA partition.

Note: Partitioning for hosts that are upgraded to ESXi 7.0 and later from earlier versions differs significantly from partitioning for new installations of ESXi. The upgrade process to ESXi 7.0 and later repartitions the boot device and consolidates the original core dump, locker, and scratch partitions into the ESX-OSData volume.

The scratch partition serves to store system logs, which you need when you create a support bundle. If the scratch partition is not present, system logs are stored in a ramdisk. If no scratch partition is created, you can configure one. You can also override the default configuration.

You can create the scratch partition on a remote SAN or NFS-mounted directory.

Set the Scratch Partition from the vSphere Client

If a scratch partition is not set up, you might want to configure one, especially if the host is low on memory. When a scratch partition is not present, system logs are stored in a ramdisk.

The directory to use for the scratch partition must exist on the host.

1. From the vSphere Client, connect to the vCenter Server.
2. Select the host in the inventory.
3. Click the **Configure** tab.
4. Select **System**.
5. Select **Advanced System Settings**.

The setting **ScratchConfig.CurrentScratchLocation** shows the current location of the scratch partition.

6. In the **ScratchConfig.ConfiguredScratchLocation** text box, enter a directory path that is unique for this host. For example, `/vmfs/volumes/DatastoreUUID/DatastoreFolder`.
7. Reboot the host for the changes to take effect.

ESXi Hardware Requirements

Make sure that the host meets the minimum hardware configurations supported by ESXi8.0.

Hardware and System Resources

To install or upgrade ESXi, your hardware and system resources must meet the following requirements:

- Supported server platform. For a list of supported platforms, see the [VMware Compatibility Guide](#).
- ESXi8.0 requires a host with at least two CPU cores.
- ESXi8.0 supports a broad range of multi-core of 64-bit x86 processors. For a complete list of supported processors, see the [VMware Compatibility Guide](#).
- ESXi8.0 requires the NX/XD bit to be enabled for the CPU in the BIOS.
- ESXi8.0 requires a minimum of 8 GB of physical RAM. Provide at least 12 GB of RAM to run virtual machines in typical production environments.
- To support 64-bit virtual machines, support for hardware virtualization (Intel VT-x or AMD RVI) must be enabled on x64 CPUs.
- One or more Gigabit or faster Ethernet controllers. For a list of supported network adapter models, see the [VMware Compatibility Guide](#).
- ESXi8.0 requires a boot disk of at least 32 GB of persistent storage such as HDD, SSD, or NVMe. A boot device must not be shared between ESXi hosts.
- SCSI disk or a local, non-network, RAID LUN with unpartitioned space for the virtual machines.
- For Serial ATA (SATA), a disk connected through supported SAS controllers or supported on-board SATA controllers. SATA disks are considered remote, not local. These disks are not used as a scratch partition by default because they are seen as remote.

Note: You cannot connect a SATA CD-ROM device to a virtual machine on an ESXi host. To use the SATA CD-ROM device, you must use IDE emulation mode.

Storage Systems

For a list of supported storage systems, see the [VMware Compatibility Guide](#). Starting with ESXi 8.0, you cannot use software adapters for Fibre Channel over Ethernet (FCoE), only hardware FCoE adapters.

ESXi Booting Requirements

In vSphere 8.0, support for legacy BIOS is limited and booting ESXi hosts from the Unified Extensible Firmware Interface (UEFI) is recommended. With UEFI, you can boot systems from hard drives, CD-ROM drives, or USB media. vSphere Auto Deploy supports network booting and provisioning of ESXi hosts with UEFI. If your system has supported data processing units (DPU), you can only use UEFI to install and boot ESXi on the DPUs. For more information on VMware plans to deprecate support for legacy BIOS in server platforms, see [Deprecation of legacy BIOS support in vSphere](#).

ESXi can boot from a disk larger than 2 TB if the system firmware and the firmware on any add-in card that you are using support it. See the vendor documentation.

Storage Requirements for ESXi8.0 Installation or Upgrade

For best performance of an ESXi8.0 installation, use a persistent storage device that is a minimum of 32 GB for boot devices. Upgrading to ESXi8.0 requires a boot device that is a minimum of 8 GB. When booting from a local disk, SAN or iSCSI LUN, at least a 32 GB disk is required to allow for the creation of system storage volumes, which include a boot partition, boot banks, and a VMFS-L based ESX-OSData volume. The ESX-OSData volume takes on the role of the legacy `/scratch` partition, locker partition for VMware Tools, and core dump destination.

Note: In ESXi8.0, the ESX-OSData volume is considered a unified partition and the separate components, such as `/scratch` and VMware Tools, are consolidated into a single persistent OSDATA partition.

Other options for best performance of an ESXi8.0 installation are the following:

- A local disk of 128 GB or larger for optimal support of ESX-OSData. The disk contains the boot partition, ESX-OSData volume and a VMFS datastore.
- A device that supports the minimum of 128 terabytes written (TBW).
- A device that delivers at least 100 MB/s of sequential write speed.
- To provide resiliency in case of device failure, a RAID 1 mirrored device is recommended.

Legacy SD and USB devices are supported with the following limitations:

- SD and USB devices are supported for boot bank partitions. The use of SD and USB devices for storing ESX-OSData partitions is being deprecated and best practice is to provide a separate persistent local device with a minimum of 32 GB to store the ESX-OSData volume. The persistent local boot device can be an industrial grade M.2 flash (SLC and MLC), SAS, SATA, HDD, SSD, or a NVMe device. The optimal capacity for persistent local devices is 128 GB.
- If you do not provide persistent storage, you see an alarm such as `Secondary persistent device not found`. Please move installation to persistent storage as support for SD-Card/USB only configuration is being deprecated.
- You must use an SD flash device that is approved by the server vendor for the particular server model on which you want to install ESXi on an SD flash storage device. You can find a list of validated devices on [partnerweb.vmware.com](#).
- See [SD card/USB boot device revised guidance](#) on updated guidance for SD card or USB-based environments.
- To choose a proper SD or USB boot device, see [Boot device guidance for low endurance media \(vSphere and vSAN\)](#).

The upgrade process to ESXi8.0 from versions earlier than 7.x repartitions the boot device and consolidates the original core dump, locker, and scratch partitions into the ESX-OSData volume.

The following events occur during the repartitioning process:

- If a custom core dump destination is not configured, then the default core dump location is a file in the ESX-OSData volume.

- If the syslog service is configured to store log files on the 4 GB VFAT scratch partition, the log files in `var/run/log` are migrated to the ESX-OSData volume.
- VMware Tools are migrated from the locker partition and the partition is wiped.
- The core dump partition is wiped. The application core dump files that are stored on the scratch partition are deleted.

Note: Rollback from ESXi 8.x to a version of ESXi earlier than 7.x is not possible due to the repartitioning process of the boot device. To use a version of ESXi earlier than 7.x after upgrading to version 8.0, you must create a backup of the boot device before the upgrade, and restore the ESXi boot device from the backup. Rollback from ESXi 8.x to 7.x is possible as long as no changes to the bootbank partitions have been made and no corrupt partition is detected.

If you use USB or SD devices to perform an upgrade, best practice is to allocate an ESX-OSData region on an available persistent disk or a SAN LUN. If persistent storage or a SAN LUN are not available, ESX-OSData is automatically created on a RAM disk. VMFS can also be used for ESX-OSData partition.

After upgrade, if ESX-OSData resides on a RAM disk and a new persistent device is found on subsequent boots, and this device has the setting `autoPartition=True`, ESX-OSData is automatically created on the new persistent device. ESX-OSData does not move between persistent storage automatically, but you can manually change the ESX-OSData location on a supported storage.

To configure the size of ESXi system partitions, you can use the `systemMediaSize` option. For more information, see [Boot option to configure the size of ESXi system partitions](#).

In Auto Deploy installations, the installer attempts to allocate a scratch region on an available local disk or datastore. If no local disk or datastore is found, installation fails.

For environments that boot from a SAN or use Auto Deploy, the ESX-OSData volume for each ESXi host must be set up on a separate SAN LUN.

Using Remote Management Applications

Remote management applications allow you to install ESXi on servers that are in remote locations.

Remote management applications supported for installation include HP Integrated Lights-Out (iLO), Dell Remote Access Card (DRAC), IBM management module (MM), and Remote Supervisor Adapter II (RSA II). For support on remote management applications, contact the vendor.

You can use remote management applications to do both interactive and scripted installations of ESXi remotely.

If you use remote management applications to install ESXi, the virtual CD might encounter corruption problems with systems or networks operating at peak capacity. If a remote installation from an ISO image fails, complete the installation from the physical CD media.

Supported Remote Management Server Models and Firmware Versions

You can use remote management applications to install or upgrade ESXi, or to manage hosts remotely.

Table 5: Supported Remote Management Server Models and Minimum Firmware Versions

Remote Management Server Model	Firmware Version	Java
Dell DRAC 9	6.0.30.00	N/A
Dell DRAC 7	1.30.30 (Build 43)	1.7.0_60-b19
Dell DRAC 6	1.54 (Build 15), 1.70 (Build 21)	1.6.0_24
Dell DRAC 5	1.0, 1.45, 1.51	1.6.0_20, 1.6.0_203
Dell DRAC 4	1.75	1.6.0_23

Remote Management Server Model	Firmware Version	Java
Fujitsu iRMC S5	1.10P	1.7.0_60-b19
Fujitsu iRMC S6	1.06S	N/A
HP ILO	1.81, 1.92	1.6.0_22, 1.6.0_23
HP ILO 2	1.8, 1.81	1.6.0_20, 1.6.0_23
HP ILO 3	1.28	1.7.0_60-b19
HP ILO 4	1.13	1.7.0_60-b19
HP ILO 5	2.72	N/A
IBM RSA 2	1.03, 1.2	1.6.0_22

Recommendations for Enhanced ESXi Performance

To enhance performance, install or upgrade ESXi on a robust system with more RAM than the minimum required and with multiple physical disks.

For ESXi system requirements, see [ESXi Hardware Requirements](#).

Table 6: Recommendations for Enhanced Performance

System Element	Recommendation
RAM	ESXi hosts require more RAM than typical servers. ESXi 8.0 requires a minimum of 8 GB of physical RAM. Provide at least 12 GB of RAM to take full advantage of ESXi features and run virtual machines in typical production environments. An ESXi host must have sufficient RAM to run concurrent virtual machines. The following examples are provided to help you calculate the RAM required by the virtual machines running on the ESXi host. Operating four virtual machines with Red Hat Enterprise Linux or Windows XP requires at least 3 GB of RAM for baseline performance. This figure includes 1024 MB for the virtual machines, 256 MB minimum for each operating system as recommended by vendors. Running these four virtual machines with 512 MB RAM requires that the ESXi host have 4 GB RAM, which includes 2048 MB for the virtual machines. These calculations do not include possible memory savings from using variable overhead memory for each virtual machine. See <i>vSphere Resource Management</i>.
Dedicated Fast Ethernet adapters for virtual machines	Place the management network and virtual machine networks on different physical network cards. Dedicated Gigabit Ethernet cards for virtual machines, such as Intel PRO 1000 adapters, improve throughput to virtual machines with high network traffic.
Disk location	Place all data that your virtual machines use on physical disks allocated specifically to virtual machines. Performance is better when you do not place your virtual machines on the disk containing the ESXi boot image. Use physical disks that are large enough to hold disk images that all the virtual machines use.

System Element	Recommendation
VMFS6 partitioning	The ESXi installer creates the initial VMFS volumes on the first blank local disk found. To add disks or modify the original configuration, use the vSphere Client. This practice ensures that the starting sectors of partitions are 64K-aligned, which improves storage performance. Note: For SAS-only environments, the installer might not format the disks. For some SAS disks, it is not possible to identify whether the disks are local or remote. After the installation, you can use the vSphere Client to set up VMFS.
Processors	Faster processors improve ESXi performance. For certain workloads, larger caches improve ESXi performance.
Hardware compatibility	Use devices in your server that are supported by ESXi drivers. See the <i>Hardware Compatibility Guide</i> at http://www.vmware.com/resources/compatibility.

Incoming and Outgoing Firewall Ports for ESXi Hosts

Open and close firewall ports for each service by using either the vSphere Client or the VMware Host Client.

includes a firewall that is enabled by default. At installation time, the firewall is configured to block incoming and outgoing traffic, except traffic for services that are enabled in the host's security profile. For the list of supported ports and protocols in the firewall, see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com/>.

The VMware Ports and Protocols Tool lists port information for services that are installed by default. If you install other VIBs on your host, additional services and firewall ports might become available. The information is primarily for services that are visible in the but the VMware Ports and Protocols Tool includes some other ports as well.

Required Free Space for System Logging

See the recommended minimum size and rotation configuration for hostd, vpxa, and fdm logs.

If you used Auto Deploy to install your ESXi8.0 host, or if you set up a log directory separate from the default location in a scratch directory on the VMFS volume, you might need to change your current log size and rotation settings to ensure that enough space is available for system logging. All vSphere components use this infrastructure. The default values for log capacity in this infrastructure vary, depending on the amount of storage available and on how you have configured system logging. Hosts that are deployed with Auto Deploy store logs on a RAM disk, which means that the amount of space available for logs is small.

If your host is deployed with Auto Deploy, reconfigure your log storage in one of the following ways:

- Redirect logs over the network to a remote collector.
- Redirect logs to a NAS or NFS store.

If you redirect logs to non-default storage, such as a NAS or NFS store, you might also want to reconfigure log sizing and rotations for hosts that are installed to disk.

You do not need to reconfigure log storage for ESXi hosts that use the default configuration, which stores logs in a scratch directory on the VMFS volume. For these hosts, ESXi8.0 configures logs to best suit your installation, and provides enough space to accommodate log messages.

Table 7: Recommended Minimum Size and Rotation Configuration for hostd, vpxa, and fdm Logs

Log	Maximum Log File Size	Number of Log Files to Preserve	Minimum Disk Space Required
Management Agent (hostd)	10 MB	10	100 MB
VirtualCenter Agent (vpxa)	5 MB	10	50 MB
vSphere HA agent (Fault Domain Manager, fdm)	5 MB	10	50 MB

You can optionally install VMware vCenter Log Insight, which provides log aggregation and analytics.

VMware Host Client System Requirements

Make sure that your browser supports the VMware Host Client.

The following guest operating systems and Web browser versions are supported for the VMware Host Client.

Supported Browsers	Mac OS	Windows 32-bit and 64-bit	Linux
Google Chrome	89+	89+	75+
Mozilla Firefox	80+	80+	60+
Microsoft Edge	90+	90+	N/A
Safari	9.0+	N/A	N/A

ESXi Passwords and Account Lockout

For ESXi hosts, you must use a password with predefined requirements. You can change the required length and the character class requirement or allow pass phrases using the `Security.PasswordQualityControl` advanced system setting. You can also set the number of passwords to remember for each user using the `Security.PasswordHistory` advanced system setting.

Note: The default requirements for ESXi passwords can change from one release to the next. You can check and change the default password restrictions by using the `Security.PasswordQualityControl` advanced system setting.

ESXi Passwords

ESXi enforces password requirements for access from the Direct Console User Interface, the ESXi Shell, SSH, or the VMware Host Client.

- By default, you must include a mix of at least three from the following four character classes: lowercase letters, uppercase letters, numbers, and special characters such as underscore or dash when you create a password.
- By default, password length is at least 7 characters and less than 40.
- Passwords must not contain a dictionary word or part of a dictionary word.
- Passwords must not contain the user name or parts of the user name.

Note: An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used. A dictionary word used inside a password reduces the overall password strength.

Example ESXi Passwords

The following password candidates illustrate potential passwords if the option is set as follows.

```
retry=3 min=disabled,disabled,disabled,7,7
```

With this setting, a user is prompted up to three times (retry=3) for a new password that is not sufficiently strong or if the password was not entered correctly twice. Passwords with one or two character classes and pass phrases are not allowed, because the first three items are deactivated. Passwords from three- and four-character classes require seven characters. See the `pam_passwdqc` man page for details on other options, such as max, passphrase, and so on.

With these settings, the following passwords are allowed.

- xQaTEhb!: Contains eight characters from three character classes.
- xQaT3#A: Contains seven characters from four character classes.

The following password candidates do not meet requirements.

- Xqat3hi: Begins with an uppercase character, reducing the effective number of character classes to two. The minimum number of required character classes is three.
- xQaTEh2: Ends with a number, reducing the effective number of character classes to two. The minimum number of required character classes is three.

ESXi Pass Phrase

Instead of a password, you can also use a pass phrase. However, pass phrases are deactivated by default. You can change the default setting and other settings by using the `Security.PasswordQualityControl` advanced system setting from the vSphere Client.

For example, you can change the option to the following.

```
retry=3 min=disabled,disabled,16,7,7
```

This example allows pass phrases of at least 16 characters and at least three words.

Changing Default Password Restrictions

You can change the default restriction on passwords or pass phrases by using the `Security.PasswordQualityControl` advanced system setting for your ESXi host. See the *vCenter Server and Host Management* documentation for information on changing ESXi advanced system settings.

You can change the default, for example, to require a minimum of 15 characters and a minimum number of four words (passphrase=4), as follows:

```
retry=3 min=disabled,disabled,15,7,7 passphrase=4
```

See the man page for `pam_passwdqc` for details.

Note: Not all possible combinations of password options have been tested. Perform testing after you change the default password settings.

This example sets the password complexity requirement to require eight characters from four character classes that enforce a significant password difference, a remembered history of five passwords, and a 90 day rotation policy:

```
min=disabled,disabled,disabled,disabled,8 similar=deny
```

ESXi Account Lockout Behavior

Account locking is supported for access through SSH and through the vSphere Web Services SDK. The Direct Console Interface (DCUI) and the ESXi Shell do not support account lockout. By default, a maximum of five failed attempts is allowed before the account is locked. The account is unlocked after 15 minutes by default.

Configuring Login Behavior

You can configure the login behavior for your ESXi host with the following advanced system settings:

- `Security.AccountLockFailures`. Maximum number of failed login attempts before a user's account is locked. Zero deactivates account locking.
- `Security.AccountUnlockTime`. Number of seconds that a user is locked out.
- `Security.PasswordHistory`. Number of passwords to remember for each user. Starting in vSphere 8.0 Update 1, the default is five. Zero deactivates password history.

See the *vCenter Server and Host Management* documentation for information on setting ESXi advanced options.

Preparing for Installing ESXi

Before you install ESXi, determine the installation option that is suitable for your environment and prepare for the installation process.

- [Download the ESXi Installer](#)
- [Options for Installing ESXi](#)
- [Media Options for Booting the ESXi Installer](#)
- [Using Remote Management Applications](#)
- [Customizing Installations with vSphere ESXi Image Builder](#)
- [Required Information for ESXi Installation](#)

Download the ESXi Installer

You can obtain the ESXi installer software either from an OEM or from the Broadcom Support Portal.

Register on the Broadcom Support Portal. For more information, see [Register for an account on the Broadcom Support Portal and Communities](#).

For product download instructions, see [Download Broadcom products and software](#).

For download of offline bundle ZIP files for ESXi patches and updates, see [Downloading Broadcom PTF files and solutions](#).

For more information, see [VMware to Broadcom Support Frequently Asked Questions](#).

Required Information for ESXi Installation

Interactive installation prompts you for the required system information, but you must supply this information in the installation script.

In an interactive installation, the system prompts you for the required system information. In a scripted installation, you must supply this information in the installation script. For future use, note the values you use during the installation. These notes are useful if you must reinstall ESXi and reenter the values that you originally selected.

Table 8: Required Information for ESXi Installation

Information	Required or Optional	Default	Comments
Keyboard layout	Required	U.S. English	
VLAN ID	Optional	None	Range: 0 through 4094
IP address	Optional	DHCP	You can allow DHCP to configure the network during installation. After installation, you can change the network settings.
Subnet mask	Optional	Calculated based on the IP address	

Information	Required or Optional	Default	Comments
Gateway	Optional	Based on the configured IP address and subnet mask	
Primary DNS	Optional	Based on the configured IP address and subnet mask	
Secondary DNS	Optional	None	
Host name	Required for static IP settings	None	The vSphere Client can use either the host name or the IP address to access the ESXi host.
Install location	Required	None	Must be at least 5 GB if you install the components on a single disk.
Root password	Required	None	The root password must contain between 8 and 40 characters. For information about passwords see the <i>vSphere Security</i> documentation.

Media Options for Booting the ESXi Installer

The ESXi installer must be accessible to the system on which you are installing ESXi.

The following boot media are supported for the ESXi installer:

- Boot from a CD/DVD. See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#).
- Boot from a USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).
- Boot from a network. See [Network Booting the ESXi Installer](#).
- Boot from a remote location using a remote management application. See [Using Remote Management Applications](#).

Download and Burn the ESXi Installer ISO Image to a CD or DVD

If you do not have an ESXi installation CD/DVD, you can create one.

You can also create an installer ISO image that includes a custom installation script. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).

1. Follow the procedure [Download the ESXi Installer](#).
2. Burn the ISO image to a CD or DVD.

Format a USB Flash Drive to Boot the ESXi Installation or Upgrade

You can format a USB flash drive to boot the ESXi installation or upgrade.

- Linux machine with superuser access to it
- USB flash drive that can be detected by the Linux machine
- The ESXi ISO image, `VMware-VMvisor-Installer-version_number-build_number.x86_64.iso`, which includes the `isolinux.cfg` file

The instructions in this procedure assume that the USB flash drive is detected as `/dev/sdb`.

Note: The `ks.cfg` file that contains the installation script cannot be located on the same USB flash drive that you are using to boot the installation or upgrade. The kickstart file does not have any dependency on BIOS or UEFI boot.

1. Boot Linux, log in, and enter superuser mode by using a `su` or `sudo root` command.
2. If your USB flash drive is not detected as `/dev/sdb`, or you are not sure how your USB flash drive is detected, determine how it is detected.
 - a) Plug in your USB flash drive.
 - b) At the command line, run the command for displaying the current log messages.

```
tail -f /var/log/messages
```

You see several messages that identify the USB flash drive in a format similar to the following message.

```
Oct 25 13:25:23 ubuntu kernel: [ 712.447080] sd 3:0:0:0: [sdb] Attached SCSI
removable disk
```

In this example, `sdb` identifies the USB device. If your device is identified differently, use that identification in place of `sdb`.

3. Overwrite the entire USB drive with the ISO image. This overwrites the partition table and any previous content on the USB drive.

```
dd bs=10M if=VMware-VMvisor-Installer-version_number-build_number.x86_64.iso of=/dev/
sdb
```

4. Eject the USB drive.

```
eject /dev/sdb
```

You can use the USB flash drive to boot the ESXi installer.

Create a USB Flash Drive to Store the ESXi Installation Script or Upgrade Script

You can use a USB flash drive to store the ESXi installation script or upgrade script that is used during scripted installation or upgrade of ESXi.

- Linux machine
- ESXi installation or upgrade script, the `ks.cfg` kickstart file
- USB flash drive

When multiple USB flash drives are present on the installation machine, the installation software searches for the installation or upgrade script on all attached USB flash drives.

The instructions in this procedure assume that the USB flash drive is detected as `/dev/sdb`.

Note: Do not store the `ks` file containing the installation or upgrade script on the same USB flash drive that you are using to boot the installation or upgrade.

1. Attach the USB flash drive to a Linux machine that has access to the installation or upgrade script.
2. Create a partition table.

```
/sbin/fdisk /dev/sdb
```

- a) Type `d` to delete partitions until they are all deleted.
- b) Type `n` to create primary partition 1 that extends over the entire disk.
- c) Type `t` to set the type to an appropriate setting for the FAT32 file system, such as `c`.
- d) Type `p` to print the partition table.

The result should be similar to the following text:

```
Disk /dev/sdb: 2004 MB, 2004877312 bytes
255 heads, 63 sectors/track, 243 cylinders
```

```
Units = cylinders of 16065 * 512 = 8225280 bytes
   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1            1           243       1951866    c   W95 FAT32 (LBA)
```

e) Type `w` to write the partition table and quit.

3. Format the USB flash drive with the FAT32 file system.

```
/sbin/mkfs.vfat -F 32 -n USB /dev/sdb1
```

4. Create a destination directory and mount the USB flash drive to it.

```
mkdir -p /usbdisk
mount /dev/sdb1 /usbdisk
```

5. Copy the ESXi installation script to the USB flash drive.

```
cp ks.cfg /usbdisk
```

6. Unmount the USB flash drive.

```
umount /usbdisk
```

The USB flash drive contains the installation or upgrade script for ESXi.

When you boot the ESXi installer, point to the location of the USB flash drive for the installation or upgrade script. See [Enter Boot Options to Run an Installation or Upgrade Script](#) and [PXELINUX Configuration Files](#).

Create an Installer ISO Image with a Custom Installation or Upgrade Script

You can customize the standard ESXi installer ISO image with your own installation or upgrade script. This customization enables you to perform a scripted, unattended installation or upgrade when you boot the resulting installer ISO image.

- Linux machine
- The ESXi ISO image `VMware-VMvisor-Installer-x.x.x-XXXXXXX.x86_64.iso`, where `x.x.x` is the version of ESXi you are installing, and `XXXXXXX` is the build number of the installer ISO image
- Your custom installation or upgrade script, the `KS_CUST.CFG` kickstart file

See also [Installing ESXi by Using a Script](#) and [About the boot.cfg File](#).

1. Download the ESXi ISO image from the Broadcom Support Portal.

2. Mount the ISO image in a folder:

```
mount -o loop VMware-VMvisor-Installer-x.x.x-XXXXXXX.x86_64.iso /esxi_cdrom_mount
```

`XXXXXXX` is the ESXi build number for the version that you are installing or upgrading to.

3. Copy the contents of `esxi_cdrom` to another folder:

```
cp -r /esxi_cdrom_mount/* /esxi_cdrom
```

4. Copy the kickstart file to `/esxi_cdrom`.

```
cp KS_CUST.CFG /esxi_cdrom
```

5. Modify the `boot.cfg` file in both `/esxi_cdrom/efi/boot/boot.cfg` (for UEFI boot) and `/esxi_cdrom/boot.cfg` (for legacy BIOS boot) to specify the location of the installation or upgrade script by using the `kernelopt` option.

You must use uppercase characters to provide the path of the script, for example,

```
kernelopt=runweasel ks=cdrom:/KS_CUST.CFG
```

The installation or upgrade becomes completely automatic, without the need to specify the kickstart file during the installation or upgrade.

6. Recreate the ISO image using the `mkisofs` or the `genisoimage` command.

Command	Syntax
mkisofs	<code>mkisofs -relaxed-filenames -J -R -o custom_esxi.iso -b ISOLINUX.BIN -c BOOT.CAT -no-emul-boot -boot-load-size 4 -boot-info-table -eltorito-alt-boot -eltorito-platform efi -b EFIBOOT.IMG -no-emul-boot /esxi_cdrom</code>
genisoimage	<code>genisoimage -relaxed-filenames -J -R -o custom_esxi.iso -b ISOLINUX.BIN -c BOOT.CAT -no-emul-boot -boot-load-size 4 -boot-info-table -eltorito-alt-boot -e EFIBOOT.IMG -no-emul-boot /esxi_cdrom</code>

You can use this ISO installer image for regular boot or UEFI secure boot. However, the vSphere Lifecycle Manager cannot verify the checksum of such an ISO image and you cannot use it for upgrades by using vSphere Lifecycle Manager workflows.

The ISO image includes your custom installation or upgrade script.

Install ESXi from the ISO image.

Customizing Installations with vSphere ESXi Image Builder

You can use VMware vSphere® ESXi™ Image Builder CLI to create ESXi installation images with a customized set of updates, patches, and drivers.

You can use vSphere ESXi Image Builder with the vSphere Client or with PowerCLI to create an ESXi installation image with a customized set of ESXi updates and patches. You can also include third-party network or storage drivers that are released between vSphere releases.

You can deploy an ESXi image created with vSphere ESXi Image Builder in either of the following ways:

- By burning it to an installation DVD.
- Through vCenter Server, using the Auto Deploy feature.


In the chapter, you can read about:

- [How the vSphere ESXi Image Builder Works](#)
- [Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects](#)
- [Configure vSphere ESXi Image Builder](#)
- [Using VMware.Image Builder Cmdlets](#)
- [ESXi Image Profile Tasks](#)
- [vSphere ESXi Image Builder Workflows with PowerCLI Cmdlets](#)

How the vSphere ESXi Image Builder Works

Create ESXi image profiles for use by vSphere Auto Deploy, add custom third-party drivers to image profiles and export them, or perform upgrades.

With vSphere ESXi Image Builder you can create ESXi image profiles for use by vSphere Auto Deploy, add custom third-party drivers to existing image profiles and export to ISO or bundle, and perform upgrades. For basic concepts related to the way VMware vSphere software is created, packaged, and distributed, see [Software Packaging Units That vSphere Lifecycle Manager Can Consume](#).

Figure 13: Image Builder Architecture

You use vSphere ESXi Image Builder cmdlets for managing the software to deploy to your ESXi hosts in several different situations.

Table 9: Cases Where You Can Use vSphere ESXi Image Builder

Use Case for vSphere ESXi Image Builder	Description
Create image profiles for use by vSphere Auto Deploy	Use vSphere ESXi Image Builder to create an image profile that defines the VIBs that vSphere Auto Deploy uses to provision hosts.
Add custom third-party drivers to existing image profile and export to ISO or bundle	When you add a third-party driver or extension custom VIBs to your ESXi hosts, use vSphere ESXi Image Builder to clone the base image provided by VMware, add the custom VIBs, and export to ISO or to offline bundle ZIP file.
Perform upgrades	If you upgrade a system that includes custom extensions or drivers, you can use vSphere ESXi Image Builder to create a custom image profile that includes vSphere 8.0 compatible VIBs for the custom extensions. Export the custom image profile to an ISO or to a ZIP to upgrade your system by using vSphere Lifecycle Manager baselines.

The vSphere ESXi Image Builder cmdlets take image profiles and VIBs as input and produce various outputs.

Table 10: Input and Output to the vSphere ESXi Image Builder Cmdlets

Parameter	Description
Input	Image profiles and VIBs that are located in a software depot are used as input to PowerCLI cmdlets running on a Windows client.

Parameter	Description
Output	PowerCLI cmdlets create custom image profiles that can be exported to an ISO image or an offline depot ZIP file. ISO images are used for installation. The ZIP depot can be used by vSphere Lifecycle Manager or by <code>esxcli software</code> commands to update or install images. Image profiles are also used in vSphere Auto Deploy rules to customize the software to provision ESXi hosts with.

Image Profiles

Image profiles define the set of VIBs that an ESXi installation or update process uses. Image profiles apply to ESXi hosts provisioned with vSphere Auto Deploy. You define and manipulate image profiles with vSphere ESXi Image Builder.

Image Profile Requirements

You can create a custom image profile from scratch or clone an existing profile and add or remove VIBs. A profile must meet the following requirements to be valid.

- Each image profile must have a unique name and vendor combination.
- Each image profile has an acceptance level. When you add a VIB to an image profile with a vSphere ESXi Image Builder cmdlet, Image Builder checks that the VIB matches the acceptance level defined for the profile.
- You cannot remove VIBs that are required by other VIBs.
- You cannot include two versions of the same VIB in an image profile. When you add a new version of a VIB, the new version replaces the existing version of the VIB.

Image Profile Validation

An image profile and its VIBs must meet several criteria to be valid.

- Image profiles must contain at least one base VIB and one bootable kernel module.
- If any VIB in the image profile depends on another VIB, that other VIB must also be included in the image profile. VIB creators store that information in the `SoftwarePackage` object's `Depends` property.
- VIBs must not conflict with each other. VIB creators store conflict information in the `SoftwarePackage` object's `Conflicts` property.
- Two VIBs with the same name, but two different versions, cannot coexist. When you add a new version of a VIB, the new version replaces the existing version of the VIB.
- No acceptance level validation issues exist.

When you make a change to an image profile, vSphere ESXi Image Builder checks that the change does not invalidate the profile.

Dependency Validation

When you add or remove a VIB, vSphere ESXi Image Builder checks that package dependencies are met. Each `SoftwarePackage` object includes a `Depends` property that specifies a list of other VIBs that VIB depends on. See [Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects](#)

Acceptance Level Validation

vSphere ESXi Image Builder performs acceptance level validation each time an image profile is created or changed. vSphere ESXi Image Builder checks the acceptance level of VIBs in the image profile against the minimum allowed acceptance level of the profile. The acceptance level of the VIB is also validated each time the signature of a VIB is validated.

VIB Validation During Export

When you export an image profile to an ISO, vSphere ESXi Image Builder validates each VIB by performing the following actions.

- Checks that no conflicts exist by checking the Conflicts property of each SoftwarePackage object.
- Performs VIB signature validation. Signature validation prevents unauthorized modification of VIB packages. The signature is a cryptographic checksum that guarantees that a VIB was produced by its author. Signature validation also happens during installation of VIBs on an ESXi host and when the vSphere Auto Deploy server uses VIBs.
- Checks that VIBs follow file path usage rules. VMware tests VMwareCertified and VMwareAccepted VIBs to guarantee those VIBs always follow file path usage rules.

Working with Acceptance Levels

Hosts, image profiles, and individual VIBs have acceptance levels. VIB acceptance levels show how the VIB was tested. Understanding what each acceptance level implies, how to change levels, and what a change implies is an important part of installation and update procedures.

Acceptance levels are set for hosts, image profiles, and individual VIBs. The default acceptance level for an ESXi image or image profile is PartnerSupported.

Host acceptance levels

The host acceptance level determines which VIBs you can install on a host. You can change a host's acceptance level with ESXCLI commands. By default, ESXi hosts have an acceptance level of PartnerSupported to allow for easy updates with PartnerSupported VIBs.

Note: VMware supports hosts at the PartnerSupported acceptance level. For problems with individual VIBs with PartnerSupported acceptance level, contact your partner's support organization.

Image profile acceptance levels

The image profile acceptance level is set to the lowest VIB acceptance level in the image profile. If you want to add a VIB with a low acceptance level to an image profile, you can change the image profile acceptance level with the `Set-ESXiImageProfile` cmdlet. See [Set the Image Profile Acceptance Level](#).

The vSphere Lifecycle Manager does not display the actual acceptance level. Use vSphere ESXi Image Builder cmdlets to retrieve the acceptance level information for VIBs and image profiles.

VIB acceptance levels

A VIB's acceptance level is set when the VIB is created. Only the VIB creator can set the acceptance level.

If you attempt to provision a host with an image profile or VIB that has a lower acceptance level than the host, an error occurs. Change the acceptance level of the host to install the image profile or VIB. See [Change the Host Acceptance Level](#). Changing the acceptance level of the host changes the support level for that host.

The acceptance level of a host, image profile, or VIB lets you determine who tested the VIB and who supports the VIB. VMware supports the following acceptance levels .

VMwareCertified

The VMwareCertified acceptance level has the most stringent requirements. VIBs with this level go through thorough testing fully equivalent to VMware in-house Quality Assurance testing for the same technology. Today, only I/O Vendor Program (IOVP) program drivers are published at this level. VMware takes support calls for VIBs with this acceptance level.

VMwareAccepted

VIBs with this acceptance level go through verification testing, but the tests do not fully test every function of the software. The partner runs the tests and VMware verifies the result. Today, CIM providers and PSA plug-ins are among the VIBs published at this level. VMware directs customers with support calls for VIBs with this acceptance level to contact the partner's support organization. VIBs with the PartnerSupported acceptance level are published by a partner that VMware trusts. The partner performs all testing. VMware does not verify the results. This level is used for a new or nonmainstream technology that partners want to enable for VMware systems. Today, driver VIB technologies such as Infiniband, ATAoE, and SSD are at this level with nonstandard hardware drivers. VMware directs customers with support calls for VIBs with this acceptance level to contact the partner's support organization.

PartnerSupported**CommunitySupported**

The CommunitySupported acceptance level is for VIBs created by individuals or companies outside of VMware partner programs. VIBs at this level have not gone through any VMware-approved testing program and are not supported by VMware Technical Support or by a VMware partner.

Change the Host Acceptance Level

You can lower the host acceptance level to match the acceptance level for a VIB or image profile you want to install.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

The acceptance level of each VIB on a host must be at least as high as the acceptance level of the host. For example, you cannot install a VIB with PartnerSupported acceptance level on a host with VMwareAccepted acceptance level. You must first lower the acceptance level of the host. For more information on acceptance levels, see [Working with Acceptance Levels](#).

Warning: Changing the host acceptance level to CommunitySupported affects the supportability of your host and might affect the security of your host.

1. Retrieve the acceptance level for the VIB or image profile.

Option	Description
View information for all VIBs	<code>esxcli --server=server_name software sources vib list --depot=depot_URL</code>
View information for a specified VIB	<code>esxcli --server=server_name software sources vib list --viburl=vib_URL</code>
View information for all image profiles	<code>esxcli --server=server_name software sources profile list --depot=depot_URL</code>
View information for a specified image profile	<code>esxcli --server=server_name software sources profile get --depot=depot_URL --profile=profile_name</code>

2. View the host acceptance level.

```
esxcli --server=server_name software acceptance get
```

3. Change the acceptance level of the host.

```
esxcli
```

```
--server=server_name software acceptance set --level=acceptance_level
```

The value for `acceptance_level` can be `VMwareCertified`, `VMwareAccepted`, `PartnerSupported`, or `CommunitySupported`. The values for `acceptance_level` are case-sensitive.

Note: If the host has a higher acceptance level than the VIB or image profile you want to add, you can run commands in the `esxcli software vib` or `esxcli software profile` namespace with the `--force` option. When you use the `--force` option, a warning appears because you enforce a VIB or image profile with lower acceptance level than the acceptance level of the host and your setup is no longer consistent. The warning is repeated when you install VIBs, remove VIBs, or perform certain other operations on the host that has inconsistent acceptance levels.

Set the Image Profile Acceptance Level

If you want to add a VIB to an image profile, and the acceptance level of the VIB is lower than that of the image profile, you can clone the image profile with a lower acceptance level or change the image profile's acceptance level.

Install PowerCLI and all prerequisite software. See [Configure vSphere ESXi Image Builder](#).

You can specify `VMwareCertified`, `VMwareAccepted`, `PartnerSupported`, or `CommunitySupported` as an acceptance level of an image profile. If you lower the acceptance level, the level of support for the image profile and hosts that you provision with it changes. For more information, see [Working with Acceptance Levels](#).

1. In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	1. Download the ZIP file to a local file system. 2. Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

2. Get the acceptance level for the image profile.

```
Get-EsxImageProfile -Name string
```

3. Set the acceptance level of the image profile.

```
Set-EsxImageProfile -Name string -AcceptanceLevel level
```

Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects

Knowing the structure of `ImageProfile`, `SoftwarePackage`, and `ImageProfileDiff` objects helps you manage deployment and upgrade processes.

ImageProfile Object Properties

The `ImageProfile` object, which is accessible with the `Get-EsxImageProfile` PowerCLI cmdlet, has the following properties.

Name	Type	Description
<code>AcceptanceLevel</code>	<code>AcceptanceLevel</code>	Determines which VIBs you can add to the profile. Levels are <code>VMwareCertified</code> , <code>VMwareAccepted</code> , <code>PartnerSupported</code> , and <code>CommunitySupported</code> . See Working with Acceptance Levels .
<code>Author</code>	<code>System.String</code>	The person who created the profile. 60 characters or fewer.

Name	Type	Description
CreationTime	System.DateTime	The timestamp of creation time.
Description	System.String	The full text description of profile. No length limit.
GUID	System.String	Globally unique ID of the image profile.
ModifiedTime	System.DateTime	The timestamp of last modification time.
Name	System.String	The name of the image profile. 80 characters or fewer.
ReadOnly	System.Boolean	When set to <code>true</code> , the profile cannot be edited. Use <code>Set-EsxImageProfile -ReadOnly</code> to make your custom image profiles read-only.
Rules	ImageProfileRule[]	Any OEM hardware requirements and restrictions that the image profile might have. vSphere Auto Deploy verifies the value of this property when deploying an image profile and deploys the profile if matching hardware is available.
Vendor	System.String	The organization that publishes the profile. 40 characters or fewer.
VibList	SoftwarePackage[]	The list of VIB IDs the image contains.

SoftwarePackage Object Properties

When preparing an image profile, you can examine software packages to decide which packages are suitable for inclusion. The `SoftwarePackage` object has the following properties.

Name	Type	Description
AcceptanceLevel	AcceptanceLevel	The acceptance level of this VIB.
Conflicts	SoftwareConstraint[]	A list of VIBs that cannot be installed at the same time as this VIB. Each constraint uses the following format: package-name [<< <= = >= <<version]
Depends	SoftwareConstraint[]	A list of VIBs that must be installed at the same time as this VIB. Same constraint format as <code>Conflicts</code> property.
Description	System.String	The long description of the VIB.
Guid	System.String	The unique ID for the VIB.
LiveInstallOk	System.Boolean	True if live installs of this VIB are supported.
LiveRemoveOk	System.Boolean	True if live removals of this VIB are supported.
MaintenanceMode	System.Boolean	True if hosts must be in maintenance mode for installation of this VIB.
Name	System.String	The name of the VIB. Usually uniquely describes the package on a running ESXi system.
Provides	SoftwareProvides	The list of virtual packages or interfaces this VIB provides. See SoftwareProvide Object Properties .
ReferenceURLs	SupportReference[]	The list of <code>SupportReference</code> objects with in-depth support information. The

Name	Type	Description
		SupportReference object has two properties, Title and URL, both of type System.String.
Replaces	SoftwareConstraint[]	The list of SoftwareConstraint objects that identify VIBs that replace this VIB or make it obsolete. VIBs automatically replace VIBs with the same name but lower versions.
ReleaseDate	System.DateTime	Date and time of VIB publication or release.
SourceUrls	System.String[]	The list of source URLs from which this VIB can be downloaded.
StatelessReady	System.Boolean	True if the package supports host profiles or other technologies that make it suitable for use in conjunction with vSphere Auto Deploy.
Summary	System.String	A one-line summary of the VIB.
Tags	System.String[]	An array of string tags for this package defined by the vendor or publisher. Tags can be used to identify characteristics of a package.
Vendor	System.String	The VIB vendor or publisher.
Version	System.String	The VIB version.
VersionObject	Software.Version	The VersionObject property is of type SoftwareVersion. The SoftwareVersion class implements a static Compare method to compare two versions of strings. See SoftwareVersion Object Properties

ImageProfileDiff Object Properties

When you run the `Compare-EsxImageProfile` cmdlet, you pass in two parameters, first the reference profile, and then the comparison profile. The cmdlet returns an `ImageProfileDiff` object, which has the following properties.

Name	Type	Description
CompAcceptanceLevel	System.String	The acceptance level for the second profile that you passed to <code>Compare-EsxImageProfile</code> .
DowngradeFromRef	System.String[]	The list of VIBs in the second profile that are downgrades from VIBs in the first profile.
Equal	System.Boolean	True if the two image profiles have identical packages and acceptance levels.
OnlyInComp	System.String	The list of VIBs found only in the second profile that you passed to <code>Compare-EsxImageProfile</code> .
OnlyInRef	System.String[]	The list of VIBs found only in the first profile that you passed to <code>Compare-EsxImageProfile</code> .
PackagesEqual	System.Boolean	True if the image profiles have identical sets of VIB packages.
RefAcceptanceLevel	System.String	The acceptance level for the first profile that you passed to <code>Compare-EsxImageProfile</code> .

Name	Type	Description
UpgradeFromRef	System.String[]	The list of VIBs in the second profile that are upgrades from VIBs in the first profile.

SoftwareVersion Object Properties

The `SoftwareVersion` object lets you compare two version strings. The object includes a `Compare` static method that accepts two strings as input and returns 1 if the first version string is a higher number than the second version string. `Compare` returns 0 if two versions strings are equal. `Compare` returns -1 if the second version string is a higher number than the first string. The object has the following properties.

Name	Type	Description
Version	System.String	The part of the version before the hyphen. This part indicates the primary version.
Release	System.String	The part of the version after the hyphen. This part indicates the release version.

SoftwareConstraint Object Properties

The `SoftwareConstraint` object implements a `MatchesProvide` method. The method accepts a `SoftwareProvides` or `SoftwarePackage` object as input and returns `True` if the constraint matches the `SoftwareProvide` or the `SoftwarePackage`, or returns `False` otherwise.

The `SoftwareConstraint` object includes the following properties.

Name	Type	Description
Name	System.String	The name of the constraint. This name should match a corresponding <code>SoftwareProvideName</code> property.
Relation	System.String	An enum, or one of the following comparison indicators: <<, <=, =, >=, >>. This property can be \$null if the constraint does not have a <code>Relation</code> and <code>Version</code> property.
Version	System.String	The version to match the constraint against. This property can be \$null if the constraint does not have a <code>Relation</code> and <code>Version</code> property.
VersionObject	SoftwareVersion	The version represented by a <code>SoftwareVersion</code> object.

SoftwareProvide Object Properties

The `SoftwareProvide` object includes the following properties.

Name	Type	Description
Name	System.String	The name of the provide.
Version	System.String	The version of the provide. Can be \$null if the provide does not specify a version.
Release	System.String	The version of the provide as represented by a <code>SoftwareVersion</code> object. See Structure of ImageProfile, SoftwarePackage, and ImageProfileDiff Objects .

Configure vSphere ESXi Image Builder

Before you can run vSphere ESXi Image Builder cmdlets, you must install PowerCLI and all prerequisite software.

Before you can run vSphere ESXi Image Builder cmdlets, you must install PowerCLI and all prerequisite software. The vSphere ESXi Image Builder snap-in is included with the PowerCLI installation. If you want to manage vSphere ESXi Image Builder with PowerCLI cmdlets, verify that Microsoft .NET Framework 4.5 or 4.5.x and Windows PowerShell 3.0 or 4.0 are installed on a Microsoft Windows system. See the PowerCLI User's Guide.

1. Open PowerShell on your workstation.
2. Download a version of PowerCLI later than PowerCLI 6.5R1 from the [PowerCLI home page](#).
3. To install all PowerCLI modules, run the command: `Install-Module VMware.PowerCLI -Scope CurrentUser`. Alternatively, you can install individual PowerCLI modules by running the `Install-Module` cmdlet with the module name. If you see a warning that you are installing modules from an untrusted repository, press **y** and then press **Enter** to confirm the installation.

You can verify that the PowerCLI module is available by using the command

```
Get-Module -Name VMware.PowerCLI* -ListAvailable.
```

Review [Using VMware.Image Builder Cmdlets](#). If you are new to PowerCLI, read the *PowerCLI User's Guide*.

Use vSphere ESXi Image Builder cmdlets and other PowerCLI cmdlets and PowerShell cmdlets to manage image profiles and VIBs. Use `Get-Help cmdlet_name` at any time for command-line help.

Configure the vSphere ESXi Image Builder

Before you can use vSphere ESXi Image Builder with the vSphere Client, you must verify that the service is enabled and running.

- Verify that you have enough storage for the vSphere Auto Deploy repository. The vSphere Auto Deploy server uses the repository to store data it needs, including the rules and rule sets you create and the VIBs and image profiles that you specify in your rules.

Best practice is to allocate 2 GB to have enough room for four image profiles and some extra space. Each image profile requires approximately 400 MB. Determine how much space to reserve for the vSphere Auto Deploy repository by considering how many image profiles you expect to use.

1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere Auto Deploy service.
2. On the **Auto Deploy** page, select your vCenter Server from the drop-down menu at the top.
3. Click **Enable Image Builder** to activate the service.
The **Software Depots** tab appears.
 - [Add a Software Depot](#).
 - [Import a Software Depot](#).
 - [Clone an Image Profile](#).
 - [Create an Image Profile](#).
 - [Prepare Your System for vSphere Auto Deploy](#).

- You can change the default configuration properties of the **Image Builder Service**. For more information, see "Configuring vCenter Server" in the *vCenter Server and Host Management* documentation.

Using VMware.Image Builder Cmdlets

With VMware.Image Builder cmdlets, you can take advantage of all PowerCLI features.

VMware.Image Builder cmdlets are implemented as Microsoft PowerShell cmdlets and included in PowerCLI. You can take advantage of all PowerCLI features by using VMware.Image Builder cmdlets. Experienced PowerShell users can use VMware.Image Builder cmdlets just like other PowerShell cmdlets. If you are new to PowerShell and PowerCLI, follow these tips.

You can type cmdlets, parameters, and parameter values in the PowerCLI shell.

- Get help for any cmdlet by running `Get-Help cmdlet_name`.
- Remember that PowerShell is not case sensitive.
- Use tab completion for cmdlet names and parameter names.
- Format any variable and cmdlet output by using `Format-List` or `Format-Table` or their short forms `fl` or `ft`. See `Get-Help Format-List`.
- Use wildcards for searching and filtering VIBs and image profiles. All wildcard expressions are supported.

Passing Parameters by Name

You can pass in parameters by name in most cases and surround parameter values that contain spaces or special characters with double quotes.

```
Add-EsxSoftwarePackage -ImageProfile profile42 -SoftwarePackage "partner package 35"
```

Passing Parameters as Objects

You can pass parameters as objects if you want to do scripting and automation. You can use the technique with cmdlets that return multiple objects or with cmdlets that return a single object.

1. Bind the output of a cmdlet that returns multiple objects to a variable.

```
$profs = Get-EsxImageProfile
```

2. When you run the cmdlet that needs the object as input, access the object by position, with the list starting with 0.

```
Add-EsxSoftwarePackage -ImageProfile $profs[4] -SoftwarePackage partner-pkg
```

The example adds the specified software package to the fifth image profile in the list returned by `Get-EsxImageProfile`.

Most of the examples in the *vCenter Server Installation and Setup* documentation pass in parameters by name. [vSphere ESXi Image Builder Workflows with PowerCLI Cmdlets](#) includes examples that pass parameters as objects.

VMware.ImageBuilder Cmdlets Overview

The VMware.Image Builder component of VMware PowerCLI provides cmdlets for managing VIBs, image profiles, and other content in software depots.

vSphere 7.0 and later introduce new ways of packaging VIBs along with legacy bulletins and patches, and software depots contain base images, vendor addons and components, along with VIBs and image profiles. VMware PowerCLI 12.0 and later provide cmdlets that work with the new content in software depots.

VMware.ImageBuilder includes the following cmdlets.

Note: When you run VMware.ImageBuilder cmdlets, provide all parameters on the command line when you invoke the cmdlet. Supplying parameters in interactive mode is not recommended.

Run `Get-Help cmdlet_name` at the PowerCLI prompt for detailed reference information.

Table 11: VMware.ImageBuilder Cmdlets Used with Legacy Content in Software Depots

Cmdlet	Description
Add-EsxSoftwareDepot	Adds the software depot or ZIP file at the specified location to your current environment. Downloads metadata from the depot and analyzes VIBs for dependencies.
Remove-EsxSoftwareDepot	Disconnects from the specified software depot.
Get-EsxSoftwareDepot	Returns a list of software depots that are in the current environment. If you want to examine and manage image profiles and VIBs, you must first add the corresponding software depot to your environment.
Get-EsxSoftwarePackage	Returns a list of software package objects (VIBs). Use this cmdlet's options to filter the results.
Get-EsxImageProfile	Returns an array of ImageProfile objects from all currently added depots.
New-EsxImageProfile	Creates a new image profile. In most cases, creating a new profile by cloning an existing profile is recommended. See Clone an Image Profile with PowerCLI Cmdlets .
Set-EsxImageProfile	Modifies a local ImageProfile object and performs validation tests on the modified profile. The cmdlet returns the modified object but does not persist it.
Export-EsxImageProfile	Exports an image profile as either an ESXi ISO image for ESXi installation, or as a ZIP file.
Compare-EsxImageProfile	Returns an ImageProfileDiff structure that shows whether the two profiles have the same VIB list and acceptance level. See Working with Acceptance Levels .
Remove-EsxImageProfile	Removes the image profile from the software depot.
Add-EsxSoftwarePackage	Adds one or more new packages (VIBs) to an existing image profile.
Remove-EsxSoftwarePackage	Removes one or more packages (VIBs) from an image profile.
Set-ESXImageProfileAssociation	Associates the specified image profile with the specified ESXi system.

Table 12: VMware.ImageBuilder Cmdlets Used with New Content in Software Depots

Cmdlet	Description
Get-DepotAddons	Retrieves an array of objects that provide basic information about addons in a software depot.
Get-DepotBaseImages	Retrieves an array of objects that provide basic information about base images in a software depot.
Get-DepotComponents	Retrieves an array of objects that provide basic information about components in a software depot.
Get-DepotInfo	Retrieves basic information about the software depot located at the specified file path or URL address.
Get-DepotVibs	Retrieves an array of objects that provide basic information about VIBs in a software depot.
New-IsoImage	Generates an ISO image by using the specified software depot and software specification at the specified file path.
New-PxeImage	Generates a PXE image by using the specified software depot and software specification at the specified file path.

ESXi Image Profile Tasks

Manipulate software depots, image profiles, and VIBs by using either the VMware.Image Builder component of VMware PowerCLI or the vSphere Client.

Add a Software Depot

Add one or more software depots to the vSphere ESXi Image Builder inventory to customize image profiles.

Before you can work with software depots and customize image profiles, you must add one or more software depots to the vSphere ESXi Image Builder inventory. You can add a software depot by using the vSphere Client.

Verify that the vSphere ESXi Image Builder service is enabled and running.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.

2. On the **Software Depots** tab, click **New**.

The **Add Software Depot** window appears.

Note:

The URLs of the official VMware online depot are now unique for each customer. To update the download sources URLs, see the following KB articles: [390121](#) and [390123](#).

3. Select the type of depot that you want to create.

Option	Action
Online Depot	1. Enter a name of the depot in the inventory. 2. Enter the URL of the online depot.
Custom Depot	Enter the name of the depot in the inventory.

The VMware online software depot is located at https://dl.broadcom.com/<Download Token>/PROD/COMP/ESX_HOST/main/vmw-depot-index.xml.

4. Click **Add**.
5. Optional: Click the **Software Packages** tab to view the contents of the selected depot and additional information about the packages.
6. Optional: If you added an **Online depot**, you can also:
 - **Check for Updates** to get the latest depot packages.
 - Click **More info** to get additional depot details.

The software depot is added to the list.

- You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
- You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
- [Edit the Image Profile Association of a Host](#)
- **Remove** a custom software depot.

Import a Software Depot

If an offline depot is located on your local file system, import the ZIP file to the vSphere ESXi Image Builder inventory.

If an offline depot is located on your local file system, you can import the ZIP file to the vSphere ESXi Image Builder inventory by using the vSphere Client.

Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder](#).

1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.
2. On the **Software Depots** tab, click **Import**.
3. Enter the name of the software depot in the inventory.
4. Click **Browse** and select a ZIP file from the local system, that contains the software depot you want to import.
5. Click **Upload**.
 - You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
 - You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
 - [Edit the Image Profile Association of a Host](#).

Clone an Image Profile

Use the vSphere Client to clone image profiles.

You can use the vSphere Client to clone image profiles. You can clone an image profile when you want to make small changes to the VIB list in a profile, or if you want to use hosts from different vendors and want to use the same basic profile, but want to add vendor-specific VIBs.

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder](#).
 - Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).
 - Verify that there is at least one custom depot in the vSphere ESXi Image Builder inventory.
1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.
 2. On the **Software Depots** tab, use the drop-down menu to select the software depot that contains the image profile that you want to work with.
 3. From the list of image profiles in the depot, select the image profile that you want to clone and click **Clone**.
 4. Enter an image profile name, vendor, and description.
You must enter a unique image profile name.
 5. From the **Software depot** drop-down menu, select in which custom depot to add the new image profile.
 6. Click **Next**.
The **Select software packages** page appears.
 7. From the drop-down menu, select an acceptance level for the image profile.
The acceptance level of the VIBs you add to the base image must be at least as high as the level of the base image. If you add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Working with Acceptance Levels](#).

8. Select the VIBs that you want to add to the image profile and deselect the ones that you want to remove, and click **Next**.

Note: The image profile must contain a bootable ESXi image to be valid.

vSphere ESXi Image Builder verifies that the change does not invalidate the profile. Some VIBs depend on other VIBs and become invalid if you include them in an image profile separately. When you add or remove a VIB, vSphere ESXi Image Builder checks whether the package dependencies are met.

9. On the **Ready to complete** page, review the summary information for the new image profile and click **Finish**.
 - You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
 - You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
 - [Edit the Image Profile Association of a Host](#).

Clone an Image Profile with PowerCLI Cmdlets

Cloning a published profile is the easiest way to create a custom image profile. Cloning a profile is especially useful if you want to remove a few VIBs from a profile, or if you want to use hosts from different vendors and want to use the same basic profile, but want to add vendor-specific VIBs. VMware partners or large installations might consider creating a new profile.

- Install the PowerCLI and all prerequisite software. See [Configure vSphere ESXi Image Builder](#).
 - Verify that you have access to the software depot that contains the image profile you want to clone.
1. In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	1. Download the ZIP file to a local file system. 2. Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

2. Optional: Run the `Get-EsxImageProfile` cmdlet to find the name of the profile that you want to clone.
You can use filtering options with `Get-EsxImageProfile`.
3. Run the `New-EsxImageProfile` cmdlet to create the new profile and use the `-CloneProfile` parameter to specify the profile you want to clone.

```
New-EsxImageProfile -CloneProfile My_Profile -Name "Test Profile 42"
```

This example clones the profile named `My_Profile` and assigns it the name `Test Profile 42`. You must specify a unique combination of name and vendor for the cloned profile.

See [Examine Depot Contents](#) for some examples of filtering.

Customize the image profile by adding or removing VIBs. See [Add VIBs to an Image Profile with PowerCLI Cmdlets](#).

Create an Image Profile

You can create a new image profile by using the vSphere Client instead of cloning an existing one.

You might consider creating a new image profile if it differs significantly from the image profiles in your inventory.

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder](#).

- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).
 - Verify that there is at least one custom depot in the vSphere ESXi Image Builder inventory.
1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.
 2. From the **Software depot** drop-down menu, select in which custom depot to add the new image profile.
 3. On the Image Profiles tab, click **New Image Profile**.
 4. Enter an image profile name, vendor, and description.
You must enter a unique image profile name.
 5. Click **Next**.
The **Select software packages** page appears.
 6. From the drop-down menu, select an acceptance level for the image profile.
The acceptance level of the VIBs you add to the base image must be at least as high as the level of the base image. If you add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Working with Acceptance Levels](#).
 7. Select the VIBs that you want to add to the image profile and deselect the ones that you want to remove, and click **Next**.
Note: The image profile must contain a bootable ESXi image to be valid.
vSphere ESXi Image Builder verifies that the change does not invalidate the profile. Some VIBs depend on other VIBs and become invalid if you include them in an image profile separately. When you add or remove a VIB, vSphere ESXi Image Builder checks whether the package dependencies are met.
 8. On the **Ready to complete** page, review the summary information for the new image profile and click **Finish**.
 - You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
 - You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
 - [Edit the Image Profile Association of a Host](#).
 - Select and **Delete** an Image profile.
 - **View Software Packages** for the selected image profile.

Create a Custom ESXi ISO Image with PowerCLI Cmdlets

With ESXi Image Builder, you can customize an ESXi image profile, but not combine content from different depots to generate an ISO image. Starting with VMware PowerCLI 12.0, you can customize ISO images by using content from multiple software depots and a custom software specification.

Install VMware PowerCLI 12.0 or later.

Verify that you have access to the software depot that contains the software specification that you want to use.

The `New-IsoImage` cmdlet preserves additional metadata required by the vSphere Lifecycle Manager, such as base image, addon and component. This additional metadata is not part of ISO images that you can export by using the legacy ESXi Image Builder cmdlets.

1. Gather the required information for the software specification that you use to create a custom ISO image.
 1. Get the base image version for the required patch or upgrade by running the `Get-DepotBaseImages` cmdlet:

```
PS C:\> Get-DepotBaseImages -Depot C:\VMware-ESXi-8.xxx-xxxxxxx-depot.zip
```

The command output is:

Version	Vendor	Release date
---------	--------	--------------

```
-----
8.0.0-0.0.xxxxxx
```

```
-----
VMware, Inc.
```

```
-----
01/01/20xx 00:00:00
```

2. Get other packages, such as OEM addons, with cmdlets used with new metadata in software depots. For example:

```
PS C:\> Get-DepotAddons -Depot C:\addon-depot.zip
```

The command output is:

Name	Version	ID	Vendor	Release
date				
----	-----	--	-----	

testaddonv1	1.0.0-1	testaddonv1:1.0.0-1	ESXLifecycle QE	02/20/20xx
18:28:23				

You can also list all components in a software depot with the `Get-DepotComponents` cmdlet:

```
PS C:\> Get-DepotComponents -Depot C:\Intel-i40en_1.12.3.0-1OEM.xxxxxxx.zip
```

The command output is:

Name	Version	ID
Vendor		
----	-----	--

Intel-i40en	1.12.3.0-1OEM.xxxxxxx	Intel-i40en:1.12.3.0-1OEM.xxxxxxx Intel

You can use any number and combination of online and offline software depots.

2. Create a software specification. For example:

```
{
  "base_image": {
    "version": "8.0.0-0.0.xxxxxxx"
  },
  "add_on": {
    "name": "testaddonv1",
    "version": "1.0.0-1"
  },
  "components": {
    "Intel-i40en": "1.12.3.0-1OEM.xxxxxxx"
  }
}
```

The software specification is a JSON file that contains information about the ESXi base image and additional packages, such as a vendor add-on.

3. Generate a custom ISO image by running the `New-IsoImage` cmdlet with the parameters `Depots`, `SoftwareSpec` and `Destination`. For example:

```
New-IsoImage -Depots "c:\temp\VMware-ESXi-8.0-xxxxxxx-depot.zip" , "c:\temp\HPE-xxxxxxx-Jan20xx-Syn-
ergy-Addon-depot.zip" -SoftwareSpec "c:\temp\HPE-80xx-custom.JSON" -Destination "c:\temp\HPE-80xx-cus-
tom.iso"
```

The depot(s) include the path to the zip files for the supported ESXi version and vendor add-on. The destination include the path and file name for the custom ISO file.

You can pass additional kernel options, create a live image, overwrite existing files, or check acceptance levels for individual VIBs used during the creation of the image. For more information about the `New-IsoImage` cmdlet, see <https://code.vmware.com/docs/11794/cmdletreference/doc/New-IsoImage.html>.

You can import the new ISO image to the vSphere Lifecycle Manager depot, so that you can create upgrade baselines, which you use for host upgrade operations.

Create a Custom PXE Image with PowerCLI Cmdlets

Starting with VMware PowerCLI 12.0, you can create a custom PXE image by using any software depot and a custom software specification.

Install VMware PowerCLI 12.0 or later.

Verify that you have access to the software depot that contains the software specification you want to use.

1. Gather the required information for the software specification that you use to create a custom PXE image.

1. Get the base image version for the required patch or upgrade by running the `Get-DepotBaseImages` cmdlet:

```
PS C:\> Get-DepotBaseImages -Depot C:\VMware-ESXi-8.xxxx-xxxxx-depot.zip
```

The command output is:

Version	Vendor	Release date
-----	-----	-----
8.x.x.xxx.xxxxx	VMware, Inc.	04/29/20xx 00:00:00

2. Get other packages, such as OEM addons, with cmdlets used with new metadata in software depots. For example:

```
PS C:\> Get-DepotAddons -Depot C:\addon-depot.zip
```

The command output is:

Name	Version	ID	Vendor	Release date
----	-----	--	-----	-----
testaddonv1	1.0.0-1	testaddonv1:1.0.0-1	ESXLifecycle QE	02/20/20xx
18:28:23				

You can also list all components in a software depot with the `Get-DepotComponents` cmdlet:

```
PS C:\> Get-DepotComponents -Depot C:\Intel-i40en_1.12.3.0-10EM.700.1.0.15843807_18058526.zip
```

The command output is:

Name	Version	ID
-----	-----	--
Intel-i40en	1.12.3.0-10EM.xxxxx	Intel-i40en:1.12.3.0-10EM.xxxxx
		Intel

You can use any number and combination of online and offline software depots.

2. Create a software specification. For example:

```
{
  "base_image": {
    "version": "8.0.xxxxx"
  },
  "add_on": {
    "name": "testaddonv1",
    "version": "1.0.0-1"
  },
  "components": {
    "Intel-i40en": "1.12.3.0-10EM.xxxxx"
  }
}
```

The software specification is a JSON file that contains information about the ESXi base image and additional packages, such as a vendor add-on.

3. Generate a custom PXE image by running the `New-PxeImage` cmdlet with the parameters `Depots`, `SoftwareSpec` and `Destination`. For example:

```
New-PxeImage -Depots "c:\temp\VMware-ESXi-8.0xxxxx-xxxxx-depot.zip" , "c:\temp\HPE-8.0xxxxx-xxx-Syner-
gy-Addon-depot.zip" -SoftwareSpec "c:\temp\HPE-xxx-custom.JSON" -Destination "C:\pxe-image"
```

The depot(s) include the path to the zip files for the supported ESXi version and vendor add-on. The destination include the path and file name for the custom PXE file.

You can pass additional kernel options, create a live image, overwrite existing files, or check acceptance levels for individual VIBs used during the creation of the image. For more information about the `New-PxeImage` cmdlet, see <https://code.vmware.com/docs/11794/cmdletreference/doc/New-PxeImage.html>.

You can use the PXE image in remediation workflows of PXE booted ESXi hosts.

Edit an Image Profile

You can edit image profiles by using the vSphere Client. You can change the name, details and VIB list of an image profile.

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder](#).
 - Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).
 - Verify that there is at least one custom depot in the vSphere ESXi Image Builder inventory.
1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.
 2. On the **Software Depots** tab, use the drop-down menu to select the software depot that contains the image profile that you want to work with.
 3. On the **Image Profiles** tab, select the image profile that you want to edit and click **Edit**.
The **Edit Image Profile** wizard appears.
 4. Optional: Change the name, vendor and description information of the image profile.
 5. Click **Next**.
The **Select software packages** page appears.
 6. From the drop-down menu, select an acceptance level for the image profile.
The acceptance level of the VIBs you add to the base image must be at least as high as the level of the base image. If you add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Working with Acceptance Levels](#).
 7. Select the VIBs that you want to add to the image profile and deselect the ones that you want to remove, and click **Next**.
Note: The image profile must contain a bootable ESXi image to be valid.
vSphere ESXi Image Builder verifies that the change does not invalidate the profile. Some VIBs depend on other VIBs and become invalid if you include them in an image profile separately. When you add or remove a VIB, vSphere ESXi Image Builder checks whether the package dependencies are met.
 8. On the **Ready to complete** page, review the summary information for the edited image profile and click **Finish**.
 - You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
 - You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
 - [Edit the Image Profile Association of a Host](#).

Add VIBs to an Image Profile with PowerCLI Cmdlets

You can add one or more VIBs to an image profile if that image profile is not set to read only. If the new VIB depends on other VIBs or conflicts with other VIBs in the profile, a message is displayed at the PowerShell prompt and the VIB is not added.

Install the PowerCLI and all prerequisite software. See [Configure vSphere ESXi Image Builder](#)

You can add VIBs from VMware or from VMware partners to an image profile. If you add VMware VIBs, vSphere ESXi Image Builder performs validation. If you add VIBs from two or more OEM partners simultaneously, no errors are reported but the resulting image profile might not work. Install VIBs from only one OEM vendor at a time.

If an error about acceptance level problems appears, change the acceptance level of the image profile and the acceptance level of the host. Consider carefully whether changing the host acceptance level is appropriate. VIB acceptance levels are set during VIB creation and cannot be changed.

You can add VIBs even if the resulting image profile is invalid.

Note: VMware can support only environments and configurations that are proven to be stable and fully functional through rigorous and extensive testing. Use only those supported configurations. You can use custom VIBs if you lower your host acceptance level, and as a result, supportability. In that case, track the changes you made, so you can revert them if you want to remove custom VIBs and restore the host acceptance level to the default (Partner Supporter) later. See [Working with Acceptance Levels](#).

1. In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	1. Download the ZIP file to a local file system. 2. Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

2. Run the `Get-EsxImageProfile` cmdlet to list all image profiles in all currently visible depots.

The cmdlet returns all available profiles. You can narrow your search by using the optional arguments to filter the output.

3. Clone the profile.

```
New-EsxImageProfile -CloneProfile My_Profile -Name "Test Profile 42" -Vendor "My Vendor"
```

Image profiles published by VMware and its partners are read only. To make changes, you must clone the image profile. The `vendor` parameter is required.

4. Run the `Add-EsxSoftwarePackage` cmdlet to add a new package to one of the image profiles.

```
Add-EsxSoftwarePackage -ImageProfile My_Profile -SoftwarePackage partner-package
```

The cmdlet runs the standard validation tests on the image profile. If validation succeeds, the cmdlet returns a modified, validated image profile. If the VIB that you want to add depends on a different VIB, the cmdlet displays that information and includes the VIB that can resolve the dependency. If the acceptance level of the VIB that you want to add is lower than the image profile acceptance level, an error occurs.

Compare Image Profiles

You can compare two image profiles by using the vSphere Client to check if they have the same VIB list, version, or acceptance level.

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder](#).

- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).
1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.
 2. On the **Software Depots** tab, use the drop-down menu to select the software depot that contains the image profile that you want to work with.
 3. On the **Image Profiles** tab, select an image profile and click **Compare To**.
The **Compare Image Profile** wizard appears.
 4. Click **Change** to select a second image profile.
The **Select Image Profile** page appears.
 5. Select a software depot from the drop-down menu and click on the second image profile.
 6. In the **Compare Image Profile** page, select a comparison option from the **Software packages** drop-down menu.
The left side of the list displays details of the VIBs that the first chosen image profile contains. The right part of the list provides information about the second image profile. The VIBs marked as *Same* are identical in both profiles. VIBs that are present in one of the image profiles are marked as *Missing* next to the image profile that they are not present in.

Compare Image Profiles with PowerCLI Cmdlets

You can compare two image profiles by using the `Compare-EsxImageProfile` cmdlet, for example, to see if they have the same VIB list or acceptance level. Comparing image profiles or their properties is also possible by using the PowerShell comparison operators.

Install the PowerCLI and all prerequisite software. See [Configure vSphere ESXi Image Builder](#).

1. In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	1. Download the ZIP file to a local file system. 2. Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

2. Optional: Run the `Get-EsxImageProfile` cmdlet to view a list of all image profiles in all available depots.
In the list, you can locate the names of the image profiles you want to compare.
3. Before comparing the image profiles, assign them to variables.
For example, you can create variables `$imageProfile1` and `$imageProfile2` to hold the names of the compared images profiles.

```

$imageProfile1
= Get-EsxImageProfile -Name "ImageProfile1"
$imageProfile2
= Get-EsxImageProfile -Name "ImageProfile2"

```
4. Compare the two image profiles by using the `Compare-EsxImageProfile` cmdlet or the `-eq` comparison operator, which returns a Boolean value.
 - Compare the two image profiles to get a full description of the differences by using the `Compare-EsxImageProfile` cmdlet.

```

Compare-EsxImageProfile -ReferenceProfile

```

```
$imageProfile1 -ComparisonProfile $imageProfile2
```

- Compare the two image profiles by VIB list and acceptance level using the `-eq` comparison operator.

```
if ($imageProfile1 -eq $imageProfile2) {
    Write-host "Successfully verified that both image profiles are equal."
} else {
    Write-host "Failed to verify that the image profiles are equal."
}
```

- Compare the two image profiles by a specific property using the `-eq` comparison operator.

```
if ($imageProfile1.vendor -eq $imageProfile2.vendor) {
    Write-host "Successfully verified that both image profiles are equal."
} else {
    Write-host "Failed to verify that the image profiles are equal."
}
```

Move an Image Profile to a Different Software Depot

You can move image profiles between custom depots by using the vSphere Client and modify it without affecting the source depot's configuration.

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).
- Verify that there is at least one custom depot in the vSphere ESXi Image Builder inventory.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.

2. On the **Software Depots** tab, use the drop-down menu to select the software depot that contains the image profile that you want to work with.
3. On the **Image Profiles** tab, select an image profile and click **Move to**.
4. From the drop-down menu, select the custom depot in which you want to move the image profile.
5. Click **OK**.

The image profile is in the new software depot.

Export an Image Profile to ISO or Offline Bundle ZIP

You can export an image profile to an ISO image or a ZIP file by using the vSphere Client.

You can export an image profile to an ISO image or a ZIP file by using the vSphere Client. You can use the ISO image as an ESXi installer or to upgrade hosts with vSphere Lifecycle Manager. The ZIP file contains metadata and the VIBs of the image profile. You can use it for ESXi upgrades or as an offline depot.

- Verify that the vSphere ESXi Image Builder service is enabled and running. See [Configure the vSphere ESXi Image Builder](#).
- Add or import a software depot to the vSphere ESXi Image Builder inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere ESXi Image Builder service.

- On the **Software Depots** tab, use the drop-down menu to select the software depot that contains the image profile that you want to work with.

- On the **Image Profiles** tab, select the image profile that you want to export and click **Export**.

The **Export Image Profile** window appears.

- Select the type of the exported file.

Option	Description
ISO	Exports the image profile to a bootable ISO image. If you want to create an ISO image that you can burn to a CD or DVD and use to start a stateless ESXi instance, select the Do not include an installer on the ISO check box.
ZIP	Exports the image profile to a ZIP file.

- Optional: If you want to bypass the acceptance level verification of the image profile, select **Skip acceptance level checking**.

- Click **Ok**.

The **Download** link starts generating in the "Download Image Profiles" column of the selected image profile.

- When the image generates successfully, click **Download** to save the exported file.

- You can associate an image profile with a new vSphere Auto Deploy rule to provision ESXi hosts. See [Create a Deploy Rule](#) or [Clone a Deploy Rule](#).
- You can associate an image profile with an ESXi host. See [Add a Host to the vSphere Auto Deploy Inventory](#).
- [Edit the Image Profile Association of a Host](#).

Export an Image Profile to an ISO or Offline Bundle ZIP with PowerCLI Cmdlets

You can export an image profile to an ISO image or a ZIP file of component files and folders. You cannot create both by running the cmdlet once. You can use the ISO image as an ESXi installer or upload the ISO into vSphere Lifecycle Manager for upgrades.

Install the PowerCLI and all prerequisite software. See [Configure vSphere ESXi Image Builder](#).

- In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	- Download the ZIP file to a local file system. - Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- Run `Export-EsxImageProfile` to export the image profile.

Export Format	Cmdlet
ISO images	<code>Export-EsxImageProfile</code> with the <code>-ExportToIso</code> parameter
Offline depot ZIP files	<code>Export-EsxImageProfile</code> with the <code>-ExportToBundle</code> parameter

For the ISO image, vSphere ESXi Image Builder validates VIB signatures, adds VIB binaries to the image, and downloads the image to the specified location. For the ZIP file, vSphere ESXi Image Builder validates VIB signatures and downloads the VIB binaries to the specified location.

Exporting an Image Profile

Follow these steps to export an image profile to an ISO image.

1. Add the software depot.

```
Add-EsxSoftwareDepot -DepotUrl url_or_file
```

2. View all available image profiles to find the name of the image profile to export.

```
Get-EsxImageProfile
```

3. Export the image profile.

```
Export-EsxImageProfile -ImageProfile "myprofile" -ExportToIso -FilePath iso_name
```

Follow these steps to export an image profile to a ZIP file of component files and folders.

1. Add the software depot.

```
Add-EsxSoftwareDepot -DepotUrl url_or_file
```

2. View all available image profiles to find the name of the image profile to export.

```
Get-EsxImageProfile
```

3. Export the image profile.

```
Export-EsxImageProfile -ImageProfile "myprofile" -ExportToBundle -FilePath C:\my_bundle.zip
```

Use the ISO image in an ESXi installation or upload the ISO image into vSphere Lifecycle Manager to perform upgrades.

Use the ZIP file to upgrade an ESXi installation.

- Import the ZIP file into vSphere Lifecycle Manager for use with patch baselines.
- Download the ZIP file to an ESXi host or a datastore and run `esxcli software vib` commands to import the VIBs in the ZIP file.

See the *vSphere Upgrade* documentation.

Regenerate an Image Profile

If you use Auto Deploy to add stateful ESXi hosts to a cluster that you manage by using an image, all hosts must have the same software specification.

When an Auto Deploy rule is created and the target cluster is managed by an image, a Preboot Execution Environment (PXE) image is created based on the image specification of the cluster. The generated PXE image is cached and is not automatically updated. As a result, if you modify the image specification in the vSphere Lifecycle Manager, you must update the PXE image manually.

For information how to create a rule to add hosts to a cluster managed by an image, see *Use Auto Deploy to Add a Host to a Cluster Managed by an Image* from the *vSphere Lifecycle Manager* documentation.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Deploy Rules** tab, select the desired rule.

You must select a rule, which matches ESXi hosts to a cluster that you manage by an image.

3. If the rule is active, first you must deactivate the rule.
 - a) Click the **Activate/Deactive Rules** tab
 - b) In the dialog box, select the rule.
 - c) In the dialog box, select **Deactivate** and click **OK**.
4. Select **Recreate Image Profile** and in the confirmation dialog box, click **Recreate**.
5. Optional: Activate the rule again.
 - a) Click the **Activate/Deactive Rules** tab
 - b) In the dialog box, select the rule.
 - c) In the dialog box, select **Activate** and click **OK**.

The PXE image is running the latest image specification.

Preserve Image Profiles Across Sessions with PowerCLI

You can export the image profile to a ZIP file software depot and add that depot in the next session.

When you create an image profile and exit the PowerCLI session, the image profile is no longer available when you start a new session. You can export the image profile to a ZIP file software depot, and add that depot in the next session.

Install the PowerCLI and all prerequisite software. See [Configure vSphere ESXi Image Builder](#).

1. In a PowerCLI session, create an image profile, for example by cloning an existing image profile and adding a VIB.
2. Export the image profile to a ZIP file by calling `Export-ESXImageProfile` with the `ExportToBundle` parameter.


```
Export-ESXImageProfile -ImageProfile "my_profile" -ExportToBundle -FilePath
"C:\isos\temp-base-plus-vib25.zip"
```

3. Exit the PowerCLI session.

4. When you start a new PowerCLI session, add the depot that contains your image profile to access it.

```
Add-ESXSoftwareDepot "C:\isos\temp-base-plus-vib25.zip"
```

Compare VIBs with PowerCLI Cmdlets

You can compare two VIBs or their properties by using the PowerShell comparison operators.

Install the PowerCLI and all prerequisite software. See [Configure vSphere ESXi Image Builder](#).

1. In a PowerCLI session, run the `Add-ESXSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-ESXSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	1. Download the ZIP file to a local file system. 2. Run <code>Add-ESXSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

2. Optional: Run the `Get-ESXSoftwarePackage` cmdlet to view all available VIBs.

In the list, you can locate the names of the VIBs you want to compare.

3. Before comparing the VIBs, assign them to variables.

For example, you can create variables `$vib1` and `$vib2` to hold the names of the compared VIBs.

```
$vib1 = Get-ESXSoftwarePackage -Name "ReferenceVIB"
$vib2 = Get-ESXSoftwarePackage -Name "ComparisonVIB"
```


4. Use a comparison operator to compare the VIBs by contents and acceptance level or by a specific property.

- Compare the two VIBs by their contents and acceptance level.

```
if ($vib1 -eq $vib2) {
    Write-host "Successfully verified that both VIBs are equal."
} else {
    Write-host "Failed to verify that the VIBs are equal."
}
```

- Compare a specific property of the VIBs by using a comparison operator such as `-eq`, `-lt`, `-le`, `-gt`, or `-ge`.

```
if ($vib1.VersionObject -lt $vib2.VersionObject) {
    Write-host "Successfully verified that both the VIBs are equal."
} else {
    Write-host "Failed to verify that the VIBs are equal."
}
```

vSphere ESXi Image Builder Workflows with PowerCLI Cmdlets

vSphere ESXi Image Builder workflows are examples for cmdlet use and do not represent actual tasks.

vSphere ESXi Image Builder workflows are examples for cmdlet use. Workflows do not represent actual tasks, but illustrate how you might explore different ways of using a cmdlet. Administrators trying out the workflows benefit from some experience with PowerCLI, Microsoft PowerShell, or both.

Examine Depot Contents

You can examine software depots and VIBs with vSphere ESXi Image Builder cmdlets by using all kinds of wildcard expressions.

Verify that PowerCLI and prerequisite software is installed. See [Configure vSphere ESXi Image Builder](#).

The workflow itself passes parameters by name. However, you can pass parameters as objects by accessing variables.

You can use filtering options and wildcard expressions to examine depot contents.

1. In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	1. Download the ZIP file to a local file system. 2. Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

2. Retrieve image profiles.

You can filter by vendor, name, and acceptance level.

- `Get-EsxImageProfiles`

Returns an array of `ImageProfile` objects from all depots you added to the session.

- `Get-EsxImageProfile -Vendor "C*"`

Returns all image profiles created by a vendor with a name that starts with the letter C.

3. Retrieve software packages by using the `Get-ExsSoftwarePackage` cmdlet.

You can filter, for example by vendor or version, and you can use the standard PowerShell wildcard characters.

- `Get-ExsSoftwarePackage -Vendor "V*"`

Returns all software packages from a vendor with a name that starts with the letter V.

- `Get-ExsSoftwarePackage -Vendor "V*" -Name "*scsi*"`

Returns all software packages with a name that contains the string `scsi` in it from a vendor with a name that starts with the letter V.

- `Get-ExsSoftwarePackage -Version "2.0*"`

Returns all software packages with a version string that starts with 2.0.

4. Use `-Newest` to find the latest package.

- `Get-ExsSoftwarePackage -Vendor "V*" -Newest`

Returns the newest package for the vendors with a name that starts with the letter V, and displays the information as a table.

- `Get-ExsSoftwarePackage -Vendor "V*" -Newest | format-list`

Returns detailed information about each software package by using a pipeline to link the output of the request for software packages to the PowerShell `format-list` cmdlet.

5. View the list of VIBs in the image profile.

```
(Get-ExsImageProfile -Name "Robin's Profile").VibList
```

`VibList` is a property of the `ImageProfile` object.

6. Retrieve software packages released before or after a certain date by using the `CreatedBefore` or `CreatedAfter` parameter.

```
Get-ExsSoftwarePackage -CreatedAfter 7/1/2010
```

Depot Content Examination Using Variables

This workflow example examines depot contents by passing in parameters as objects accessed by position in a variable, instead of passing in parameters by name. You can run the following commands in sequence from the PowerCLI prompt. Replace names with names that are appropriate in your installation.

```
Get-ExsSoftwarePackage -Vendor "V*"
Get-ExsSoftwarePackage -Vendor "V*" -Name "r*"
Get-ExsSoftwarePackage -Version "2.0*"
$ip1 = Get-ExsImageProfile -name ESX-5.0.0-123456-full
$ip1.VibList
Get-ExsSoftwarePackage -CreatedAfter 7/1/2010
```

Create Image Profiles by Cloning Workflow

Use vSphere ESXi Image Builder cmdlets to check available depots, add a depot, view image profile information, and to clone a new image profile.

You can use vSphere ESXi Image Builder cmdlets to check which depots are available, to add a depot, to view image profile information, and to create a new image profile by cloning one of the available image profiles.

Verify that PowerCLI and prerequisite software is installed. See [Configure vSphere ESXi Image Builder](#).

Published profiles are usually read-only and cannot be modified. Even if a published profile is not read-only, cloning instead of modifying the profile is a best practice, because modifying the original profile erases the original. You cannot revert to the original, unmodified profile except by reconnecting to a depot.

A profile cloning workflow might include checking the current state of the system, adding a software depot, and cloning the profile.

1. In a PowerShell window, check whether any software depots are defined for the current session.

```
$DefaultSoftwareDepots
```

PowerShell returns the currently defined depots, or nothing if you just started PowerShell.

2. If the depot containing the profile that you want to clone does not appear in the results, add it to the current session.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl depot_url</code> .
ZIP file	1. Download the ZIP file to a local file path. 2. Run <code>Add-EsxSoftwareDepot -DepotUrl C:\file_path\offline-bundle.zip</code>

PowerShell adds the specified depot to your current session and lists all current depots.

3. Optional: Check the `$DefaultSoftwareDepots` variable, which now returns the newly added depot.
4. View all available image profiles.

```
Get-EsxImageProfile
```

5. To clone an image profile, enter its name, a new name for the new profile, and a name of the vendor.

```
$ip = New-EsxImageProfile -CloneProfile base-tbd-v1 -Name "Test Profile 42" -Vendor "Vendor20"
```

6. Optional: View the newly created image profile, `$ip`.

PowerShell returns the information about the image profile in tabular format.

Name	Vendor	Last Modified	Acceptance Level
----	-----	-----	-----
Test Profile 42	Vendor20	9/15/2010 5:45:43...	PartnerSupported

Creating Image Profile by Cloning Using Variables

This workflow example repeats the steps of this workflow by passing in parameters as objects accessed by position in a variable, instead of passing in parameters by name. You can run the following cmdlets in sequence from the PowerCLI prompt.

```
$DefaultSoftwareDepots
Add-EsxSoftwareDepot -DepotUrl depot_url
$DefaultSoftwareDepots
$profs = Get-EsxImageProfile
$profs
$ip = New-EsxImageProfile -CloneProfile $profs[2] -Name "new_profile_name" -Vendor "my_vendor"
$ip
```

Create New Image Profiles Workflow

You can clone an existing image profile or create a new image profile, for which you must define dependencies and acceptance levels.

In most situations, you create an image profile by cloning an existing profile. Some VMware customers or partners might need to create a new image profile. Pay careful attention to dependencies and acceptance levels if you create an image profile from scratch.

- PowerCLI and prerequisite software is installed. See [Configure vSphere ESXi Image Builder](#).

- You have access to a depot that includes a base image and one or more VIBs. VMware and VMware partners have public depots, accessible by a URL. VMware or VMware partners can create a ZIP file that you can unzip to your local environment and access by using a file path.

The system expects that the acceptance level of the VIBs you add to the base image is at least as high as the level of the base image. If you have to add a VIB with a lower acceptance level to the image profile, you must lower the image profile acceptance level. For more information, see [Set the Image Profile Acceptance Level](#).

As an alternative to specifying the parameters on the command line, you can use the PowerShell prompting mechanism to specify string parameters. Prompting does not work for other parameters such as objects.

- In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	- Download the ZIP file to a local file system. - Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- Run the `Get-EsxImageProfile` cmdlet to list all image profiles in all currently visible depots. You can narrow your search by using the optional arguments to filter the output.

```
Get-EsxSoftwarePackage -CreatedAfter 7/1/2010
```

- Create a new profile, assign it a name and vendor, and add a base package.

```
New-EsxImageProfile -NewProfile -Name "Test #2" -vendor "Vendor42" -SoftwarePackage esx-base[0],esx-xlibs[0]
```

The example uses the `esx-base` package. In most cases, you include the `esx-base` package when you create a new image profile. Names that contain spaces are surrounded by quotes.

- Use a pipeline to pass the new image profile to `format-list` for detailed information about the new package.

```
(Get-EsxImageProfile -Name "Test #2").VibList | format-list
```

Creating Image Profiles from Scratch Using Variables

This command sequence repeats the steps of the workflow, but passes parameters as objects, accessed by position in a variable, instead of passing parameters by name. You can run the following commands in sequence at the PowerCLI prompt.

```
Add-EsxSoftwareDepot depoturl
$pkgs = Get-EsxSoftwarePackage -CreatedAfter 7/1/2010
$ip2 = New-EsxImageProfile -NewProfile -Name "Test #2" -vendor "Vendor42" -SoftwarePackage $pkgs[0]
$ip2.VibList | format-list
```

Edit Image Profiles Workflow

You can create a custom image by cloning and editing an image profile by using PowerCLI.

You can create a custom image by cloning and editing an image profile by using PowerCLI. You can add or remove one or more VIBs in the existing profile. If adding or removing VIBs prevents the image profile from working correctly, an error occurs.

- PowerCLI and prerequisite software is installed. See [Configure vSphere ESXi Image Builder](#).

- You have access to a depot that includes a base image and one or more VIBs. VMware and VMware partners make public depots, accessible by a URL, available. VMware or VMware partners can create a ZIP file that you can download to your local environment and access by using a file path.

- In a PowerCLI session, run the `Add-EsxSoftwareDepot` cmdlet for each depot you want to work with.

Option	Action
Remote depot	Run <code>Add-EsxSoftwareDepot -DepotUrl <depot_url></code> .
ZIP file	- Download the ZIP file to a local file system. - Run <code>Add-EsxSoftwareDepot -DepotUrl C:\<file_path>\<offline-bundle>.zip</code>

The cmdlet returns one or more `SoftwareDepot` objects.

- Use a pipeline to pass the image profile you intend to edit to `format-list` to see detailed information. In this example, the image profile created in [Create New Image Profiles Workflow](#) contains only the base image. A newly created image profile is not included in the depot. Instead, you access the image profile by name or by binding it to a variable.

```
Get-EsxImageProfile "Test #2" | format-list
```

PowerShell returns the information.

```
Name           : Test #2
Vendor         : Vendor42
...
VibList        : {esx-base 5.0.0.-...,}
```

- Optional: If you are adding a VIB with a lower acceptance level than that of the image profile, change the acceptance level of the image profile.

```
Set-EsxImageProfile -ImageProfile "Test #2" -AcceptanceLevel VMwareAccepted
```

PowerShell returns the information about the changed profile in tabular format.

```
Name           Vendor           Last Modified      Acceptance Level
----           -
Test #2        Vendor42           9/22/2010 12:05:... VMwareAccepted
```

- Add a software package (VIB) to the image profile. You can add the package by name.

```
Add-EsxSoftwarePackage -ImageProfile "Test #2"
                        -SoftwarePackage NewPack3
```

PowerShell returns the information about the image profile in tabular format.

```
Name           Vendor           Last Modified      Acceptance Level
----           -
Test #2        Vendor42           9/22/2010 12:05:... VMwareAccepted
```

Note: If an error occurs when you add the software package, you might have a problem with acceptance levels, see [Working with Acceptance Levels](#)

- View the image profile again.

```
Get-EsxImageProfile "Test #2" | format-list
```

The VIB list is updated to include the new software package and the information is displayed.

```
Name           : Test #2
```

```
Vendor      : Vendor42
...
VibList     : {esx-base 5.0.0.-..., NewPack3}
```

Editing Image Profiles by Using Variables

This cmdlet sequence repeats the steps of the workflow but passes parameters as objects, accessed by position in a variable, instead of passing parameters by name. You can run the following cmdlets in sequence from the PowerCLI prompt.

```
Add-ESXSoftwareDepot -DepotUrl depot_url
$ip2 = Get-ESXImageProfile -name "Test #2"
$ip2 | format-list
Set-ESXImageProfile -ImageProfile $ip2 -AcceptanceLevel VMwareAccepted
Add-ESXImageSoftwarePackage -ImageProfile $ip2 -SoftwarePackage NewPack3
$ip2 | format-list
```

Installing ESX

You can install ESXi interactively, with a scripted installation, or with vSphere Auto Deploy.

In this chapter, you can read about:

- [Installing ESXi Interactively](#)
- [Installing ESXi by Using a Script](#)
- [How to Boot an ESXi Host from a Network Device](#)
- [Installing ESXi Using vSphere Auto Deploy](#)
- [Troubleshooting vSphere Auto Deploy](#)

Installing ESXi Interactively

Use the interactive installation option for small deployments of fewer than five hosts.

In a typical interactive installation, you boot the ESXi installer and respond to the installer prompts to install ESXi to the local host disk. The installer reformats and partitions the target disk and installs the ESXi boot image. If you have not installed ESXi on the target disk before, all data on the drive is overwritten, including hardware vendor partitions, operating system partitions, and associated data.

Note: To ensure that you do not lose any data, migrate the data to another machine before you install ESXi.

If you are installing ESXi on a disk that contains a previous installation of ESXi or ESX, or a VMFS datastore, the installer provides you with options for upgrading. See the *vSphere Upgrade* documentation.

Interactive ESXi Installation

You boot the server from an attached vSphere ESXi CD or DVD, or from a bootable USB device, or by PXE booting the server from a location on the network. You follow the prompts in the installation wizard to install ESXi to disk.

- [Interactive ESXi Installation](#)
- [Install ESXi Interactively](#)
- [Install ESXi on a Software iSCSI Disk](#)

Install ESXi Interactively

You use the ESXi CD/DVD or a USB flash drive to install the ESXi software onto a SAS, SATA, SCSI hard drive, or USB drive.

- You must have the ESXi installer ISO in one of the following locations:
 - On CD or DVD. If you do not have the installation CD/DVD, you can create one. See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#)
 - On a USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).

Note: You can also PXE boot the ESXi installer to run an interactive installation or a scripted installation. See [Network Booting the ESXi Installer](#).

- Verify that the server hardware clock is set to UTC. This setting is in the system BIOS or UEFI.
- Verify that a keyboard and monitor are attached to the machine on which the ESXi software is installed. Alternatively, use a remote management application. See [Using Remote Management Applications](#).
- Consider disconnecting your network storage. This action decreases the time it takes the installer to search for available disk drives. When you disconnect network storage, any files on the disconnected disks are unavailable at installation.

Do not disconnect a LUN that contains an existing ESX or ESXi installation. Do not disconnect a VMFS datastore that contains the Service Console of an existing ESX installation. These actions can affect the outcome of the installation.

- Gather the information required by the ESXi installation wizard. See [Required Information for ESXi Installation](#).
- Verify that ESXi Embedded is not present on the host machine. ESXi Installable and ESXi Embedded cannot exist on the same host.

1. Insert the ESXi installer CD/DVD into the CD/DVD-ROM drive, or attach the Installer USB flash drive and restart the machine.
2. Set the BIOS or UEFI to boot from the CD-ROM device or the USB flash drive.

Note: If your system has data processing units (DPUs), you can only use UEFI to install and boot ESXi on the DPUs. See your hardware vendor documentation for information on changing boot order.

3. On the welcome screen, press Enter to continue.
4. Accept the End User License Agreement by pressing **Enter**.

Starting with ESXi 8.0 Update 3, after the scanning for available devices completes, if your system has DPUs, you see them automatically listed with their respective PCI slots. You no longer select a slot. The DPU devices must be identical: same vendor, same hardware version and same firmware.

5. On the **Select a Disk to Install or Upgrade ESXi** screen, select the drive on which to install ESXi and press **Enter**. Press F1 for information about the selected disk.

Note: Do not rely on the disk order in the list to select a disk. The disk order is determined by the BIOS or UEFI and might be out of order. This might occur on systems where drives are continuously being added and removed.

If you select a disk that contains data, the **Confirm Disk Selection** page appears.

If you are installing on a disc with a previous ESXi or ESX installation or VMFS datastore, the installer provides several choices.

Important: If you are upgrading or migrating an existing ESXi installation, see the *VMware ESXi Upgrade* documentation.

If you select a disk that is in vSAN disk group, the resulting installation depends on the type of disk and the group size:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group are wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD is wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD is wiped.

6. Select the keyboard type for the host.

You can change the keyboard type after installation in the direct console.

7. Enter the root password for the host.

You can change the password after installation in the direct console.

8. In the **Confirm Install** screen, if your vSphere system has DPUs, you see each listed on a separate row. Press **F11** to confirm the start of the installation.

Starting with 8.0 Update 3, if your systems has DPUs, you see a single progress bar for the ESXi and DPU installation, with dynamic updates to the label showing what stage of the installer is being run.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

If you select an SD or USB device, you see a warning that prompts you to select a persistent disk to store the ESXi-OSData partition. In the **Select a Disk to store ESX OSData** screen, select a persistent storage device with minimum 32 GB available space.

9. When the installation is complete, remove the installation CD, DVD, or USB flash drive.

Starting with 8.0 Update 3, if the installation is not successful on any of the targets, you see the **Operation Failed** screen. Record the error log in the screen to help troubleshooting.

10. Press **Enter** to reboot the host.

11. Set the first boot device to be the drive on which you installed ESXi in step 5.

For information about changing boot order, see your hardware vendor documentation.

Note: UEFI systems might require additional steps to set the boot device. See [Host Fails to Boot After You Install ESXi in UEFI Mode](#)

After the installation is complete, you can migrate existing VMFS data to the ESXi host.

You can boot a single machine from each ESXi image. Booting multiple devices from a single shared ESXi image is not supported.

Set up basic administration and network configuration for ESXi. See [After You Install and Set Up ESXi](#).

Install ESXi on a Software iSCSI Disk

When you install ESXi to a software iSCSI disk, you must configure the target iSCSI qualified name (IQN).

- Verify that the target IQN is configured in the iBFT BIOS target parameter setting. This setting is in the option ROM of the network interface card (NIC) to be used for the iSCSI LUN. See the vendor documentation for your system.
- Deactivate the iBFT adapter option to boot to the iSCSI target. This action is necessary to make sure that the ESXi installer boots, rather than the iSCSI target. When you start your system, follow the prompt to log in to your iBFT adapter and deactivate the option to boot to the iSCSI target. See the vendor documentation for your system and iBFT adapter. After you finish the ESXi installation, you can reenable the option to boot from the LUN you install ESXi on.

During system boot, the system performs a Power-On Self Test (POST), and begins booting the adapters in the order specified in the system BIOS. When the boot order comes to the iSCSI Boot Firmware Table (iBFT) adapter, the adapter attempts to connect to the target, but does not boot from it. See [Prerequisites](#).

If the connection to the iSCSI target is successful, the iSCSI boot firmware saves the iSCSI boot configuration in the iBFT. The next adapter to boot must be the ESXi installation media, either a mounted ISO image or a physical CD-ROM.

1. Start an interactive installation from the ESXi installation CD/DVD or mounted ISO image.
2. On the Select a Disk screen, select the iSCSI target you specified in the iBFT BIOS target parameter setting.

If the target does not appear in this menu, make sure that the TCP/IP and initiator iSCSI IQN settings are correct. Check the network Access Control List (ACL) and confirm that the adapter has adequate permissions to access the target.

3. Follow the prompts to complete the installation.
4. Reboot the host.
5. In the host BIOS settings, enter the iBFT adapter BIOS configuration, and change the adapter parameter to boot from the iSCSI target.

See the vendor documentation for your system.

On your iBFT adapter, reenable the option to boot to the iSCSI target, so the system will boot from the LUN you installed ESXi on.

Installing ESXi by Using a Script

The installation/upgrade script is a text file, for example `ks.cfg`, that contains supported commands.

The command section of the script contains the ESXi installation options. This section is required and must appear first in the script.


Scripted ESXi Installation

Running a script is an efficient way to deploy multiple ESXi hosts with an unattended installation.

The installation script contains the host configuration settings. You can use the script to configure multiple hosts with the same settings. See [Installing or Upgrading Hosts by Using a Script](#).

The installation script must be stored in a location that the host can access by HTTP, HTTPS, FTP, NFS, CDROM, or USB. You can PXE boot the ESXi installer or boot it from a CD/DVD or USB drive.

Figure 14: Scripted Installation



Approaches for Scripted Installation

You can install ESXi on multiple machines using a single script for all of them or a separate script for each machine.

For example, because disk names vary from machine to machine, one of the settings that you might want to configure in a script is the selection for the disk to install ESXi on.

Table 13: Scripted Installation Choices

Option	Action
Always install on the first disk on multiple machines.	Create one script.
Install ESXi on a different disk for each machine.	Create multiple scripts.

For information about the commands required to specify the disk to install on, see [Installation and Upgrade Script Commands](#).

Enter Boot Options to Run an Installation or Upgrade Script

You can start an installation or upgrade script by typing boot options at the ESXi installer boot command line.

At boot time you might need to specify options to access the kickstart file. You can enter boot options by pressing Shift+O in the boot loader. For a PXE boot installation, you can pass options through the `kernelopts` line of the `boot.cfg` file. See [About the boot.cfg File](#) and [Network Booting the ESXi Installer](#).

To specify the location of the installation script, set the `ks=filepath` option, where `filepath` indicates the location of your kickstart file. Otherwise, a scripted installation or upgrade cannot start. If `ks=filepath` is omitted, the text installer is run.

Supported boot options are listed in [Boot Options](#).

1. Start the host.
2. When the ESXi installer window appears, press Shift+O to edit boot options.



3. At the `runweasel` command prompt, type `ks=location of installation script plus boot command-line options`.

Boot Option

You type the following boot options:

```
ks=http://00.00.00.00/kickstart/ks-osdc-pdp101.cfg nameserver=00.00.0.0 ip=00.00.00.000 net-
mask=255.255.255.0 gateway=00.00.00.000
```

Boot Options

When you perform a scripted installation, you might need to specify options at boot time to access the kickstart file.

Supported Boot Options

Table 14: Boot Options for ESXi Installation

Boot Option	Description
<code>BOOTIF=hwtype-MAC address</code>	Similar to the <code>netdevice</code> option, except in the PXELINUX format as described in the IPAPPEND option under SYSLINUX at the syslinux.org site.
<code>gateway=ip address</code>	Sets this network gateway as the default gateway to be used for downloading the installation script and installation media.
<code>ip=ip address</code>	Sets up a static IP address to be used for downloading the installation script and the installation media. Note: the PXELINUX format for this option is also supported. See the IPAPPEND option under SYSLINUX at the syslinux.org site.
<code>ks=cdrom:/path</code>	Performs a scripted installation with the script at <code>path</code> , which resides on the CD in the CD-ROM drive. Each CDROM is mounted and checked until the file that matches the path is found. Important: If you have created an installer ISO image with a custom installation or upgrade script, you must use uppercase characters to provide the path of the script, for example, <code>ks=cdrom:/KS_CUST.CFG</code> .
<code>ks=file://path</code>	Performs a scripted installation with the script at <code>path</code> .
<code>ks=protocol://serverpath</code>	Performs a scripted installation with a script located on the network at the given URL. <code>protocol</code> can be <code>http</code> , <code>https</code> , <code>ftp</code> , or <code>nfs</code> . An example using NFS protocol is <code>ks=nfs://host/porturl-path</code> . The format of an NFS URL is specified in RFC 2224.
<code>ks=usb</code>	Performs a scripted installation, accessing the script from an attached USB drive. Searches for a file named <code>ks.cfg</code> . The file must be located in the root directory of the drive. If multiple USB flash drives are attached, they are searched until the <code>ks.cfg</code> file is found. Only FAT16 and FAT32 file systems are supported.
<code>ks=usb:/path</code>	Performs a scripted installation with the script file at the specified path, which resides on USB.
<code>ksdevice=device</code>	Tries to use a network adapter <code>device</code> when looking for an installation script and installation media. Specify as a MAC address, for example, <code>00:50:56:C0:00:01</code> . This location can also be a <code>vmnicNN</code> name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.
<code>nameserver=ip address</code>	Specifies a domain name server to be used for downloading the installation script and installation media.
<code>netdevice=device</code>	Tries to use a network adapter <code>device</code> when looking for an installation script and installation media. Specify as a MAC address, for example, <code>00:50:56:C0:00:01</code> . This location can also be a <code>vmnicNN</code> name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.

Boot Option	Description
<code>netmask=subnet mask</code>	Specifies subnet mask for the network interface that downloads the installation script and the installation media.
<code>vlanid=vlanid</code>	Configure the network card to be on the specified VLAN.
<code>systemMediaSize=small</code>	Limits the size of system storage partitions on the boot media. The selected value must fit the purpose of your system. You can select from the following values: <code>min</code> (32 GB, for single disk or embedded servers) <code>small</code> (64 GB, for servers with at least 512 GB RAM) <code>default</code> (128 GB) <code>max</code> (consume all available space, for multi-terabyte servers)

For more information on ESXi booting options post installation, see VMware knowledge base article [77009](#).

Installation and Upgrade Scripts Used for ESXi Installation

You can use a default script to perform a standard ESXi installation to the first detected disk.

During an ESXi installation, you can use a default script to perform a standard installation to the first detected disk, and the boot loader configuration file to specify the kernel, the kernel options, and the boot modules that the mboot.c32 or mboot.efi boot loader uses for the installation.

About the Default ks.cfg Installation Script

The ESXi installer includes a default installation script that performs a standard installation to the first detected disk.

The default `ks.cfg` installation script is located in the initial RAM disk at `/etc/vmware/weasel/ks.cfg`. You can specify the location of the default `ks.cfg` file with the `ks=file:///etc/vmware/weasel/ks.cfg` boot option. See [Enter Boot Options to Run an Installation or Upgrade Script](#).

When you install ESXi using the `ks.cfg` script, the default root password is `myp@ssw0rd`.

You cannot modify the default script on the installation media. After the installation, you can use the vSphere Client to log in to the vCenter Server that manages the ESXi host and modify the default settings.

With vSphere 8.0, if your system has data processing units (DPU), you use the `ks.cfg` script also to install ESXi on DPUs.

The default script contains the following commands:

```
#
# Sample scripted installation file
#

# Accept the VMware End User License Agreement
vmaccepteula

# Set the root password for the DCUI and Tech Support Mode
rootpw myp@ssw0rd

# Install on the first local disk available on machine
install --firstdisk --overwritevmfs
  In case you system has DPUs, you also specify a PCI slot:
install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>
```

```
# Set the network to DHCP on the first network adapter
network --bootproto=dhcp --device=vmnic0

# A sample post-install script
%post --interpreter=python --ignorefailure=true
import time
stampFile = open('/finished.stamp', mode='w')
stampFile.write( time.asctime() )
```

About the boot.cfg File

The boot loader configuration file `boot.cfg` specifies the kernel, the kernel options, and the boot modules that the `mboot.c32` or `mboot.efi` boot loader uses in an ESXi installation.

The `boot.cfg` file is provided in the ESXi installer. You can modify the `kernelopt` line of the `boot.cfg` file to specify the location of an installation script or to pass other boot options.

The `boot.cfg` file has the following syntax:

```
# boot.cfg -- mboot configuration file
#
# Any line preceded with '#' is a comment.

title=STRING
prefix=DIRPATH
kernel=FILEPATH
kernelopt=STRING
modules=FILEPATH1---FILEPATH2... ---FILEPATHn

# Any other line must remain unchanged.
```

The commands in `boot.cfg` configure the boot loader.

Table 15: Commands in boot.cfg.

Command	Description
<code>title=STRING</code>	Sets the boot loader title to <i>STRING</i> .
<code>prefix=STRING</code>	(Optional) Adds <i>DIRPATH</i> / in front of every <i>FILEPATH</i> in the <code>kernel=</code> and <code>modules=</code> commands that do not already start with / or with <code>http://</code> .
<code>kernel=FILEPATH</code>	Sets the kernel path to <i>FILEPATH</i> .
<code>kernelopt=STRING</code>	Appends <i>STRING</i> to the kernel boot options.
<code>modules=FILEPATH1 --- FILEPATH2... --- FILEPATHn</code>	Lists the modules to be loaded, separated by three hyphens (<code>---</code>).

See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#) and [Network Booting the ESXi Installer](#).

Locations Supported for Installation or Upgrade Scripts

In scripted installations and upgrades, the ESXi installer can access the installation or upgrade script, also called the kickstart file, from several locations.

The following locations are supported for the installation or upgrade script:

- CD/DVD. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).
- USB Flash drive. See [Create a USB Flash Drive to Store the ESXi Installation Script or Upgrade Script](#).
- A network location accessible through the following protocols: NFS, HTTP, HTTPS, FTP

Path to the Installation or Upgrade Script

You can specify the path to an installation or upgrade script.

`ks=http://XXX.XXX.XXX.XXX/kickstart/KS.CFG` is the path to the ESXi installation script, where `XXX.XXX.XXX.XXX` is the IP address of the machine where the script resides. See [Installing ESXi by Using a Script](#).

To start an installation script from an interactive installation, you enter the `ks=` option manually. See [Enter Boot Options to Run an Installation or Upgrade Script](#).

Installation and Upgrade Script Commands

To modify the default installation or upgrade script or to create your own script, use supported commands. Use supported commands in the installation script, which you specify with a boot command when you boot the installer.

To determine which disk to install or upgrade ESXi on, the installation script requires one of the following commands: `install`, `upgrade`, or `installorupgrade`. The `install` command creates the default partitions, including a VMFS datastore that occupies all available space after the other partitions are created.

With vSphere 8.0, if your system has supported data processing units (DPU), always consider the installation, re-installation or upgrade of ESXi on the DPUs along with ESXi on hosts. ESXi update and upgrade on DPUs is not supported by the interactive or scripted method, you can only use vSphere Lifecycle Manager.

Note: The use of SD and USB devices for storing ESX-OSData partitions is being deprecated. You can use SD and USB devices only to create boot bank partitions, `boot-bank 0` and `boot-bank 1`. Additionally, you can provide a persistent disk of minimum 32 GB on which to install the ESX-OSData partition. You define such disks by using the parameter `systemDisk` in the `install` command.

accepteula or vmaccepteula (Required)

Accepts the ESXi license agreement.

clearpart (Optional)

Clears any existing partitions on the disk. Requires the `install` command to be specified. Carefully edit the `clearpart` command in your existing scripts.

--drives=

Remove partitions on the specified drives.

--alldrives

Ignores the `--drives=` requirement and allows clearing of partitions on every drive.

--ignoredrives=

Removes partitions on all drives except those specified. Required unless the `--drives=` or `--alldrives` flag is specified.

--overwritevmfs

Allows overwriting of VMFS partitions on the specified drives. By default, overwriting VMFS partitions is not allowed.

--firstdisk=
disk-type1
 [*disk-type2*,...]

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot: : `install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>` . For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (*local*)
2. Network storage (*remote*)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including *esx* for the first disk with ESXi installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the *mptsas* driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use *localesx* for local storage that contains ESXi image or *remoteesx* for remote storage that contains ESXi image.

dryrun (Optional)

Parses and checks the installation script. Does not perform the installation.

install

Specifies that this is a fresh installation. Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

--disk= or --drive=

Specifies the disk to partition. In the command `--disk=diskname` , the *diskname* can be a disk name or a full disk filesystem path in ESXi, for example:

- Disk name: `--disk=naa.6d09466044143600247aee55ca2a6405` or
- Device path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`

For accepted disk name formats, see [Disk Device Names](#).

--firstdisk=
disk-type1,
 [*disk-type2*,...]

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot: : `install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>` . For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (*local*)
2. Network storage (*remote*)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including *esx* for the first disk with ESXi installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the *mptsas* driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use *localesx* for local storage that contains ESXi image or *remoteesx* for remote storage that contains ESXi image.

--ignoressd

Excludes solid-state disks from eligibility for partitioning. This option can be used with the `install` command and the `--firstdisk` option. This option takes precedence over the `--firstdisk` option. This option is invalid with the `--drive` or `--disk` options and with the `upgrade` and `installorupgrade` commands. See the *vSphere Storage* documentation for more information about preventing SSD formatting during auto-partitioning.

--overwritevsan

You must use the `--overwritevsan` option when you install ESXi on a disk, either SSD or HDD (magnetic), that is in a vSAN disk group. If you use this option and no vSAN partition is on the selected disk, the installation fails. When you install ESXi on a disk that is in vSAN disk group, the result depends on the disk that you select:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group is wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD is wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD is wiped.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

--overwritevmfs

Required to overwrite an existing VMFS datastore on the disk before installation.

--preservevmfs

Preserves an existing VMFS datastore on the disk during installation.

--novmfsondisk

Prevents a VMFS partition from being created on this disk. Must be used with `--overwritevmfs` if a VMFS partition exists on the disk.

--systemdisk

If you use an USB or SD device, `systemDisk` specifies local persistent disk on which to install the ESX-OSData partition. For example, `install --firstdisk = usb --systemDisk=<diskID>`. As a result, boot bank partitions are placed on the USB device, while the OSData partition is on the disk specified in the `systemDisk` parameter.

--repartitionssystemdisk

If you use an USB or SD device and the local disk that you specify with the `systemDisk` parameter is not empty or contains a datastore, you can use `repartitionSystemDisk` to make sure that the persistent disk is repartitioned before use.

Note: If a local persistent disk is not available or the disk size is less than 32GB, you see warning messages, but installation continues.

--forceunsupportedinstall

Blocks the installation of deprecated CPUs.

installorupgrade

Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

--disk= or --drive=

Specifies the disk to partition. In the command `--disk=diskname`, the `diskname` can be a disk name or a full disk filesystem path in ESXi, for example:

- Disk name: `--disk=naa.6d09466044143600247aee55ca2a6405` or
- Device path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`

For accepted disk name formats, see [Disk Device Names](#).

--firstdisk=
disk-type1,
[disk-type2,...]

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot: :
`install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>` . For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (`local`)
2. Network storage (`remote`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESX installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the `mptsas` driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `localeSX` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

--overwritevsan

You must use the `--overwritevsan` option when you install ESXi on a disk, either SSD or HDD (magnetic), that is in a vSAN disk group. If you use this option and no vSAN partition is on the selected disk, the installation fails. When you install ESXi on a disk that is in a vSAN disk group, the result depends on the disk that you select:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group is wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD is wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD is wiped.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

--overwritevmfs

Install ESXi if a VMFS partition exists on the disk, but no ESX or ESXi installation exists. Unless this option is present, the installer fails if a VMFS partition exists on the disk, but an ESX or ESXi installation is missing.

keyboard (Optional)

Sets the keyboard type for the system.

keyboardType

Specifies the keyboard map for the selected keyboard type. *keyboardType* must be one of the following types.

- Belgian
- Brazilian
- Croatian
- Czechoslovakian
- Danish
- Estonian
- Finnish
- French
- German
- Greek
- Icelandic
- Italian
- Japanese

- Latin American
- Norwegian
- Polish
- Portuguese
- Russian
- Slovenian
- Spanish
- Swedish
- Swiss French
- Swiss German
- Turkish
- Ukrainian
- United Kingdom
- US Default
- US Dvorak

serialnum or vmserialnum (Optional)

The command is supported in ESXi version 5.1 and later. Configures licensing. If not included, ESXi installs in evaluation mode.

--esx=<license-key>

Specifies the vSphere license key to use. The format is 5 five-character groups (XXXXX-XXXXX-XXXXX-XXXXX-XXXXX).

network (Optional)

Specifies a network address for the system.

--bootproto=[dhcp|static]

Specifies whether to obtain the network settings from DHCP or set them manually.

--device=

Specifies either the MAC address of the network card or the device name, in the form `vmnicNN`, as in `vmnic0`. This option refers to the uplink device for the virtual switch.

--ip=

Sets an IP address for the machine to be installed, in the form `xxx.xxx.xxx.xxx`. Required with the `--bootproto=static` option and ignored otherwise.

--gateway=

Designates the default gateway as an IP address, in the form `xxx.xxx.xxx.xxx`. Used with the `--bootproto=static` option.

--nameserver=

Designates the primary name server as an IP address. Used with the `--bootproto=static` option. Omit this option if you do not intend to use DNS.

The `--nameserver` option can accept two IP addresses. For example: `--nameserver="10.126.87.104[,10.126.87.120]"`

--netmask=

Specifies the subnet mask for the installed system, in the form `255.xxx.xxx.xxx`. Used with the `--bootproto=static` option.

--hostname=

Specifies the host name for the installed system.

--vlanid= *vlanid*

Specifies which VLAN the system is on. Used with either the `--bootproto=dhcp` or `--bootproto=static` option. Set to an integer from 1 to 4096.

--addvmportgroup=(0|1)

Specifies whether to add the VM Network port group, which is used by virtual machines. The default value is 1.

paranoid (Optional)

Causes warning messages to interrupt the installation. If you omit this command, warning messages are logged.

part or partition (Optional)

Creates an extra VMFS datastore on the system. Only one datastore per disk can be created. Cannot be used on the same disk as the `install` command. Only one partition can be specified per disk and it can only be a VMFS partition.

datastore name

Specifies where the partition is to be mounted.

--ondisk= or --ondrive=

Specifies the disk or drive where the partition is created.

--onfirstdisk=

disk-type1,

[*disk-type2*,...]

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot: : `install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>`. For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (*local*)
2. Network storage (*remote*)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including *esx* for the first disk with ESX installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the *mptsas* driver rather than a normal local disk, the argument is `--onfirstdisk=ST3120814A,mptsas,local`. You can use *localesx* for local storage that contains ESXi image or *remoteesx* for remote storage that contains ESXi image.

reboot (Optional)

Reboots the machine after the scripted installation is complete.

<--noeject>

The CD is not ejected after the installation.

rootpw (Required)

Sets the root password for the system.

--iscrypted

Specifies that the password is encrypted.

password

Specifies the password value.

upgrade

Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

--disk= or --drive=

Specifies the disk to partition. In the command `--disk=diskname`, the *diskname* can be a disk name or a full disk filesystem path in ESXi, for example:

- Disk name: `--disk=naa.6d09466044143600247aee55ca2a6405` or
- Device path: `--disk=/vmfs/devices/disks/mpx.vmhbal:C0:T0:L0`

For accepted disk name formats, see [Disk Device Names](#).

--firstdisk=

disk-type1,

[*disk-type2*,...]

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (`local`)
2. Network storage (`remote`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESX installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the `mptsas` driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `localesx` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

%include or include (Optional)

Specifies another installation script to parse. This command is treated similarly to a multiline command, but takes only one argument.

filename

For example: `%include part.cfg`

%pre (Optional)

Specifies a script to run before the kickstart configuration is evaluated. For example, you can use it to generate files for the kickstart file to include.

--interpreter

=[python|busybox]

Specifies an interpreter to use. The default is `busybox`.

%post (Optional)

Runs the specified script after package installation is complete. If you specify multiple `%post` sections, they run in the order that they appear in the installation script.

--interpreter

=[python|busybox]

Specifies an interpreter to use. The default is `busybox`.

--timeout=secs

Specifies a timeout for running the script. If the script is not finished when the timeout expires, the script is forcefully stopped.

--ignorefailure

=[true|false]

If true, the installation is considered a success even if the `%post` script stops with an error.

%firstboot

Creates an `init` script that runs only during the first boot. The script has no effect on subsequent boots. If multiple `%firstboot` sections are specified, they run in the order that they appear in the kickstart file.

Note: You cannot check the semantics of `%firstboot` scripts until the system is booting for the first time. A `%firstboot` script might contain potentially catastrophic errors that are not exposed until after the installation is complete.

Important: The `%firstboot` script does not run, if secure boot is enabled on the ESXi host.

--interpreter

=[python|busybox]

Specifies an interpreter to use. The default is `busybox`.

Note: You cannot check the semantics of the `%firstboot` script until the system boots for the first time. If the script contains errors, they are not exposed until after the installation is complete.

Disk Device Names

The `install`, `upgrade`, and `installorupgrade` installation script commands require the use of disk device names.

Table 16: Disk Device Names

Format	Example	Description
NAA	naa.6d09466044143600247aee55ca2a6405	SCSI INQUIRY identifier
EUI	eui.3966623838646463	SCSI INQUIRY identifier
T10	t10.SanDisk00Cruzer_Blade000000004C530001171118101244	SCSI INQUIRY identifier
VML	vml.00025261	Legacy VMkernel identifier
MPX	mpx.vmhba0:C0:T0:L0	Path-based identifier

For more information on storage device names, see *Storage Device Names and Identifiers* in the *vSphere Storage* documentation.

Configuring External Entropy Sources During Scripted Installation

Starting with ESXi 8.0 Update 1, you can configure external entropy sources in the kickstart file for scripted installation.

You can configure ESXi in a highly secure environment to consume entropy from external entropy sources, such as a Hardware Security Module (HSM), and align with standards such as BSI Common criteria, EAL4, and NIST FIPS CMVP, by using the scripted installation method.

ESXi 8.0 Update 1 introduces an entropy daemon, `entropyd`, that creates a vAPI endpoint to provide REST API to query and send entropy data to ESXi hosts. The entropy daemon has several configurable parameters: `in-memory-cache size`, `in-storage-cache size`, `in-memory-low watermark`, and `entropy-lost-timeout`. Entropy sources,

internal and external, are collected in the entropy-mixer module. The entropy daemon forwards entropy sources from the mixer to the entropy pool of the kernel.

If you do not need to select external entropy sources, you do not need to change anything in your existing scripts.

You can configure external entropy sources only for new installations. The entropy daemon uses entropy data passed during installation at first boot.

Note: You cannot use any other method, such as interactive install, a cluster image, or Auto Deploy, to configure external entropy sources. You cannot enable the feature on an existing ESXi host.

Before you update the kickstart file, you must get legitimate binary entropy data from an external source such as HSM and save it as a file of size between 512 KB to 10 MB, for example, `entropy_data.dat`. You then encode the contents of the `entropy_data` file in a valid base64 format to create another file, for example `entropy_data.b64`.


In the kickstart file, you provide the following new parameters as part of the `entropy` command:

- `data` - the base64-encoded entropy data from an external source.
- `sources` - This bitmask value (0=default, 1=interrupts, 2=RD RAND, 4=RDSEED, 8=entropyd.) is set in the kernel settings during install time. If RDSEED is supported, the default is FIPS compliance. Otherwise, the default is all entropy sources, except `entropyd`. For EAL4 compliance, 8 is the only required value, but you can use other sources in the decimal range 0–15. Source values persist as a kernel setting during install.

Note: If you configure an ESXi host with external entropy sources only, that is when the `sources` value is set to 8, you must keep supplying external entropy to the host by using the entropy API. If external entropy is exhausted in the host, the host becomes unresponsive and it might require a hard reboot or re-installation to recover the host from such a situation.

A sample kickstart with the entropy parameters:

```
vmaccepteula
rootpw xxxxxxxx
entropy --sources=8 --data=xxxxxxxx/xx/xxxxxx/xx/xxxx...
install --firstdisk --overwritevmfs
network --bootproto=dhcp
```

Figure 15: Scripted installation workflow for adding external entropy sources

After installation completes, you can log in to the ESXi host and define some parameters for the entropy daemon from the shell by using the following ESXCLI commands:

ESXCLI commands

1. Get Commands

##	esxcli system entropyd get	Command Description
1	<no argument>	Get currently configured and default values of all entropyd parameters
2	--default-values	Get default values

2. Set Commands

##	esxcli system entropyd set	Command Description
1	--help	Print details of "esxcli system entropyd set" command and it's arguments
2	--reset=all --reset=memory-cache-size --reset=memory-cache-low-watermark --reset=storage-cache-size --reset=external-entropy-lost-timeout	Reset a parameter or all parameters to its default values.
3	--memory-cache-size= <value in KiB>	Set memory cache size in KiB.
4	--memory-cache-low-watermark= <value in %>	Set memory cache low water mark in percentage.
5	--storage-cache-size=<value in KiB>	Set storage cache size in KiB.
6	--external-entropy-lost-timeout= <value in seconds>	Set external entropy lost timeout in seconds.

Example:

```
$ esxcli system entropyd set --external-entropy-lost-timeout=70 --memory-cache-low-watermark=30 --memory-cache-size=612 --storage-cache-size=5096
$ esxcli system entropyd get
External Entropy Lost Timeout Seconds: 70
Memory Cache Low Watermark Percentage: 30
Memory Cache Size Kibibytes: 612
Storage Cache Size Kibibytes: 5096

$ esxcli system entropyd get --default-values
External Entropy Lost Timeout Seconds: 60
Memory Cache Low Watermark Percentage: 20
Memory Cache Size Kibibytes: 512
Storage Cache Size Kibibytes: 4096
```

During installation, the ESXi installer updates the entropy sources value in the kernel settings so that it persists in the ConfigStore after the installation. This change does not affect secure booting of ESXi hosts.

For more information, see vSphere Security and vSphere Automation SDKs Programming guides.

Install or Upgrade ESXi from a CD or DVD by Using a Script


You can install or upgrade ESXi from a CD-ROM or DVD-ROM drive by using a script that specifies the installation or upgrade options.

Before you run the scripted installation or upgrade, verify that the following prerequisites are met:

- The system on which you are installing or upgrading meets the hardware requirements. See [ESXi Hardware Requirements](#).
- You have the ESXi installer ISO on an installation CD or DVD . See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#).
- The default installation or upgrade script (`ks.cfg`) or a custom installation or upgrade script is accessible to the system. See [Installing ESXi by Using a Script](#).
- You have selected a boot command to run the scripted installation or upgrade. See [Enter Boot Options to Run an Installation or Upgrade Script](#). For a complete list of boot commands, see [Boot Options](#).

You can start the installation or upgrade script by entering a boot option when you start the host. You can also create an installer ISO image that includes the installation script. With an installer ISO image, you can perform a scripted, unattended installation when you boot the resulting installer ISO image. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).

1. Boot the ESXi installer from the local CD-ROM or DVD-ROM drive.
2. When the ESXi installer window appears, press Shift+O to edit boot options.



3. Type a boot option that calls the default installation or upgrade script or an installation or upgrade script file that you created.

The boot option has the form `ks=` .

4. Press Enter.

The installation, upgrade, or migration runs, using the options that you specified.

Install or Upgrade ESXi from a USB Stick by Using a Script

You can install or upgrade ESXi from a USB flash drive by using a script that specifies the installation or upgrade options.


Before running the scripted installation or upgrade, verify that the following prerequisites are met:

- The system that you are installing or upgrading to ESXi meets the hardware requirements for the installation or upgrade. See [ESXi Hardware Requirements](#).
- You have the ESXi installer ISO on a bootable USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).
- The default installation or upgrade script (`ks.cfg`) or a custom installation or upgrade script is accessible to the system. See [Installing ESXi by Using a Script](#).

- You have selected a boot option to run the scripted installation, upgrade, or migration. See [Enter Boot Options to Run an Installation or Upgrade Script](#).

Supported boot options are listed in [Boot Options](#).

1. Boot the ESXi installer from the USB flash drive.
2. When the ESXi installer window appears, press Shift+O to edit boot options.



3. Type a boot option that calls the default installation or upgrade script or an installation or upgrade script file that you created.

The boot option has the form `ks=`.

4. Press Enter.

The installation, upgrade, or migration runs, using the options that you specified.

Network Boot the ESXi Installer for a Scripted Installation or Upgrade

ESXi8.0 provides many options for booting the installer over a network and using an installation or upgrade script.

- For information about setting up a network infrastructure, see [Network Booting the ESXi Installer](#).
- For information about creating and locating an installation script, see [Installing ESXi by Using a Script](#).
- For specific procedures to network boot the ESXi installer and use an installation script, see one of the following topics:
 - [Boot the ESXi Installer by Using Native UEFI HTTP](#)
 - [Boot the ESXi Installer by Using iPXE and HTTP](#)
 - [Boot the ESXi Installer by Using PXE and TFTP](#)

How to Boot an ESXi Host from a Network Device

Network Booting the ESXi Installer

You can use preboot execution environment (PXE) to boot an ESXi host from a network device, if your host uses legacy BIOS or UEFI.

Alternatively, if your ESXi host supports native UEFI HTTP, you can use hypertext transfer protocol (HTTP) to boot the host from a network device. ESXi is distributed in an ISO format that is used to install to flash memory or to a local hard drive. You can extract the files and boot them over a network interface.

PXE uses Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP) to boot an operating system over a network.

PXE booting requires some network infrastructure and a machine with a PXE-capable network adapter. Most machines that can run ESXi have network adapters that can PXE boot.

Native UEFI HTTP uses DHCP and HTTP to boot over a network. UEFI HTTP boot requires a network infrastructure, UEFI firmware version on the ESXi host that includes HTTP boot feature, and a network adapter that supports UEFI networking.


Booting by using HTTP is faster and more reliable than using TFTP. This is due to the capabilities of the TCP protocol that underlies the HTTP, such as built-in streaming and lost packet recovery. If your ESXi hosts do not support native UEFI HTTP, you can use iPXE HTTP for the boot process.

Note: Network booting with legacy BIOS firmware is possible only over IPv4. Network booting with UEFI BIOS firmware is possible over IPv4 or IPv6.

Overview of the Network Boot Installation Process

You can boot an ESXi host from a network interface. The network boot process varies depending on whether the target host is using legacy BIOS or UEFI firmware, and whether the boot process uses PXE TFTP, iPXE HTTP, or UEFI HTTP.

When you boot a target host, it interacts with the different servers in the environment to get a network adapter, boot loader, kernel, IP address for the kernel, and finally an installation script. When all components are in place, installation starts, as shown in the following illustration.

Figure 16: Overview of PXE Boot Installation Process

The interaction between the ESXi host and other servers proceeds as follows:

1. The user boots the target ESXi host.
2. The target ESXi host makes a DHCP request.
3. The DHCP server responds with the IP information, the location of the TFTP or HTTP server, and the filename or URL of the initial network boot loader.
4. The ESXi host contacts the TFTP or HTTP server and requests the filename or URL that the DHCP server specified.
5. The TFTP or HTTP server sends the network boot loader, and the ESXi host runs it. The initial boot loader might load additional boot loader components from the server.
6. The boot loader searches for a configuration file on the TFTP or HTTP server, downloads the kernel and other ESXi components as specified in the configuration file, and boots the kernel on the ESXi host.
7. The installer runs interactively or using a kickstart script, as specified in the configuration file.

Network Boot Background Information

Understanding the network boot process can help you during troubleshooting.

TFTP Server

Trivial File Transfer Protocol (TFTP) is similar to the FTP service, and is typically used only for network booting systems or loading firmware on network devices such as routers. TFTP is available on Linux and Windows.

- Most Linux distributions include a copy of the `tftp-hpa` server. If you require a supported solution, purchase a supported TFTP server from your vendor of choice. You can also acquire a TFTP server from one of the packaged appliances on the VMware Marketplace.
- If your TFTP server runs on a Microsoft Windows host, use `tfptd32` version 2.11 or later. See <http://tftpd32.jounin.net/>.

SYSLINUX and PXELINUX

If you are using PXE in a legacy BIOS environment, you must understand the different boot environments.

- SYSLINUX is an open-source boot environment for machines that run legacy BIOS firmware. The ESXi boot loader for BIOS systems, `mboot.c32`, runs as a SYSLINUX plugin. You can configure SYSLINUX to boot from several types of media, including disk, ISO image, and network. You can find the SYSLINUX package at <http://www.kernel.org/pub/linux/utils/boot/syslinux/>.
- PXELINUX is a SYSLINUX configuration for booting from a TFTP server according to the PXE standard. If you use PXELINUX to boot the ESXi installer, the `pxelinux.0` binary file, `mboot.c32`, the configuration file, the kernel, and other files are transferred by TFTP.

Note: VMware builds the `mboot.c32` plugin to work with SYSLINUX version 3.86 and tests PXE booting only with that version. Other versions might be incompatible. *The Open Source Disclosure Package for VMware vSphere Hypervisor* includes bug fixes for SYSLINUX version 3.86.

iPXE

iPXE is open-source software that provides an implementation of HTTP. You can use the software to perform an initial boot. For more information, see <https://ipxe.org/>.

VMware includes a build of iPXE as part of Auto Deploy. The source tree for this build is available in *The Open Source Disclosure Package for VMware vCenter Server*.

UEFI PXE and UEFI HTTP

Most UEFI firmware natively includes PXE support that allows booting from a TFTP server. The firmware can directly load the ESXi boot loader for UEFI systems, `mboot.efi`. Additional software such as PXELINUX is not required.

Some UEFI firmware support native UEFI HTTP boot. The feature is introduced in version 2.5 of the UEFI specification. The firmware can load the ESXi boot loader from an HTTP server, without additional software, such as iPXE.

Note: Apple Macintosh products do not include PXE boot support. They include support for network booting through an Apple-specific protocol instead.

Alternative Approaches to Network Booting

Alternative approaches to network booting different software on different hosts are also possible, for example:

- Configuring the DHCP server to provide different initial boot loader filenames to different hosts depending on MAC address or other criteria. See your DHCP server's documentation.
- Approaches using iPXE as the initial bootloader with an iPXE configuration file that selects the next bootloader based on the MAC address or other criteria.

PXELINUX Configuration Files

You need a PXELINUX configuration file to boot the ESXi installer on a legacy BIOS system. The configuration file defines the menu displayed to the target ESXi host as it starts.

This section gives general information about PXELINUX configuration files.

For syntax details, see the SYSLINUX website at <http://www.syslinux.org/>.

Required Files

In the PXE configuration file, you must include paths to the following files:

- `mboot.c32` is the boot loader.
- `boot.cfg` is the boot loader configuration file.

See [About the boot.cfg File](#)

Filename for the PXE Configuration File

For the filename of the PXE configuration file, select one of the following options:

- `01-mac_address_of_target_ESXi_host`. For example, `01-23-45-67-89-0a-bc`
- The target ESXi host IP address in a hexadecimal notation.
- `default`

The initial boot file `pxelinux.0` tries to load a PXE configuration file in the following order:

1. It tries with the MAC address of the target ESXi host, prefixed with its ARP type code, which is 01 for Ethernet.
2. If that attempt fails, it tries with the hexadecimal notation of target ESXi system IP address.
3. Ultimately, it tries to load a file named `default`.

File Location for the PXE Configuration File

Save the file in `/tftpboot/pxelinux.cfg/` on the TFTP server.

For example, you might save the file on the TFTP server at `/tftpboot/pxelinux.cfg/01-00-21-5a-ce-40-f6`. The MAC address of the network adapter on the target ESXi host is `00-21-5a-ce-40-f6`.

Boot the ESXi Installer by Using PXE and TFTP

You can use a TFTP server to PXE boot the ESXi installer. The process differs slightly depending on whether you use UEFI or boot from a legacy BIOS.

Because most environments include ESXi hosts that support UEFI boot and hosts that support only legacy BIOS, this topic discusses prerequisites and steps for both types of hosts.

Verify that your environment meets the following prerequisites.

- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with PXE support on the target ESXi host.
- DHCP server that you can configure for PXE booting. See [Sample DHCP Configurations](#).
- TFTP server.
- Network security policies to allow TFTP traffic (UDP port 69).
- For legacy BIOS, you can use only IPv4 networking. For UEFI PXE boot, you can use IPv4 or IPv6 networking.
- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

For legacy BIOS systems, obtain version 3.86 of the SYSLINUX package. For more information, see [Network Boot Background Information](#).

- For legacy BIOS machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `pxelinux.0` initial boot loader for all target machines, but potentially different PXELINUX configuration files depending on the target machine's MAC address.
- For UEFI machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

1. If your ESXi host runs legacy BIOS firmware only, obtain and configure PXELINUX.

- Obtain SYSLINUX version 3.86, unpack it, and copy the `pxelinux.0` file to the top-level `/tftpboot` directory on your TFTP server.
- Create a PXELINUX configuration file using the following code model.

`ESXi-8.x.x-XXXXXX` is the name of the TFTP subdirectory that contains the ESXi installer files.

```
DEFAULT install
NOHALT 1
LABEL install
  KERNEL ESXi-8.x.x-XXXXXX/mboot.c32
  APPEND -c ESXi-8.x.x-XXXXXX/boot.cfg
  IPAPPEND 2
```

- Save the PXELINUX file in the `/tftpboot/pxelinux.cfg` directory on your TFTP server with a filename that will determine whether all hosts boot this installer by default:

Option	Description
Same installer	Name the file <code>default</code> if you want all host to boot this ESXi installer by default.
Different installers	Name the file with the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>) if you want only a specific host to boot with this file, for example, <code>01-23-45-67-89-0a-bc</code> .

2. If your ESXi host runs UEFI firmware, copy the `efi/boot/bootx64.efi` and `efi/boot/crypto64.efi` files from the ESXi installer ISO image to the `/tftpboot` folder on your TFTP server.

3. Rename the `efi/boot/bootx64.efi` file to `mboot.efi`.

Note: Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.

4. Configure the DHCP server.

5. Create a subdirectory of your TFTP server's top-level `/tftpboot` directory and name it after the version of ESXi it will hold, for example, `/tftpboot/ESXi-8.x.x-xxxxx`.

6. Copy the contents of the ESXi installer image to the newly created directory.

7. Modify the `boot.cfg` file

- Add the following line:

```
prefix=ESXi-7.x.x-xxxxxx
```

Here, `ESXi-7.x.x-xxxxxx` is the pathname of the installer files relative to the TFTP server's root directory.

- b) If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (/) character, delete that character.
- c) If the `kernelopt=` line contains the string `cdromBoot`, remove the string only.

8. Optional: For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.

Use the following code as a model, where `XXX.XXX.XXX.XXX` is the IP address of the server where the installation script resides, and `esxi_ksFiles` is the directory that contains the `ks.cfg` file.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```

9. If your ESXi host runs UEFI firmware, specify whether you want all UEFI hosts to boot the same installer.

Option	Description
Same installer	Copy or link the <code>boot.cfg</code> file to <code>/tftpboot/boot.cfg</code>
Different installers	1. Create a subdirectory of <code>/tftpboot</code> named after the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), for example, <code>01-23-45-67-89-0a-bc</code>. 2. Place a copy of (or a link to) the host's <code>boot.cfg</code> file in that directory, for example, <code>/tftpboot/01-23-45-67-89-0a-bc/boot.cfg</code>.

Boot the ESXi Installer by Using iPXE and HTTP

You can use iPXE to boot the ESXi installer from an HTTP server.

The prerequisites and steps depend on the support of UEFI boot or legacy BIOS only. You can use iPXE to boot the ESXi installer from an HTTP server. The following topic discusses prerequisites and steps for ESXi hosts that support UEFI boot and hosts that support legacy BIOS only.

Verify that your environment has the following components:

- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with PXE support on the target ESXi host.
- DHCP server that you can configure for PXE booting. See [Sample DHCP Configurations](#).
- TFTP server.
- Network security policies to allow TFTP traffic (UDP port 69).
- For legacy BIOS, you can use only IPv4 networking. For UEFI PXE boot, you can use IPv4 or IPv6 networking.
- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

Verify that your environment also meets the following prerequisites required for PXE boot using an HTTP Server:

- Verify that the HTTP server is accessible by your target ESXi hosts.
- If your ESXi host runs legacy BIOS firmware only, obtain version 3.86 of the SYSLINUX package. For more information, see [Network Boot Background Information](#).
- For legacy BIOS machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `pxelinux.0` initial boot loader for all target machines, but potentially different PXELINUX configuration files depending on the target machine's MAC address.

- For UEFI machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

1. Obtain and configure iPXE.

- Obtain the iPXE source code.
- On the iPXE download page, follow the build instructions, but run one of the following commands.
 - For ESXi hosts that run legacy BIOS firmware only, run `make bin/undionly.kpxe`.
 - For ESXi hosts that run UEFI firmware, run `make bin-x86_64-efi/snponly.efi`.
- Copy the `undionly.kpxe` or `snponly.efi` file to the `/tftpboot` directory on your TFTP server.

2. If your ESXi host runs legacy BIOS firmware only, obtain and configure PXELINUX.

- Obtain SYSLINUX version 3.86, unpack it, and copy the `pxelinux.0` file to the `/tftpboot` directory on your TFTP server.
- Create a PXELINUX configuration file using the following code model.

`ESXi-8.x.x-XXXXXX` is the name of the TFTP subdirectory that contains the ESXi installer files.

```
DEFAULT install
NOHALT 1
LABEL install
    KERNEL ESXi-8.x.x-XXXXXX/mboot.c32
    APPEND -c ESXi-8.x.x-XXXXXX/boot.cfg
    IPAPPEND 2
```

- Save the PXELINUX file in the `/tftpboot/pxelinux.cfg` directory on your TFTP server.

The filename determines whether all hosts boot this installer by default.

Option	Description
Same installer	Name the file <code>default</code> if you want all host to boot this ESXi installer by default.
Different installers	Name the file with the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), if only a specific host must boot this file. For example, <code>01-23-45-67-89-0a-bc</code> .

- If your ESXi host runs UEFI firmware, copy the `efi/boot/bootx64.efi` file from the ESXi installer ISO image to the `/tftpboot` folder on your TFTP server, and rename the file to `mboot.efi`.

Note: Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.

4. Configure the DHCP server.

- Create a directory on your HTTP server with the same name as the version of ESXi it will hold. For example, `/var/www/html/ESXi-8.x.x-XXXXXX`.

- Copy the contents of the ESXi installer image to the newly created directory.

7. Modify the `boot.cfg` file

- Add the following line:

```
prefix=http://XXX.XXX.XXX.XXX/ESXi-8.x.x-XXXXXX
```

where `http://XXX.XXX.XXX.XXX/ESXi-8.x.x-XXXXXX` is the location of the installer files on the HTTP server.

- b) If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (/) character, delete that character.
 - c) If the `kernelopt=` line contains the string `cdromBoot`, remove the string only.
8. Optional: For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.

Use the following code as a model, where `XXX.XXX.XXX.XXX` is the IP address of the server where the installation script resides, and `esxi_ksFiles` is the directory that contains the `ks.cfg` file.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```

9. If your ESXi host runs UEFI firmware, specify whether you want all UEFI hosts to boot the same installer.

Option	Description
Same installer	Copy or link the <code>boot.cfg</code> file to <code>/tftpboot/boot.cfg</code>
Different installers	1. Create a subdirectory of <code>/tftpboot</code> named after the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), for example, <code>01-23-45-67-89-0a-bc</code>. 2. Place a copy of (or a link to) the host's <code>boot.cfg</code> file in that directory, for example, <code>/tftpboot/01-23-45-67-89-0a-bc/boot.cfg</code>.

Boot the ESXi Installer by Using Native UEFI HTTP

You can boot the ESXi installer directly from an HTTP server, without additional software to support the process.

Verify that your environment has the following components:

- ESXi host with UEFI firmware that supports the HTTP boot feature.
- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with UEFI networking support on the target ESXi host.
- DHCP server that you can configure for UEFI HTTP booting. See [Sample DHCP Configurations](#)
- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

UEFI HTTP supports booting multiple versions of the ESXi installer. You use the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

Note: Do not mix IPv4 or IPv6 networking during the boot process. Use either IPv4 or IPv6 networking.

1. Copy the `efi/boot/bootx64.efi` file from the ESXi installer ISO image to a directory on your HTTP server and rename the file to `mboot.efi`. For example, `http://www.example.com/esxi/mboot.efi`.

Note: Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.

2. Configure the DHCP server.
3. Create a directory on your HTTP server with the same name as the version of ESXi it will hold. For example, `http://www.example.com/esxi/ESXi-8.x.x-XXXXXX`.
4. Copy the contents of the ESXi installer image to the newly created directory.
5. Modify the `boot.cfg` file.
 - a) Add the following line with the URL of the newly created directory.

```
prefix=http://www.example.com/esxi/ESXi-8.x.x-XXXXXX
```
 - b) If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (/) character, delete that character.
 - c) If the `kernelopt=` line contains the string `cdromBoot`, remove the string only.
6. Optional: For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.
 For example, `kernelopt=ks=http://www.example.com/esxi_ksFiles/ks.cfg`
7. Optional: You can use the virtual machine configuration parameters `networkBootProtocol` and `networkBootUri` to specify from where a virtual machines can boot. The setting `networkBootProtocol` specifies the boot protocol, IPv4 or IPv6. For example, `networkBootProtocol = httpv4`. The setting `networkBootUri` specifies the HTTP URL to the ESXi bootloader (`bootx64.efi`). For example, `networkBootUri = http://xxx.xxx.xx.x/esxi80uc1/efi/boot/bootx64.efi`.
8. Specify whether you want all UEFI hosts to boot the same installer.

Option	Description
Same installer	Add the <code>boot.cfg</code> file to the same directory as <code>mboot.efi</code> . For example, <code>http://www.example.com/esxi/boot.cfg</code>
Different installers	1. Create a subdirectory of the directory that contains the <code>mboot.efi</code> file. Name the directory as the MAC address of the target host machine (01-mac_address_of_target_ESXi_host), for example, 01-23-45-67-89-0a-bc. 2. Add the custom <code>boot.cfg</code> file in the directory. For example, <code>http://www.example.com/esxi/01-23-45-67-89-0a-bc/boot.cfg</code>.

You can use both installer types. ESXi hosts without custom `boot.cfg` file on your HTTP server, boot from the default `boot.cfg` file.

Sample DHCP Configurations

The DHCP server must send the address of the TFTP or HTTP server and the filename of the initial boot loader to the ESXi host.

When the target machine first boots, it broadcasts a packet across the network requesting information to boot itself. The DHCP server responds. The DHCP server must be able to determine whether the target machine is allowed to boot and the location of the initial boot loader binary. For PXE boot, the location is a file on a TFTP server. For UEFI HTTP boot, the location is a URL.

CAUTION: Do not set up a second DHCP server if your network already has one. If multiple DHCP servers respond to DHCP requests, machines can obtain incorrect or conflicting IP addresses, or can fail to receive the proper boot information. Talk to a network administrator before setting up a DHCP server. For support on configuring DHCP, contact your DHCP server vendor.

There are many DHCP servers that you can use. The following examples are for an ISC DHCP server. If you are using a version of DHCP for Microsoft Windows, see the DHCP server documentation to determine how to pass the next-server and filename arguments to the target machine.

Example of Booting Using PXE and TFTP with IPv4

This example shows how to configure an ISC DHCP server to PXE boot ESXi using a TFTP server at IPv4 address xxx.xxx.xxx.xxx.

```
#
# ISC DHCP server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server xxx.xxx.xxx.xxx;
    if option client-system-arch = 00:07 or option client-system-arch = 00:09 {
        filename = "mboot.efi";
    } else {
        filename = "pxelinux.0";
    }
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `pxelinux.0` or `mboot.efi` binary file on the TFTP server.

Example of Booting Using PXE and TFTP with IPv6

This example shows how to configure an ISC DHCPv6 server to PXE boot ESXi using a TFTP server at IPv6 address xxxx:xxxx:xxxx:xxxx::xxxx.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option dhcp6.bootfile-url code 59 = string;
option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx::xxxx]/mboot.efi";
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `mboot.efi` binary file on the TFTP server.

Example of Booting Using iPXE and HTTP with IPv4

This example shows how to configure an ISC DHCP server to boot ESXi by loading iPXE from a TFTP server at IPv4 address xxx.xxx.xxx.xxx.

```
#
# ISC DHCP server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
```

```
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server xxx.xxx.xxx.xxx;
    if option client-system-arch = 00:07 or option client-system-arch = 00:09 {
        if exists user-class and option user-class = "iPXE" {
            # Instruct iPXE to load mboot.efi as secondary bootloader
            filename = "mboot.efi";
        } else {
            # Load the snponly.efi configuration of iPXE as initial bootloader
            filename = "snponly.efi";
        }
    } else {
        if exists user-class and option user-class = "iPXE" {
            # Instruct iPXE to load pxelinux.0 as secondary bootloader
            filename = "pxelinux.0";
        } else {
            # Load the undionly configuration of iPXE as initial bootloader
            filename = "undionly.kpxe";
        }
    }
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `undionly.kpxe` or `snponly.efi` binary file on the TFTP server. In the legacy BIOS case, iPXE then asks the DHCP server for the next file to load, and the server returns `pxelinux.0` as the filename. In the UEFI case, iPXE then asks the DHCP server for the next file to load, and this time the server returns `mboot.efi` as the filename. In both cases, iPXE is resident and the system has HTTP capability. As a result, the system can load additional files from an HTTP server.

Example of Booting Using iPXE and HTTP with IPv6

This example shows how to configure an ISC DHCPv6 server to boot ESXi by loading iPXE from a TFTP server at IPv6 address `xxxx:xxxx:xxxx:xxxx:xxxx`.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;

option dhcp6.bootfile-url code 59 = string;
if exists user-class and option user-class = "iPXE" {
    # Instruct iPXE to load mboot.efi as secondary bootloader
    option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx:xxxx]/mboot.efi";
} else {
    # Load the snponly.efi configuration of iPXE as initial bootloader
    option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx:xxxx]/snponly.efi";
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `snponly.efi` (iPXE) binary file on the TFTP server. iPXE then asks the DHCP server for the next file to load, and this time the server returns `mboot.efi` as the filename. iPXE is resident and the system has HTTP capability. As a result, the system can load additional files from an HTTP server.

Example of Booting Using UEFI HTTP with IPv4

This example shows how to configure an ISC DHCP server to boot ESXi by using native UEFI HTTP over IPv4 from Web server `www.example.com`.

```
#
# ISC DHCP server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "httpclients" {
    match if substring(option vendor-class-identifier, 0, 10) = "HTTPClient";
    option vendor-class-identifier "HTTPClient";

    if option client-system-arch = 00:10 {
        # x86_64 UEFI HTTP client
        filename = http://www.example.com/esxi/mboot.efi;
    }
}
```

Example of Booting Using UEFI HTTP with IPv6

This example shows how to configure an ISC DHCPv6 server to boot ESXi by using native UEFI HTTP over IPv6 from Web server `www.example.com`.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;

option dhcp6.bootfile-url code 59 = string;
option dhcp6.user-class code 15 = { integer 16, string };
option dhcp6.vendor-class code 16 = { integer 32, integer 16, string };

if option dhcp6.client-arch-type = 00:10 {
    # x86_64 HTTP clients
    option dhcp6.vendor-class 0 10 "HTTPClient";
    option dhcp6.bootfile-url "http://www.example.com/esxi/mboot.efi";
}
```

Installing ESXi Using vSphere Auto Deploy

vSphere Auto Deploy lets you provision hundreds of physical hosts with ESXi software.

Using Auto Deploy, experienced system administrators can manage large deployments efficiently. Hosts are network-booted from a central Auto Deploy server. Optionally, hosts are configured with a host profile of a reference host. The host profile can be set up to prompt the user for input. After boot up and configuration complete, the hosts are managed by vCenter Server just like other ESXi hosts.

Auto Deploy can also be used for stateless caching or stateful installs.

Important: Auto Deploy requires a secure separation between the production network and the management or deployment networks as discussed in [vSphere Auto Deploy Security Considerations](#). Using Auto Deploy without this separation is insecure.

Stateless caching

By default, Auto Deploy does not store ESXi configuration or state on the host disk. Instead, an image profile defines the image that the host is provisioned with, and other host attributes are managed through host profiles. A host that uses Auto Deploy for stateless caching still needs to connect to the Auto Deploy server and the vCenter Server.

Stateful installs

You can provision a host with Auto Deploy and set up the host to store the image to disk. On subsequent boots, the host boots from disk.

Understanding vSphere Auto Deploy

vSphere Auto Deploy can provision hundreds of physical hosts with ESXi software.

You can specify the image to deploy and the hosts to provision with the image. Optionally, you can specify host profiles to apply to the hosts, a vCenter Server location (datacenter, folder or cluster), and assign a script bundle for each host.

Introduction to vSphere Auto Deploy

vSphere Auto Deploy uses PXE boot infrastructure with host profiles, a desired image, or configuration on a cluster level to provision ESXi hosts.

State Information for ESXi Hosts

Note: You cannot use Auto Deploy on ESXi hosts configured with DPUs as part of the vSphere Distributed Services Engine feature.

When you start a physical host that is set up for vSphere Auto Deploy, vSphere Auto Deploy uses PXE boot infrastructure in conjunction with vSphere host profiles, a desired image, or configuration on a cluster level to provision and customize that host. No state is stored on the host itself. Instead, the vSphere Auto Deploy server manages state information for each host. vSphere Auto Deploy stores the information for the ESXi hosts to be provisioned in different locations. Information about the location of image profiles, host profiles, or clusters that you manage by either a single image or by a configuration on a cluster level is initially specified in the rules that map machines to image profiles and host profiles.

Table 17: vSphere Auto Deploy Stores Information for Deployment


Information Type	Description	Source of Information
Image state	The executable software to run on an ESXi host.	Image profile, created with vSphere ESXi Image Builder or a vSphere Lifecycle Manager image.
Configuration state	The configurable settings that determine how the host is configured, for example, virtual switches and their settings, driver settings, boot parameters, and so on.	Host profile, created by using the host profile UI, or a configuration that you create when setting up a cluster that manages all ESXi host settings at a cluster level in the Inventory UI.

Information Type	Description	Source of Information
Dynamic state	The runtime state that is generated by the running software, for example, generated private keys or runtime databases.	Host memory, lost during reboot.
Virtual machine state	The virtual machines stored on a host and virtual machine autostart information (subsequent boots only).	Virtual machine information sent by vCenter Server to vSphere Auto Deploy must be available to supply virtual machine information to vSphere Auto Deploy.
User input	State that is based on user input, for example, an IP address that the user provides when the system starts up, cannot automatically be included in the host profile.	Host customization information, stored by vCenter Server during first boot. You can create a host profile that requires user input for certain values. When vSphere Auto Deploy applies a host profile that requires user provided information, the host is placed in maintenance mode. Use the host profile UI to check the host profile compliance, and respond to the prompt to customize the host.

vSphere Auto Deploy Architecture

The vSphere Auto Deploy infrastructure consists of several components.

Figure 17: vSphere Auto Deploy Architecture



vSphere Auto Deploy server

vSphere Auto Deploy rules engine

Serves images and host profiles to ESXi hosts.

Apart from legacy image profiles that you create by using the VMware Image Builder and host profiles, you can also create

Image profiles

vSphere Auto Deploy rules to deploy ESXi by using a single vSphere Lifecycle Manager image or a configuration on a cluster level.

Define the set of VIBs to boot ESXi hosts with.

- VMware and VMware partners make image profiles and VIBs available in public depots. Use vSphere ESXi Image Builder to examine the depot and use the vSphere Auto Deploy rules engine to specify which image profile to assign to which host.
- You use vSphere Lifecycle Manager images to apply software and firmware updates to the ESXi hosts in a cluster. Using a single image to manage all hosts in a cluster ensures cluster-wide host image homogeneity.
- With ESXi 8.0, you can set up a cluster that manages all ESXi host settings at a cluster level.

Host profiles

Host customization

Stores information that the user provides when host profiles are applied to the host. Host customization might contain an IP address or other information that the user supplied for that host. For more information about host customizations, see the *vSphere Host Profiles* documentation.

Host customization was called answer file in earlier releases of vSphere Auto Deploy.

Auto Deploy Certificates

By default, the Auto Deploy server provisions each host with certificates that are signed by the VMware Certificate Authority (VMware CA). For more information, see [Managing Certificates for ESXi Hosts](#).

Alternatively, if your corporate policy requires that you use custom certificates, you can set up the Auto Deploy server to provision all hosts with custom certificates that are not signed by VMware CA. The Auto Deploy server becomes a subordinate certificate authority of your third-party CA. In the Custom Certificate Authority mode, you are responsible for managing the certificates. You cannot refresh and renew certificates from the vSphere Client. In this mode, you also cannot select only a set of hosts to provision with custom certificates, and you can manually sign custom certificates only for stateful hosts. For more information, see [Use Custom Certificates with Auto Deploy](#).

With ESXi8.0, Auto Deploy provides a third option that allows you to generate a certificate outside vSphere and become independent of the certificate management in vCenter Server. For example, you can generate a custom certificate by using a custom script or by using a provider of domain name registry services such as Verisign. You can use custom certificates for only a set of ESXi hosts. You can provide custom certificates for stateless hosts as well. ESXi hosts are identified by the MAC address of the NIC used for network booting, or the BIOS UUID of the ESXi host. You update the VMware Endpoint Certificate Store (VECS) with the custom certificate by using PowerCLI. For more information on the new PowerCLI cmdlets, see [vSphere Auto Deploy PowerCLI Cmdlet Overview](#). The VMware CA must trust the custom ESXi certificates so you must add the CA public certificate for the custom certificates to the TRUSTED_ROOTS store in VECS. Auto Deploy also stores the custom certificates and when it recognizes a booting host with the respective MAC address of the NIC used for network booting, or the BIOS UUID of the ESXi host, it automatically provides the custom certificate. You do not need to stop or restart Auto Deploy or vCenter Server when you add a custom certificate to VECS, only restart the host for which you upload a custom certificate. For more information, see [Use Custom Certificates with Auto Deploy](#).

Rules and Rule Sets

You specify the behavior of the vSphere Auto Deploy server by using a set of rules. The vSphere Auto Deploy rules engine checks the rule set for matching host patterns to decide which items (image profile, host profile, vCenter Server location, or script object) to provision each host with.

The rules engine maps software and configuration settings to hosts based on the attributes of the host. For example, you can deploy image profiles or host profiles to two clusters of hosts by writing two rules, each matching on the network address of one cluster.

For hosts that have not yet been added to a vCenter Server system, the vSphere Auto Deploy server checks with the rules engine before serving image profiles, host profiles, and inventory location information to hosts. For hosts that are managed by a vCenter Server system, the image profile, host profile, and inventory location that vCenter Server has stored in the host object is used. If you make changes to rules, you can use the vSphere Client or vSphere Auto Deploy cmdlets in a PowerCLI session to test and repair rule compliance. When you repair rule compliance for a host, that host's image profile and host profile assignments are updated.

The rules engine includes rules and rule sets.

Rules

Rules can assign image profiles and host profiles to a set of hosts, or specify the location (folder or cluster) of a host on the target vCenter Server system. A rule can identify target hosts by boot MAC address, SMBIOS information, BIOS UUID, Vendor, Model, or fixed DHCP IP address. In most cases, rules apply to multiple hosts. You create rules by using the vSphere Client or vSphere Auto Deploy cmdlets in a PowerCLI session. After you create a rule, you must add it to a rule set. Only two rule sets, the active rule set and the working rule set, are supported. A rule can belong to both sets, the default, or only to the working rule set. After you add a rule to a rule set, you can no longer change the rule. Instead, you copy the rule and replace items or patterns in the copy. If you are managing vSphere Auto Deploy with the vSphere Client, you can edit a rule if it is in inactive state.

You can specify the following parameters in a rule.

Parameter	Description
Name	Name of the rule, specified with the <code>-Name</code> parameter.
Item	One or more items, specified with the <code>-Item</code> parameter. An item can be an image profile, a host profile, a vCenter Server inventory location (datacenter, folder, cluster) for the target host, or a custom script. You can specify multiple items separated by commas.
Pattern	The pattern specifies the host or group of hosts to which the rule applies. vendor Machine vendor name. model Machine model name. serial Machine serial number. hostname Machine hostname. domain Domain name.

Parameter	Description
ipv4	IPv4 address of the machine.
ipv6	IPv6 address of the machine.
	PXE booting with BIOS firmware is possible only with IPv4, PXE booting with UEFI firmware is possible with either IPv4 or IPv6.
mac	Boot NIC MAC address.
asset	Machine asset tag.
oemstring	OEM-specific strings in the SMBIOS.
	You can specify <code>-AllHosts</code> to apply the item or items to all hosts.

Active Rule Set

When a newly started host contacts the vSphere Auto Deploy server with a request for an image profile, the vSphere Auto Deploy server checks the active rule set for matching rules. The image profile, host profile, vCenter Server inventory location, and script object that are mapped by matching rules are then used to boot the host. If more than one item of the same type is mapped by the rules, the vSphere Auto Deploy server uses the item that is first in the rule set.

Working Rule Set

The working rule set allows you to test changes to rules before making the changes active. For example, you can use vSphere Auto Deploy cmdlets for testing compliance with the working rule set. The test verifies that hosts managed by a vCenter Server system are following the rules in the working rule set. By default, cmdlets add the rule to the working rule set and activate the rules. Use the `NoActivate` parameter to add a rule only to the working rule set.

You use the following workflow with rules and rule sets.

1. Make changes to the working rule set.
2. Test the working rule set rules against a host to make sure that everything is working correctly.
3. Refine and retest the rules in the working rule set.
4. Activate the rules in the working rule set.

If you add a rule in a PowerCLI session and do not specify the `NoActivate` parameter, all rules that are currently in the working rule set are activated. You cannot activate individual rules.

See the PowerCLI command-line help and [Managing vSphere Auto Deploy with PowerCLI Cmdlets](#) for more information on using vSphere Auto Deploy with PowerCLI cmdlets. See [Managing vSphere Auto Deploy with the vSphere Client](#) for more information on using vSphere Auto Deploy with the vSphere Client.

Install and Configure vSphere Auto Deploy

Before you can start using vSphere Auto Deploy, you must prepare your environment in several steps.

You start with server setup and hardware preparation. You must configure the vSphere Auto Deploy service startup type in the vCenter Server system that you plan to use for managing the hosts you provision, and install PowerCLI.

vSphere Auto Deploy Preinstallation Checklist

Before you can start the tasks in this vSphere Auto Deploy scenario, make sure that your environment meets the hardware and software requirements, and that you have the necessary permissions for the components included in the setup.

Table 18: Preinstallation Checklist

Required Software and Hardware	Details
vCenter Server	The vSphere Auto Deploy server is part of vCenter Server. You must enable and start the vSphere Auto Deploy service on the vCenter Server system. You can perform many of the setup tasks by logging in to vCenter Server. See Prepare Your System for vSphere Auto Deploy .
Storage	Storage for ESXi datastores NFS, iSCSI, or Fibre Channel, with servers and storage arrays that are configured so the servers can detect the LUNs. - A list of target IP addresses for NFS or iSCSI. - A list of target volume information for NFS or iSCSI.
Host information (for four ESXi hosts)	A list of target IP addresses for NFS or iSCSI. A list of target volume information for NFS or iSCSI. - Default route, net mask, and primary and secondary DNS server IP addresses. - IP address and net mask for the VMkernel primary management network. - IP address and net mask for other VMkernel networks such as storage, vSphere FT, or VMware vMotion. vSphere Auto Deploy does not overwrite existing partitions by default.
PowerCLI	See Install PowerCLI .
ESXi software depot	The location of the ESXi software depot on the Downloads page of the VMware website. You use a URL to point to the image profile stored at that location, or you download a ZIP file to work with a local depot. Do not download the ESXi image.
TFTP server	TFTP installer software such as WinAgents TFTP server.
DHCP server	The DHCP server is included in the vSphere supported Windows Server versions.
DNS server	A working DNS server. You must add entries in both Forward (A Record) and Reverse (PTR Record) Zones for each target host.

You also need information about and administrator privileges to the core servers of the environment, including the ActiveDirectory server, DNS server, DHCP server, NTP server, and so on.

You must have complete control of the broadcast domain of the subnet in which you deploy the setup. Ensure that no other DHCP, DNS, or TFTP server are on this subnet.

Prepare Your System for vSphere Auto Deploy

Before you can PXE boot an ESXi host with vSphere Auto Deploy, you must install prerequisite software and set up the DHCP and TFTP servers that vSphere Auto Deploy interacts with.

- Verify that the hosts that you plan to provision with vSphere Auto Deploy meet the hardware requirements for ESXi. See [ESXi Hardware Requirements](#).
- Verify that the ESXi hosts have network connectivity to vCenter Server and that all port requirements are met. See *vCenter Server Upgrade*.
- Verify that you have a TFTP server and a DHCP server in your environment to send files and assign network addresses to the ESXi hosts that Auto Deploy provisions. See [vSphere PowerCLI Scenario for vSphere Auto Deploy](#).
- Verify that the ESXi hosts have network connectivity to DHCP, TFTP, and vSphere Auto Deploy servers.
- If you want to use VLANs in your vSphere Auto Deploy environment, you must set up the end to end networking properly. When the host is PXE booting, the firmware driver must be set up to tag the frames with proper VLAN IDs. You must do this set up manually by making the correct changes in the UEFI/BIOS interface. You must also correctly configure the ESXi port groups with the correct VLAN IDs. Ask your network administrator how VLAN IDs are used in your environment.
- Verify that you have enough storage for the vSphere Auto Deploy repository. The vSphere Auto Deploy server uses the repository to store data it needs, including the rules and rule sets you create and the VIBs and image profiles that you specify in your rules.

Best practice is to allocate 2 GB to have enough room for four image profiles and some extra space. Each image profile requires approximately 400 MB. Determine how much space to reserve for the vSphere Auto Deploy repository by considering how many image profiles you expect to use.

- Obtain administrative privileges to the DHCP server that manages the network segment you want to boot from. You can use a DHCP server already in your environment, or install a DHCP server. For your vSphere Auto Deploy setup, replace the `gpxelinux.0` filename with `snponly64.efi.vmw-hardwired` for UEFI or `undionly.kpxe.vmw-hardwired` for BIOS. For more information on DHCP configurations, see [Sample DHCP Configurations](#).
- Secure your network as for any other PXE-based deployment method. vSphere Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or the vSphere Auto Deploy server is not checked during a PXE boot.
- If you want to manage vSphere Auto Deploy with PowerCLI cmdlets, verify that Microsoft .NET Framework 4.5 or 4.5.x and Windows PowerShell 3.0 or 4.0 are installed on a Windows machine. See the *vSphere PowerCLI User's Guide*.
- Set up a remote Syslog server. See the *vCenter Server and Host Management* documentation for Syslog server configuration information. Configure the first host you boot to use the remote Syslog server and apply that host's host profile to all other target hosts. Optionally, install and use VMware vCenter Log Insight, which provides log aggregation and analytics for VMware and non-VMware products, virtual and physical, with near real-time search and analytics of log events.
- If the hosts that you plan to provision with vSphere Auto Deploy are with legacy BIOS, verify that the vSphere Auto Deploy server has an IPv4 address. PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

If you want to manage vSphere Auto Deploy with PowerCLI cmdlets, see *Set Up vSphere Auto Deploy and Provision Hosts with vSphere PowerCLI*.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Auto Deploy** page, select your vCenter Server from the drop-down menu at the top.
3. Click **Enable Auto Deploy and Image Builder** to activate the service.
If the **Image Builder** service is already active, select the **Configure** tab and click **Enable Auto Deploy Service**.
The **Software Depot** page appears.
4. Configure the TFTP server.
 - a) Click the **Configure** tab.
 - b) Click **Download TFTP Boot Zip** to download the TFTP configuration file and unzip the file to the directory in which your TFTP server stores files.
 - c) Optional: To use a proxy server, click **Add** on the *Auto Deploy Runtime Summary* pane and enter a proxy server URL in the text box.
Using reverse proxy servers can offload the requests made to the vSphere Auto Deploy server.
5. Set up your DHCP server to point to the TFTP server on which the TFTP ZIP file is located.
 - a) Specify the TFTP Server's IP address in DHCP option 66, frequently called next-server.
 - b) Specify the boot filename, which is `snponly64.efi.vmw-hardwired` for UEFI or `undionly.kpxe.vmw-hardwired` for BIOS in the DHCP option 67, frequently called `boot-filename`.
6. Set each host you want to provision with vSphere Auto Deploy to network boot or PXE boot, following the manufacturer's instructions.
7. Optional: If you set up your environment to use Thumbprint mode, you can use your own Certificate Authority (CA) by replacing the OpenSSL certificate `rbd-ca.crt` and the OpenSSL private key `rbd-ca.key` with your own certificate and key file.
The files are in `/etc/vmware-rbd/ssl/`.
By default, vCenter Server uses VMware Certificate Authority (VMCA).

When you start an ESXi host that is set up for vSphere Auto Deploy, the host contacts the DHCP server and is directed to the vSphere Auto Deploy server, which provisions the host with the image profile specified in the active rule set.

- You can change the default configuration properties of the **Auto Deploy Service**. For more information, see "Configuring vCenter Server" in the *vCenter Server and Host Management* documentation.
- You can change the default configuration properties of the **Image Builder Service**. For more information, see "Configuring vCenter Server" in the *vCenter Server and Host Management* documentation.
- Configure the first host that you provision as a reference host. Use the storage, networking, and other settings you want for your target hosts to share. Create a host profile for the reference host and write a rule that assigns both the already tested image profile and the host profile to target hosts.
- If you have to configure host-specific information, set up the host profile of the reference host to prompt for user input. For more information about host customizations, see the *vSphere Host Profiles* documentation.

Using vSphere Auto Deploy Cmdlets

vSphere Auto Deploy cmdlets are implemented as Microsoft PowerShell cmdlets and included in PowerCLI. Users of vSphere Auto Deploy cmdlets can take advantage of all PowerCLI features.

Experienced PowerShell users can use vSphere Auto Deploy cmdlets just like other PowerShell cmdlets. If you are new to PowerShell and PowerCLI, the following tips might be helpful.

You can type cmdlets, parameters, and parameter values in the PowerCLI shell.

- Get help for any cmdlet by running `Get-Help cmdlet_name`.
- Remember that PowerShell is not case sensitive.
- Use tab completion for cmdlet names and parameter names.

- Format any variable and cmdlet output by using `Format-List` or `Format-Table`, or their short forms `fl` or `ft`. For more information, run the `Get-Help Format-List cmdlet`.

Passing Parameters by Name

You can pass in parameters by name in most cases and surround parameter values that contain spaces or special characters with double quotes.

```
Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile
```

Most examples in the *vCenter Server Installation and Setup* documentation pass in parameters by name.

Passing Parameters as Objects

You can pass parameters as objects if you want to perform scripting and automation. Passing in parameters as objects is useful with cmdlets that return multiple objects and with cmdlets that return a single object. Consider the following example.

1. Bind the object that encapsulates rule set compliance information for a host to a variable.

```
$tr = Test-DeployRuleSetCompliance MyEsxi42
```

2. View the `itemlist` property of the object to see the difference between what is in the rule set and what the host is currently using.

```
$tr.itemlist
```

3. Remediate the host to use the revised rule set by using the `Repair-DeployRuleSetCompliance` cmdlet with the variable.

```
Repair-DeployRuleSetCompliance $tr
```

The example remediates the host the next time you boot the host.

Set Up Bulk Licensing

You can use the vSphere Client or ESXi Shell to specify individual license keys, or you can set up bulk licensing by using PowerCLI cmdlets. Bulk licensing works for all ESXi hosts, but is especially useful for hosts provisioned with vSphere Auto Deploy.

[#unique_239](#).

Assigning license keys through the vSphere Client and assigning licensing by using PowerCLI cmdlets function differently.

Assign license keys with the vSphere Client

You can assign license keys to a host when you add the host to the vCenter Server system or when the host is managed by a vCenter Server system.

Assign license keys with LicenseDataManager PowerCLI

You can specify a set of license keys to be added to a set of hosts. The license keys are added to the vCenter Server database. Each time a host is added to the vCenter Server system or reconnects to it, the host is assigned a license key. A license key that is assigned through PowerCLI is treated as a default license key. When an unlicensed host is added or reconnected, it is assigned the default license key. If a host is already licensed, it keeps its license key.

The following example assigns licenses to all hosts in a data center. You can also associate licenses with hosts and clusters.

The following example is for advanced PowerCLI users who know how to use PowerShell variables.

1. In a PowerCLI session, connect to the vCenter Server system you want to use and bind the associated license manager to a variable.

```
Connect-VIServer -Server 192.XXX.X.XX -User username -Password password
$licenseDataManager = Get-LicenseDataManager
```

2. Run a cmdlet that retrieves the data center in which the hosts for which you want to use the bulk licensing feature are located.

```
$hostContainer = Get-Datacenter -Name Datacenter-X
```

You can also run a cmdlet that retrieves a cluster to use bulk licensing for all hosts in a cluster, or retrieves a folder to use bulk licensing for all hosts in a folder.

3. Create a `LicenseData` object and a `LicenseKeyEntry` object with associated type ID and license key.

```
$licenseData = New-Object VMware.VimAutomation.License.Types.LicenseData
$licenseKeyEntry = New-Object VMware.VimAutomation.License.Types.LicenseKeyEntry
$licenseKeyEntry.TypeId = "vmware-vmware"
$licenseKeyEntry.LicenseKey = "XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX"
```

4. Associate the `LicenseKeys` attribute of the `LicenseData` object you created in step 3 with the `LicenseKeyEntry` object.

```
$licenseData.LicenseKeys += $licenseKeyEntry
```

5. Update the license data for the data center with the `LicenseData` object and verify that the license is associated with the host container.

```
$licenseDataManager.UpdateAssociatedLicenseData($hostContainer.Uid, $licenseData)
$licenseDataManager.QueryAssociatedLicenseData($hostContainer.Uid)
```

6. Provision one or more hosts with vSphere Auto Deploy and assign them to the data center or to the cluster that you assigned the license data to.

7. You can use the vSphere Client to verify that the host is successfully assigned to the default license `XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX`.

All hosts that you assigned to the data center are now licensed automatically.

How vSphere Auto Deploy Works with PowerCLI

You can manage vSphere Auto Deploy with PowerCLI cmdlets to create rules that associate ESXi hosts with various objects on the vCenter target.

You can manage vSphere Auto Deploy with PowerCLI cmdlets to create rules that associate hosts with image profiles, host profiles, custom scripts and locations on the vCenter Server target. You can also update hosts by testing rule compliance and repairing compliance issues.

Managing vSphere Auto Deploy with PowerCLI Cmdlets

You can manage vSphere Auto Deploy with PowerCLI cmdlets to create rules that associate hosts with image profiles, host profiles, custom scripts and locations on the vCenter Server target. You can also update hosts by testing rule compliance and repairing compliance issues.

Overview of the vSphere Auto Deploy Process by Using PowerCLI

Getting started with vSphere Auto Deploy requires that you learn how vSphere Auto Deploy works, install the vSphere Auto Deploy server, install PowerCLI, write PowerCLI rules that provision hosts, and power on your hosts to be booted with the image profile you specify. You can customize of the image profile, host profile, and vCenter Server location.

See [vSphere PowerCLI Scenario for vSphere Auto Deploy](#) for a step-by-step exercise that helps you set up your first vSphere Auto Deploy environment.

To provision the hosts in your environment with vSphere Auto Deploy successfully, you can follow these steps.

1. Deploy vCenter Server.
The vSphere Auto Deploy server is included.
2. Configure the vSphere Auto Deploy service startup type.
See [Prepare Your System for vSphere Auto Deploy](#).
3. Install PowerCLI, which includes vSphere Auto Deploy and vSphere ESXi Image Builder cmdlets.
See [Configure vSphere ESXi Image Builder](#), [Using vSphere Auto Deploy Cmdlets](#), and [Using VMware Image Builder Cmdlets](#).
4. Find the image profile that includes the VIBs that you want to deploy to your hosts.
 - Usually, you add the depots containing the required software to your PowerCLI session, and then select an image profile from one of those depots.
 - To create a custom image profile, use vSphere ESXi Image Builder cmdlets to clone an existing image profile and add the custom VIBs to the clone. Add the custom image profile to the PowerCLI session.

You must use vSphere ESXi Image Builder for customization only if you have to add or remove VIBs. In most cases, you can add the depot where VMware hosts the image profiles to your PowerCLI session as a URL.

5. Start a PowerCLI session and connect to the vCenter Server system that vSphere Auto Deploy is registered with.
6. Use the `New-DeployRulePowerCLI` cmdlet to write a rule that assigns the image profile to one host, to multiple hosts specified by a pattern, or to all hosts.

```
New-DeployRule -Name "testrule" -Item image-profile -AllHosts
```

See [Assign an Image Profile to Hosts](#).

Note: vSphere Auto Deploy is optimized for provisioning hosts that have a fixed MAC address to IP address mapping in DHCP (sometimes called DHCP reservations). If you want to use static IP addresses, you must set up the host profile to prompt for host customization. For more information, see the *vSphere Host Profiles* documentation.

7. Power on the hosts that you want to provision.
8. Set up the host you provisioned as a reference host for your host profile.
You can specify the reference host syslog settings, firewall settings, storage, networking, and so on.
9. Set up the host you provisioned as a reference host for your host profile.
You can specify the reference host syslog settings, firewall settings, storage, networking, and so on. See [Setting Up a vSphere Auto Deploy Reference Host](#).
10. Create and export a host profile for the reference host.
See the *Host Profiles* documentation.
11. To provision multiple hosts with the host profile, use the `Copy-DeployRule` cmdlet to edit the previously created rule.
You can revise the rule to assign not only an image profile but also a host profile, a vCenter Server location, and a custom script bundle.

```
Copy-DeployRule -DeployRule "testrule" -ReplaceItem my_host_profile_from_reference_host,my_target_cluster
-ReplacePattern "ipv4=192.XXX.1.10-192.XXX.1.20"
```

Where `my_host_profile_from_reference_host` is the name of the reference host profile, and `my_target_cluster` is the name of the target cluster.

12. Perform the test and repair compliance operations to remediate the hosts.
See [Test and Repair Rule Compliance](#).

13. Verify that the hosts you provisioned meet the following requirements.

- Each host is connected to the vCenter Server system.
- The hosts are not in maintenance mode.
- The hosts have no compliance failures.
- Each host with a host profile that requires user input has up-to-date host customization information.

Remediate host associations and compliance problems and reboot hosts until all hosts meet the requirements.

Read for an introduction to the boot process, differences between first and subsequent boots, and an overview of using host customization.

Using vSphere Auto Deploy Cmdlets

vSphere Auto Deploy cmdlets are implemented as Microsoft PowerShell cmdlets and included in PowerCLI. Users of vSphere Auto Deploy cmdlets can take advantage of all PowerCLI features.

Experienced PowerShell users can use vSphere Auto Deploy cmdlets just like other PowerShell cmdlets. If you are new to PowerShell and PowerCLI, the following tips might be helpful.

You can type cmdlets, parameters, and parameter values in the PowerCLI shell.

- Get help for any cmdlet by running `Get-Help cmdlet_name`.
- Remember that PowerShell is not case sensitive.
- Use tab completion for cmdlet names and parameter names.
- Format any variable and cmdlet output by using `Format-List` or `Format-Table`, or their short forms `fl` or `ft`. For more information, run the `Get-Help Format-List cmdlet`.

Passing Parameters by Name

You can pass in parameters by name in most cases and surround parameter values that contain spaces or special characters with double quotes.

```
Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile
```

Most examples in the *vCenter Server Installation and Setup* documentation pass in parameters by name.

Passing Parameters as Objects

You can pass parameters as objects if you want to perform scripting and automation. Passing in parameters as objects is useful with cmdlets that return multiple objects and with cmdlets that return a single object. Consider the following example.

1. Bind the object that encapsulates rule set compliance information for a host to a variable.

```
$tr = Test-DeployRuleSetCompliance MyEsxi42
```

2. View the `itemlist` property of the object to see the difference between what is in the rule set and what the host is currently using.

```
$tr.itemlist
```

3. Remediate the host to use the revised rule set by using the `Repair-DeployRuleSetCompliance` cmdlet with the variable.

```
Repair-DeployRuleSetCompliance $tr
```

The example remediates the host the next time you boot the host.

vSphere Auto Deploy PowerCLI Cmdlet Overview

You specify the rules that assign image profiles and host profiles to hosts using a set of PowerCLI cmdlets that are included in PowerCLI.

If you are new to PowerCLI, read the PowerCLI documentation and review [Using vSphere Auto Deploy Cmdlets](#). You can get help for any command at the PowerShell prompt.

- Basic help: `Get-Help cmdlet_name`
- Detailed help: `Get-Help cmdlet_name-Detailed`

Note: When you run vSphere Auto Deploy cmdlets, provide all parameters on the command line when you invoke the cmdlet. Supplying parameters in interactive mode is not recommended.

Table 19: Rule Engine PowerCLI Cmdlets

Command	Description
<code>Get-DeployCommand</code>	Returns a list of vSphere Auto Deploy cmdlets.
<code>New-DeployRule</code>	Creates a new rule with the specified items and patterns.
<code>Set-DeployRule</code>	Updates an existing rule with the specified items and patterns. You cannot update a rule that is part of a rule set.
<code>Get-DeployRule</code>	Retrieves the rules with the specified names.
<code>Copy-DeployRule</code>	Clones and updates an existing rule.
<code>Add-DeployRule</code>	Adds one or more rules to the working rule set and, by default, also to the active rule set. Use the <code>NoActivate</code> parameter to add a rule only to the working rule set.
<code>Remove-DeployRule</code>	Removes one or more rules from the working rule set and from the active rule set. Run this command with the <code>-Delete</code> parameter to completely delete the rule.
<code>Set-DeployRuleset</code>	Explicitly sets the list of rules in the working rule set.
<code>Get-DeployRuleset</code>	Retrieves the current working rule set or the current active rule set.
<code>Switch-ActiveDeployRuleset</code>	Activates a rule set so that any new requests are evaluated through the rule set.
<code>Get-VMHostMatchingRules</code>	Retrieves rules matching a pattern. For example, you can retrieve all rules that apply to a host or hosts. Use this cmdlet primarily for debugging.
<code>Test-DeployRulesetCompliance</code>	Checks whether the items associated with a specified host are in compliance with the active rule set.
<code>Repair-DeployRulesetCompliance</code>	Given the output of <code>Test-DeployRulesetCompliance</code> , this cmdlet updates the image profile, host profile, and location for each host in the vCenter Server inventory. The cmdlet might apply image profiles, apply host profiles, or move hosts to prespecified folders or clusters on the vCenter Server system.
<code>Apply-EsxImageProfile</code>	Associates the specified image profile with the specified host.
<code>Get-VMHostImageProfile</code>	Retrieves the image profile in use by a specified host. This cmdlet differs from the <code>Get-EsxImageProfile</code> cmdlet in vSphere ESXi Image Builder.
<code>Repair-DeployImageCache</code>	Use this cmdlet only if the vSphere Auto Deploy image cache is accidentally deleted.
<code>Get-VMHostAttributes</code>	Retrieves the attributes for a host that are used when the vSphere Auto Deploy server evaluates the rules.

Command	Description
<code>Get-DeployMachineIdentity</code>	Returns a string value that vSphere Auto Deploy uses to logically link an ESXi host in vCenter Server to a physical machine.
<code>Set-DeployMachineIdentity</code>	Logically links a host object in the vCenter Server database to a physical machine. Use this cmdlet to add hosts without specifying rules.
<code>Get-DeployOption</code>	Retrieves the vSphere Auto Deploy global configuration options. This cmdlet currently supports the <code>vlan-id</code> option, which specifies the default VLAN ID for the ESXi Management Network of a host provisioned with vSphere Auto Deploy. vSphere Auto Deploy uses the value only if the host boots without a host profile.
<code>Set-DeployOption</code>	Sets the value of a global configuration option. Currently supports the <code>vlan-id</code> option for setting the default VLAN ID for the ESXi Management Network.
<code>Add-ProxyServer</code>	Adds a proxy server to the vSphere Auto Deploy database. Run the command with the <code>-Address</code> parameter to specify the IPv4 or IPv6 address. The address can include a port number.
<code>List-ProxyServer</code>	Lists the proxy servers that are currently registered with vSphere Auto Deploy.
<code>Delete-ProxyServer</code>	Deletes one or more proxy servers from the list of proxy servers that are registered with vSphere Auto Deploy. You can run the command with the <code>-id</code> parameter from the list of proxy servers or with the <code>-Address</code> parameter by specifying the IPv4 or IPv6 address of the proxy server you want to delete.
<code>Add-ScriptBundle</code>	Adds one or more script bundles to the vSphere Auto Deploy server.
<code>Get-ScriptBundle</code>	Retrieves the list of script bundles available on the vSphere Auto Deploy server and the scripts they contain.
<code>Remove-ScriptBundle</code>	Removes a script bundle from vSphere Auto Deploy. Applicable for vSphere version 6.7 and later.
<code>Get-CustomCertificate</code>	Retrieves the custom host certificate uploaded into AutoDeploy. You must run the command with the <code>-HostId [MAC_Address BIOS_UUID]</code> parameter. The first time you add custom certificates, you don't see any certificates returned by this cmdlet.
<code>List-CustomCertificates</code>	Retrieves information about all custom host certificates used by Auto Deploy. The list provides details for the name of the certificate, <code>Host ID</code> , and <code>Associated Host Name</code> , which reflects the name of the vCenter Server for the Auto Deploy server.
<code>Add-CustomCertificate</code>	Adds a custom certificate to the VMware Endpoint Certificate Store and associates it with an ESXi host. The certificate becomes active upon host reboot. You can use the <code>Get-CustomCertificate</code> cmdlet to retrieve the custom host certificate key. You can run the command with the <code>-HostId [MAC_Address BIOS_UUID]</code> parameter to associate the certificate to the host, specifying a <code>-Key [file:///path/to/key.key]</code> and <code>-Cert [file:///path/to/cert.crt]</code> . Using this cmdlet requires the AutoDeploy.Rule.Create privilege on the root folder of vCenter Server.

Command	Description
Remove-CustomCertificate	Removes a set of custom host certificates from Auto Deploy. The certificate entries are deleted from the database and the certificate files are removed from the filestore. Hosts that have already booted with a custom certificate must be rebooted to receive a new certificate. You must provide at least one of <code>-Cert</code> or <code>-HostId</code> parameters. Using this cmdlet requires the AutoDeploy.Rule.Create privilege on the root folder of vCenter Server.

Assign an Image Profile to Hosts

Before you can provision a host, you must create rules that assign an image profile to each host that you want to provision by using vSphere Auto Deploy.

vSphere Auto Deploy extensibility rules enforce that VIBs at the `CommunitySupported` level can only contain files from certain predefined locations, such as the ESXCLI plug-in path, jumpstart plug-in path, and so on. If you add a VIB that is in a different location to an image profile, a warning results. You can override the warning by using the `force` option.

If you call the `New-DeployRule` cmdlet on an image profile that includes VIBs at the `CommunitySupported` level which violate the rule, set `$DeployNoSignatureCheck = $true` before adding the image profile. With that setting, the system ignores signature validation and does not perform the extensibility rules check.

Note: Image profiles that include VIBs at the `CommunitySupported` level are not supported on production systems.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Determine the location of a public software depot, or define a custom image profile by using vSphere ESXi Image Builder.
3. Run `Add-EsxSoftwareDepot` to add the software depot that contains the image profile to the PowerCLI session.

Depot Type	Cmdlet
Remote depot	Run <code>Add-EsxSoftwareDepot depot_url</code> .
ZIP file	1. Download the ZIP file to a local file path. 2. Run <code>Add-EsxSoftwareDepot C:\file_path\my_offline_depot.zip</code>.

4. In the depot, find the image profile that you want to use by running the `Get-EsxImageProfile` cmdlet.
By default, the ESXi depot includes one base image profile that includes VMware tools and has the string `standard` in its name, and one base image profile that does not include VMware tools.
5. Define a rule in which hosts with certain attributes, for example a range of IP addresses, are assigned to the image profile.

```
New-DeployRule -Name "testrule" -Item "My Profile25" -Pattern "vendor=Acme,Zven", "ipv4=192.XXX.1.10-192.XXX.1.20"
```

Double quotes are required if a name contains spaces, optional otherwise. Specify `-AllHosts` instead of a pattern to apply the item to all hosts.

The cmdlet creates a rule named `testrule`. The rule assigns the image profile named `My Profile25` to all hosts with a vendor of `Acme` or `Zven` that also have an IP address in the specified range.

6. Add the rule to the rule set.

```
Add-DeployRule testrule
```

By default, the rule is added to both the working rule set and the active rule set. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

When the host boots from iPXE, it reports attributes of the machine to the console. Use the same format of the attributes when writing deploy rules.

```
*****
* Booting through VMware AutoDeploy...
*
* Machine attributes:
* . asset=No Asset Tag
* . domain=vmware.com
* . hostname=myhost.mycompany.com
* . ipv4=XX.XX.XXX.XXX
* . mac=XX:XX:XX:XX:XX:XX
* . model=MyVendorModel
* . oemstring=Product ID: XXXXXX-XXX
* . serial=XX XX XX XX XX XX...
* . uuid=XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX
* . vendor=MyVendor
*****
```

- For hosts already provisioned with vSphere Auto Deploy, perform the compliance testing and repair operations to provision them with the new image profile. See [Test and Repair Rule Compliance](#).
- Turn on unprovisioned hosts to provision them with the new image profile.

Write a Rule and Assign a Host Profile to Hosts

vSphere Auto Deploy can assign a host profile to one or more ESXi hosts.

The host profile might include information about storage configuration, network configuration, or other characteristics of the host. If you add a host to a cluster, that cluster's host profile is used.

- Install PowerCLI and all prerequisite software. For information, see [vCenter Server Installation and Setup](#).
- Export the host profile that you want to use.

In many cases, you assign a host to a cluster instead of specifying a host profile explicitly. The host uses the host profile of the cluster.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Using the vSphere Client, set up a host with the settings you want to use and create a host profile from that host.
3. Find the name of the host profile by running `Get-VMhostProfilePowerCLI` cmdlet, passing in the ESXi host from which you create a host profile.
4. At the PowerCLI prompt, define a rule in which host profiles are assigned to hosts with certain attributes, for example a range of IP addresses.

```
New-DeployRule -Name "testrule2" -Item my_host_profile -Pattern "vendor=Acme,Zven",
"ipv4=192.XXX.1.10-192.XXX.1.20"
```

The specified item is assigned to all hosts with the specified attributes. This example specifies a rule named `testrule2`. The rule assigns the specified host profile `my_host_profile` to all hosts with an IP address inside the specified range and with a manufacturer of Acme or Zven.

5. Add the rule to the rule set.

```
Add-DeployRule testrule2
```

By default, the working rule set becomes the active rule set, and any changes to the rule set become active when you add a rule. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

- Assign a host already provisioned with vSphere Auto Deploy to the new host profile by performing compliance test and repair operations on those hosts. For more information, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the host profile.

Write a Rule and Assign a Host to a Folder or Cluster

vSphere Auto Deploy can assign a host to a folder or cluster. When the host boots, vSphere Auto Deploy adds it to the specified location on the vCenter Server. Hosts assigned to a cluster inherit the cluster's host profile.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
 - Verify that the folder you select is in a data center or in a cluster. You cannot assign the host to a standalone top-level folder.
1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Define a rule in which hosts with certain attributes, for example a range of IP addresses, are assigned to a folder or a cluster.

```
New-DeployRule -Name testrule3 -Item "my_folder" -Pattern "vendor=Acme,Zven",
"ipv4=192.XXX.1.10-192.XXX.1.20"
```

This example passes in the folder by name. You can instead pass in a folder, cluster, or data center object that you retrieve with the `Get-Folder`, `Get-Cluster`, or `Get-Datacenter` cmdlet.

3. Add the rule to the rule set.

```
Add-DeployRule testrule3
```

By default, the working rule set becomes the active rule set, and any changes to the rule set become active when you add a rule. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

- Assign a host already provisioned with vSphere Auto Deploy to the new folder or cluster location by performing test and repair compliance operation. See [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to add them to the specified vCenter Server location.

Configure a Stateless System by Running a Custom Script

You can use vSphere Auto Deploy to configure one or more hosts by associating custom scripts with a vSphere Auto Deploy rule.

- Verify that the script bundle you want to associate with a vSphere Auto Deploy rule is in `.tgz` format, with a maximum size of 10 MB, and written in Python or BusyBox ash scripting language.

The scripts run in alphabetical order after the initial ESXi boot workflow of the host.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Run the `Add-ScriptBundle` cmdlet to add the script bundle that contains the necessary scripts to the vSphere Auto Deploy inventory.

```
Add-ScriptBundle c:/temp/MyScriptBundle.tgz
```

The name of the script bundle without the `.tgz` extension is the name identifier or object of the script bundle item. You can update an existing script bundle by using the `-Update` parameter with the `Add-ScriptBundle` cmdlet.

3. Optional: Run the `Get-ScriptBundle` cmdlet to verify that the script bundle is added to the vSphere Auto Deploy inventory.
4. Define a rule in which hosts with certain attributes, for example a range of IP addresses, are assigned to the script bundle.

```
New-DeployRule -Name "testrule4" -Item "MyScriptBundle" -Pattern "ven-  
dor=Acme,Zven", "ipv4=192.XXX.1.10-192.XXX.1.20"
```

Double quotes are required if a name contains spaces, optional otherwise. Specify `-AllHosts` instead of a pattern to apply the item to all hosts.

You create a rule named `testrule4`. The rule assigns the script bundle named `My Script Bundle` to all hosts with a vendor of `Acme` or `Zven` that also have an IP address in the specified range. You can use the name identifier of the script bundle or the object returned by the `Get-ScriptBundle` cmdlet to identify the script bundle you want to associate with the rule.

5. Add the rule to the rule set.

```
Add-DeployRule testrule4
```

By default, the rule is added to both the working rule set and the active rule set. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

- For hosts already provisioned with vSphere Auto Deploy, perform the compliance testing and repair operations to provision them with the new scripts. See [Test and Repair Rule Compliance](#).
- Turn on unprovisioned hosts to provision them with the new scripts.

Test and Repair Rule Compliance

Test new or modified rules for compliance and repair accordingly, as changes in the vSphere Auto Deploy rule set are not updated automatically.

When you add a rule to the vSphere Auto Deploy rule set or modify one or more rules, hosts are not updated automatically. vSphere Auto Deploy applies the new rules only when you test their rule compliance and perform remediation.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).

- Verify that your infrastructure includes one or more ESXi hosts provisioned with vSphere Auto Deploy, and that the host on which you installed PowerCLI can access those ESXi hosts.
1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Use PowerCLI to check which vSphere Auto Deploy rules are currently available.

```
Get-DeployRule
```

The system returns the rules and the associated items and patterns.

3. Modify one of the available rules.

For example, you can change the image profile and the name of the rule.

```
Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile
```

You cannot edit a rule already added to the active rule set. Instead, you can copy the rule and replace the item or pattern you want to change.

4. Verify that you can access the host for which you want to test rule set compliance.

```
Get-VMHost -Name MyEsxi42
```

5. Run the cmdlet that tests rule set compliance for the host, and bind the return value to a variable for later use.

```
$tr = Test-DeployRuleSetCompliance MyEsxi42
```

6. Examine the differences between the contents of the rule set and configuration of the host.

```
$tr.itemlist
```

If the host for which you want to test the new rule set compliance is compliant with the active rule set, the system returns a table of current and expected items.

CurrentItem	ExpectedItem
-----	-----
My Profile 25MyNewProfile	

7. Remediate the host to use the revised rule set the next time you boot the host.

```
Repair-DeployRuleSetCompliance $tr
```

If the rule you changed specified the inventory location, the change takes effect when you repair compliance. For all other changes, reboot your host to have vSphere Auto Deploy apply the new rule and to achieve compliance between the rule set and the host.

Register a Caching Proxy Server Address with vSphere Auto Deploy

Simultaneously booting large number of stateless hosts places a significant load on the vSphere Auto Deploy server. You can load balance the requests between the vSphere Auto Deploy server and one or more proxy servers that you register with vSphere Auto Deploy.

- Prepare your system for vSphere Auto Deploy. For more information, see *vSphere Installation and Setup*.
1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Register a caching proxy server addresses with vSphere Auto Deploy by running the `Add-ProxyServer` cmdlet.

```
Add-ProxyServer -Address 'https://proxy_server_ip_address:port_number'
```

You can run the cmdlet multiple times to register multiple proxy servers. The address can contain a port number.

3. Optional: Run the `List-ProxyServer` cmdlet to verify that the caching proxy server is registered with vSphere Auto Deploy.

vSphere Auto Deploy Tasks by Using the vSphere Client

Managing vSphere Auto Deploy with the vSphere Client

You can add ESXi hosts to the vSphere Auto Deploy inventory, create, monitor, and manage rules, and host associations by using the vSphere Client.

Overview of the vSphere Auto Deploy Process by Using the vSphere Client

Getting started with vSphere Auto Deploy requires that you learn how vSphere Auto Deploy works, start the vSphere Auto Deploy and vSphere ESXi Image Builder vCenter Server services, create deploy rules that provision hosts, and power on your hosts to be booted with the image profile you specify.

The workflow for provisioning the hosts in your environment with vSphere Auto Deploy includes the following tasks:

1. Deploy vCenter Server.
The vSphere Auto Deploy server is included.
2. Configure the vSphere Auto Deploy and vSphere ESXi Image Builder service startup types.
See [Prepare Your System for vSphere Auto Deploy](#) and [Configure the vSphere ESXi Image Builder](#).
3. Add or import a software depot to the vSphere Auto Deploy inventory.
See [Add a Software Depot](#) or [Import a Software Depot](#).
4. (Optional) If you want to create a custom image profile, clone, or create an image profile by using the vSphere Client.
See [Clone an Image Profile](#) or [Create an Image Profile](#).
5. Create a deploy rule that assigns the image profile to one host, to multiple hosts specified by a pattern, or to all hosts.
See [Create a Deploy Rule](#).
Note: vSphere Auto Deploy is optimized for provisioning hosts that have a fixed MAC address to IP address mapping in DHCP (sometimes called DHCP reservations). If you want to use static IP addresses, you must set up the host profile to prompt for host customization. For more information, see the *vSphere Host Profiles* documentation.
6. Power on the hosts that you want to provision.
7. Set up the host you provisioned as a reference host for your host profile.
You can specify the reference host syslog settings, firewall settings, storage, networking, and so on.
8. Extract a host profile from the reference host.
See the *Host Profiles* documentation.
9. To provision multiple hosts with the host profile, clone or edit the previously created rule by using the vSphere Client.
See [Clone a Deploy Rule](#) or [Edit a Deploy Rule](#).
10. Activate the new rule and deactivate the old one.
See [Activate, Deactivate, and Reorder Deploy Rules](#).
11. Remediate the host associations to apply the new rule to the host.
See [Remediate a Non-compliant Host](#).

12. Verify that the hosts you provisioned meet the following requirements.

- Each host is connected to the vCenter Server system.
- The hosts are not in maintenance mode.
- The hosts have no compliance failures.
- Each host with a host profile that requires user input has up-to-date host customization information.

Remediate host associations and compliance problems and reboot hosts until all hosts meet the requirements.

Read for an introduction to the boot process, differences between first and subsequent boots, and an overview of using host customization.

Create a Deploy Rule

Before you provision ESXi hosts with vSphere Auto Deploy, you must create rules that assign host locations, image and host profiles to the hosts.

An ESXi host can match more than one vSphere Auto Deploy rule criteria, when this is the case, the rule order is considered.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
 - If you want to include an image profile to the rule, verify that the software depot you need is added to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).
1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere Auto Deploy service.
 2. On the **Deploy Rules** tab, click **New Deploy Rule**.
The **New Deploy Rule** wizard appears.
 3. On the **Name and hosts** page of the wizard, enter a name for the new rule.
 4. Select to either apply the rule to all hosts in the inventory or only to hosts that match a specific pattern.
You can select one or more patterns.
For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.
 5. On the **Configuration** page of the wizard, you can optionally include items in the rule.
Each enabled item adds a new page to the wizard.

Option	Action
Host Location	Add the hosts that match the criteria of the rule to a specific location.
Image Profile	Assign an image profile to the hosts that match the rule criteria.
Host Profile	Assign a host profile to the hosts that match the rule criteria.
Script Bundle	Assign a script bundle to the host that match the rule criteria.

6. Optional: On the **Select host location** page of the wizard, select a data center, folder, or cluster as host location for the hosts that match the rule.
7. Optional: On the **Select image profile** page of the wizard, use the drop-down menu to select a software depot and choose an image profile from the list.
If you want to bypass the acceptance level verification for the image profile, select the **Skip image profile signature check** check box.

8. Optional: On the **Select host profile** page of the wizard, select a host profile from the list.
9. Optional: On the **Select script bundle** page of the wizard, select a script bundle from the list.
10. On the **Ready to complete** page, review the summary information for the new rule.

You can view the newly created rule listed on the **Deploy Rules** tab.

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- Edit a vSphere Auto Deploy rule. See [Edit a Deploy Rule](#).
- Clone a vSphere Auto Deploy rule. See [Clone a Deploy Rule](#).
- View the host location, image profile, host profile, and added script bundles. See [View Host Associations](#).
- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).
- Change the image profile association of a host. See [Edit the Image Profile Association of a Host](#).

Clone a Deploy Rule

You can use a vSphere Auto Deploy rule as a template and modify only parts of the rule instead of creating a new one.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- If you want to include an image profile to the rule, verify that the software depot you need is added to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

You can clone an existing vSphere Auto Deploy rule by using the Clone Deploy Rule wizard.

1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere Auto Deploy service.
2. On the **Deploy Rules** tab, select a rule from the list.
3. Click **Clone**.
The **Clone Deploy Rule** wizard appears.
4. On the **Name and hosts** page of the wizard, enter a name for the new rule.
5. Select to either apply the rule to all hosts in the inventory or only to hosts that match a specific pattern.
You can select one or more patterns.
For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.
6. On the **Configuration** page of the wizard, you can optionally include items in the rule.
Each enabled item adds a new page to the wizard.

Option	Action
Host Location	Add the hosts that match the criteria of the rule to a specific location.
Image Profile	Assign an image profile to the hosts that match the rule criteria.
Host Profile	Assign a host profile to the hosts that match the rule criteria.
Script Bundle	Assign a script bundle to the host that match the rule criteria.

7. On the **Select host location** page of the wizard, select a location for the hosts that match the rule.

Option	Action
If you want to keep the host location used in the cloned rule	Select the Same Host location check box.
If you want to select a new location for the selected hosts	1. Select the Browse for Host location check box. 2. Select a data center, folder, or cluster as host location. 3. Click Next.

8. On the **Select image profile** page of the wizard, select an image profile.

Option	Action
If you do not want to change the image profile	Select the Same image profile check box.
If you want to assign a new image profile to the selected hosts	1. Select the Browse for Image Profile check box. 2. Select a software depot from the drop-down menu. 3. Select an image profile from the list. 4. (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

9. On the **Select host profile** page of the wizard, select a host profile.

Option	Action
If you want to keep the host profile used in the cloned rule	Select the Same Host profile check box.
If you want to assign a new host profile to the selected hosts	1. Select the Browse for Host Profile check box. 2. Select a host profile from the list and click Next.

10. On the **Select script bundle** page of the wizard, select a script bundle from the list.

11. On the **Ready to complete** page, review the summary information for the new rule.

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- Edit a vSphere Auto Deploy rule. See [Edit a Deploy Rule](#).

Edit a Deploy Rule

You can edit the name of an inactive Auto Deploy rule, its matching hosts, the host location, the image profile, and the host profile.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Deploy Rules** tab, from the list of rules in the inventory select the rule and click **Edit**.

The Edit Deploy Rule dialog box appears.

3. Optional: On the **Name and hosts** page of the wizard, enter a new name for the rule.

4. Select to either apply the rule to all hosts in the inventory or only to hosts that match a specific pattern.

You can select one or more patterns.

For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.

5. On the **Configuration** page of the wizard, you can optionally include items in the rule.
Each enabled item adds a new page to the wizard.

Option	Action
Host Location	Add the hosts that match the criteria of the rule to a specific location.
Image Profile	Assign an image profile to the hosts that match the rule criteria.
Host Profile	Assign a host profile to the hosts that match the rule criteria.
Script Bundle	Assign a script bundle to the host that match the rule criteria.

6. On the **Select host location** page of the wizard, select a location for the hosts that match the rule.

Option	Action
If you want to keep the host location used in the cloned rule	Select the Same Host location check box.
If you want to select a new location for the selected hosts	1. Select the Browse for Host location check box. 2. Select a data center, folder, or cluster as host location. 3. Click Next.

7. On the **Select image profile** page of the wizard, select an image profile.

Option	Action
If you do not want to change the image profile	Select the Same image profile check box.
If you want to assign a new image profile to the selected hosts	1. Select the Browse for Image Profile check box. 2. Select a software depot from the drop-down menu. 3. Select an image profile from the list. 4. (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

8. On the **Select host profile** page of the wizard, select a host profile.

Option	Action
If you want to keep the host profile used in the cloned rule	Select the Same Host profile check box.
If you want to assign a new host profile to the selected hosts	1. Select the Browse for Host Profile check box. 2. Select a host profile from the list and click Next.

9. On the **Select script bundle** page of the wizard, select a script bundle from the list.

10. On the **Ready to complete** page, review the summary information for the new rule.

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- Clone a vSphere Auto Deploy rule. See [Clone a Deploy Rule](#).

Activate, Deactivate, and Reorder Deploy Rules

After you create a vSphere Auto Deploy rule, the rule is in inactive state. You must activate the rule for it to take effect.

You can use the Activate and Reorder wizard to activate, deactivate, and change the order of the rules.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).

- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).

The upper list on the **Activate and Reorder** page of the wizard displays the rules in the active rule set. The lower list displays the inactive rules.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Deploy Rules** tab, click **Activate/Deactivate rules**.

The Activate and Reorder wizard appears.

3. Optional: If you want to deactivate an active rule, select the rule from the active rules list and click the **Deactivate** button.
4. From the list of inactive rules, select the rule that you want to activate and click the **Activate** button.
5. Optional: If you want to reorder the rules in the active rule list, select a rule that you want to move up or down in the list and click **Move up** or **Move down** above the list of active rules.
The rules are listed by priority. For example, if two or more rules apply to the same host but are set to provision the host with different host locations, image profiles, and host profiles, the rule that is highest in the list takes effect on the host.
6. Optional: If you want to test an inactive rule before activation, click **Test rules before activation**.
 - a) Select a host from the list and click **Check Compliance** to view the current status of the host and the changes that are expected after the activation of the rule.
If the host is compliant with the rule, you do not need to remediate the host after you activate the rule.
 - b) Optional: If you want to remediate the selected hosts after the rule activation, enable the toggle button or select the **Remediate all host associations after rule activation** check box to remediate all hosts.
7. Review the list of active rules and click **OK**.

On the **Deploy Rules** tab, the rule is listed as active in the Status column.

- View the host location, image profile, host profile, and added script bundles. See [View Host Associations](#).
- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).

View Host Associations

Some ESXi hosts in the vSphere Auto Deploy inventory might not be compliant with the active deploy rules and you must check for compliance.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. Check the host associations compliance.

The **Check Host Associations Compliance** window displays the status of the host and whether the host is compliant with the active rule set. You can view the currently assigned host location, image profile, host profile, script bundle, and the associations that will take effect after a remediation of the host.

Option	Steps
If you want to check the host associations compliance of a single host	1. On the Deployed Hosts tab, select an ESXi host. 2. Click Check Host Associations Compliance.

Option	Steps
	3. Check if the host associations are compliant with the current active rule set. 4. If you want to remediate the host, click Remediate. 5. Close the Check Host Associations Compliance window.
If you want to check the host associations compliance of multiple hosts	1. On the Deployed Hosts tab, select multiple ESXi hosts. 2. Click Check Host Associations Compliance. 3. Confirm that you want to check the compliance of all selected hosts. 4. Review the compliance status of the hosts in the left pane. 5. (Optional) Select a host to view the compliance status details. 6. (Optional) Select a host and click Remediate. 7. (Optional) Select the Remediate all host associations after rule activation check box to remediate all hosts. 8. Close the Check Host Associations Compliance window.

- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).
- Edit the image profile association of a host. See [Edit the Image Profile Association of a Host](#).
- Edit a vSphere Auto Deploy rule. See [Edit a Deploy Rule](#).

Edit the Image Profile Association of a Host

You can edit the image profile association of a single host if the host is not associated with a vSphere Auto Deploy rule.

Alternatively, you might not want to change the image profile association of multiple hosts by editing a rule.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Deployed Hosts** tab, select an ESXi host.

3. Click **Edit Image Profile Association**.

The Edit Image Profile Association dialog box appears.

4. Edit the image profile association of the host.

Option	Action
If you do not want to change the image profile	Select the Same image profile check box.
If you want to assign a new image profile to the selected hosts	1. Select the Browse for Image Profile check box. 2. Select a software depot from the drop-down menu. 3. Select an image profile from the list. 4. (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

5. Click **OK**.

The new image profile is listed in the Associated Image Profile column after a refresh of the page.

- View the host location, image profile, host profile, and added script bundles. See [View Host Associations](#).

- If the host is associated with a rule and you want to revert to the image profile defined in the rule, remediate the host. See [Remediate a Non-compliant Host](#).

Remediate a Non-compliant Host

Remediate the ESXi host associations when you add a rule to the vSphere Auto Deploy active rule set or make changes to one or more rules.

When you add a rule to the vSphere Auto Deploy active rule set or make changes to one or more rules, hosts are not updated automatically. You must remediate the host associations to apply the new rules to the host.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- Create a vSphere Auto Deploy rule. See [Create a Deploy Rule](#).
- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- If the remediation of a host, results in a change in its location, the host must be placed in maintenance mode.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Deployed Hosts** tab, select a single or multiple ESXi hosts.

3. Click **Remediate Host Associations**.

If you remediate a host that has an edited image profile association, the host reverts to the settings defined in the rule that it matches.

You can monitor the progress of the remediation process in the Recent Tasks pane.

- View the host location, image profile, host profile, and added script bundles. See [View Host Associations](#).
- Change the image profile association of a host. See [Edit the Image Profile Association of a Host](#).

Add a Host to the vSphere Auto Deploy Inventory

You can view the hosts that do not match any vSphere Auto Deploy rule and manually add a host to the vSphere Auto Deploy inventory.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
- To assign an image profile to the host, add the software depot that you need to the inventory. See [Add a Software Depot](#) or [Import a Software Depot](#).

To add a host to the current vSphere Auto Deploy inventory of deployed hosts, you can create a new rule or edit an existing rule to include a host that is not deployed with vSphere Auto Deploy and associate it with a specific host location, image profile, host profile, and script bundle. Alternatively, you can manually add a host to the inventory by assigning it a host location, image profile, host profile, and script bundle.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Discovered Hosts** tab, select one or more hosts that you want to provision with a host location, image profile, and host profile.

3. Select **Add to Inventory**.

Alternatively, click **Remove** to dismiss the selected hosts from the **Discovered Hosts** tab.

The Add to Inventory wizard appears.

4. On the **Select host location** page of the wizard, select a data center, folder, or cluster as host location for the hosts that match the rule.
5. On the **Select image profile** page of the wizard, use the drop-down menu to select a software depot and choose an image profile from the list.
If you want to bypass the acceptance level verification for the image profile, select the **Skip image profile signature check** check box.
6. On the **Select host profile** page of the wizard, select a host profile from the list.
7. On the **Select host profile** page of the wizard, use the **Filter** to search the host profiles list or select the **Do not include a host profile** check box to continue without adding a host profile.
8. On the **Select script bundle** page of the wizard, select a script bundle from the list.
9. On the Ready to complete page, review the selected host associations.
 - Edit a vSphere Auto Deploy rule. See [Edit a Deploy Rule](#).
 - Clone a vSphere Auto Deploy rule. See [Clone a Deploy Rule](#).
 - View the host location, image profile, host profile, and added script bundles. See [View Host Associations](#).
 - Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).

Add a Host to a Cluster That Uses a Single Image

Create a rule in Auto Deploy that assigns a cluster that you manage by a single image as the host location for newly added ESXi hosts.

To add ESXi hosts to a cluster that you manage by a single image, you create a rule in Auto Deploy that assigns such a cluster as the host location for newly added hosts. Each host inherits the same image, which enables hardware compatibility checking, cluster-wide remediation, and easier upgrades.

- Prepare your system and install the Auto Deploy Server. For more information, see *Prepare Your System for vSphere Auto Deploy* in the *VMware ESXi Installation and Setup* documentation.
- Verify that each host is version 8.0 or later.
- Verify that all hosts in the cluster are stateful and have a physical storage attached.
- Verify that solutions, which are not integrated with vSphere Lifecycle Manager are not enabled for the cluster.

By creating an Auto Deploy rule, where the host target location is a cluster managed by an image, you can transition stateful ESXi hosts to the cluster. Based on host identification mechanisms, the rule adds the hosts to the target cluster.

Such Auto Deploy rule is not allowed to contain an Image Profile or a Host Profile, as the image specification and configuration of hosts in the target cluster are created automatically.

1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere Auto Deploy service.
2. On the **Deploy Rules** tab, click **New Deploy Rule**.
The **New Deploy Rule** wizard appears.
3. On the **Name and hosts** page of the wizard, enter a name for the new rule.
4. Select to either apply the rule to all hosts in the inventory or only to hosts that match a specific pattern.
You can select one or more patterns.
For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.

5. On the **Configuration** page of the wizard, select the **Host Location** check box, and click **Next**.
You can optionally include a script bundle to the ESXi hosts that match the rule criteria by selecting the **Script Bundle** check box.
Note: The image profile and host profile attached to the target cluster are used.
6. On the **Select host location** page of the wizard, select a cluster that uses a single image.
7. On the **Select script bundle** page of the wizard, select a script bundle from the list.
8. On the **Ready to complete** page, review the summary information for the new rule.
9. Click **Finish**.

You can view the newly created rule listed on the **Deploy Rules** tab.

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- For more information on stateless caching and stateful installs, see [Use vSphere Auto Deploy for Stateless Caching and Stateful Installs](#).

Add a Host to a Cluster Managed by vSphere Configuration Profiles

Create a rule in Auto Deploy that assigns newly added hosts to a cluster that manages ESXi configuration collectively with vSphere Configuration Profiles.

To add ESXi hosts to a cluster that manages ESXi configuration at a cluster level, you create a rule in Auto Deploy that assigns such a cluster as the host location for newly added hosts, which inherit the same settings and do not require manual configuration.

- Prepare your system and install the Auto Deploy Server. For more information, see [Install and Configure vSphere Auto Deploy](#).
- Verify that each ESXi host is version 8.0 or later.

Note:

- You cannot provision bare-metal servers by using AutoDeploy if the desired configuration for a cluster managed by vSphere Configuration Profiles contains vSphere Distributed Switch or NSX settings. You can add vSphere Distributed Switch settings to the desired configuration of a cluster after you provision and add hosts, but both your vCenter Server instance and all ESXi hosts in the cluster must be of version 8.0 Update 3 or later.

By creating an Auto Deploy rule, where the host target location is a cluster in which you manage the configuration of all hosts collectively by using vSphere Configuration Profiles, you remove the need to use Host Profiles, or make any manual configurations, while keeping the flexibility to define custom settings per host or override the cluster-level settings for a group of hosts. For more information on how to set up clusters that you manage with a configuration at a cluster level, and vSphere Configuration Profiles, see [Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level](#).

Note: Once you set up a cluster that manages ESXi configuration at a cluster level, you cannot roll back to using Host Profiles or single images, and you cannot use the Quickstart option for such clusters. However, you can switch from a cluster that you manage with a single image to a cluster that you manage with a configuration at a cluster level just by selecting the **Set Up Host Settings** option under **Configure > Desired State > Host Settings**.

1. Navigate to **Home > Auto Deploy**.
By default, only the administrator role has privileges to use the vSphere Auto Deploy service.
2. On the **Deploy Rules** tab, click **New Deploy Rule**.
The **New Deploy Rule** wizard appears.

3. On the **Name and hosts** page of the wizard, enter a name for the new rule.
4. Select to either apply the rule to all hosts in the inventory or only to hosts that match a specific pattern.
You can select one or more patterns.
For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.
5. On the **Configuration** page of the wizard, select the **Host Location** check box, and click **Next**.
You can ignore the **Image**, **Host Settings**, and **Script Bundle** check boxes. Image specification and configuration of hosts in the target cluster, and any post-installation settings are created automatically. Even if you select any of the check boxes, the rule ignores the selections.
6. On the **Select host location** page of the wizard, select a cluster that manages ESXi configuration at a cluster level.
7. On the **Ready to complete** page, review the summary information for the new rule.
8. Click **Finish**.

You can view the newly created rule listed on the **Deploy Rules** tab.

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- For more information on stateless caching and stateful installs, see [Use vSphere Auto Deploy for Stateless Caching and Stateful Installs](#).

Working with Script Bundles

You can add a custom script for additional post-deployment host configuration. The script runs after you provision an ESXi host with Auto Deploy. For example, you can create a custom ESXi firewall rule and other configurations not available with Host Profiles.

- Verify that you can run the script in the ESXi Shell.

Since vSphere 6.7 Update 1, you can add or remove a custom script by using the vSphere Client. A script bundle can include multiple scripts and must be delivered as a single compressed file with the `.tgz` extension. After uploaded to the vCenter Server, you can include the script bundle to an Auto Deploy rule.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. Select the **Script Bundles** tab.
3. Click **Upload**.
4. Browse to a script bundle file and select **Upload**.

The script is present in the **Script Bundles** list.

5. Optional: Select a script bundle, click **Remove**, and confirm the selection.
The script bundle is deleted from the list.

- Activate a vSphere Auto Deploy rule. See [Activate, Deactivate, and Reorder Deploy Rules](#).
- Edit a vSphere Auto Deploy rule. See [Edit a Deploy Rule](#).
- Clone a vSphere Auto Deploy rule. See [Clone a Deploy Rule](#).
- View the host location, image profile, host profile, and added script bundles. See [View Host Associations](#).
- Remediate non-compliant hosts. See [Remediate a Non-compliant Host](#).
- Change the image profile association of a host. See [Edit the Image Profile Association of a Host](#).

Download vSphere Auto Deploy Logs

You can use the vSphere Auto Deploy logging information from the vSphere Client to resolve problems that you encounter with vSphere Auto Deploy.

Use the vSphere Client to log in to the vCenter Server instance that vSphere Auto Deploy is registered with.

1. Navigate to **Home > Administration** and select **Deployment > System Configuration**.
2. Select one of the nodes for which you want to retrieve a support bundle. The support bundle holds the services logs.
3. Click **Export Support Bundle**.
4. Select only **VirtualAppliance > Auto Deploy**.
5. Click the **Export Support Bundle** button to download the log files.

Start, Stop, or Restart the vSphere Auto Deploy Service

You can start, stop, or restart the Auto Deploy service in the vCenter Server Management Interface.

Verify that you have a root access to the vCenter Server Management Interface.

To start, stop, and restart services in vCenter Server, you use the vCenter Server Management Interface.

1. Log in to the vCenter Server Management Interface, <https://IP-address-or-FQDN:5480>.
2. Click **Services**.
The **Services** pane displays a table of all installed services. You can sort them by name, startup type, health, and state.
3. Select the **Auto Deploy** service and select your action.
The available actions depend on whether the Auto Deploy service is already running or not.
 - Click **Restart** to restart the service.
Restarting the service requires confirmation and might lead to the Auto Deploy functionality becoming temporarily unavailable.
 - Click **Start** to start the service.
 - Click **Stop** to stop the service.
Stopping the service requires confirmation.

Provision ESXi Hosts with vSphere Auto Deploy

Use vSphere Auto Deploy to provision or reprovision hundreds of physical hosts with ESXi software.

You can use vSphere Auto Deploy to provision hundreds of physical hosts with ESXi software either for the first time (first boot), or reboot hosts, or reprovision hosts with a different image profile, host profile, custom script, or folder, or cluster location. You can also select to provision the hosts with an image profile that does not contain VMware Tools binaries.

Provisioning ESXi Systems with vSphere Auto Deploy

vSphere Auto Deploy can provision hundreds of physical hosts with ESXi software, either for first boot, reboot, or reprovisioning.

You can provision hosts that did not previously run ESXi software (first boot), reboot hosts, or reprovision hosts with a different image profile, host profile, custom script, or folder, or cluster location. The vSphere Auto Deploy process differs depending on the state of the host and on the changes that you want to make.

vSphere Auto Deploy Boot Process

When you boot a host that you want to provision or reprovision with vSphere Auto Deploy, the vSphere Auto Deploy infrastructure supplies the image profile and, optionally, a host profile, a vCenter Server location, and script bundle for that host.

The boot process is different for hosts that have not yet been provisioned with vSphere Auto Deploy (first boot) and for hosts that have been provisioned with vSphere Auto Deploy and added to a vCenter Server system (subsequent boot).

First Boot Prerequisites

Before a first boot process, you must set up your system. Setup includes the following tasks, which are discussed in more detail in [Install and Configure vSphere Auto Deploy](#).


- Set up a DHCP server that assigns an IP address to each host upon startup and that points the host to the TFTP server to download the iPXE boot loader from.
- If the hosts that you plan to provision with vSphere Auto Deploy are with legacy BIOS, verify that the vSphere Auto Deploy server has an IPv4 address. PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.
- Identify an image profile to be used in one of the following ways.
 - Choose an ESXi image profile in a public depot.
 - Create a custom image profile by using vSphere ESXi Image Builder, and place the image profile in a depot that the vSphere Auto Deploy server can access. The image profile must include a base ESXi VIB.
- If you have a reference host in your environment, export the host profile of the reference host and define a rule that applies the host profile to one or more hosts. See [Setting Up a vSphere Auto Deploy Reference Host](#).
- Specify rules for the deployment of the host and add the rules to the active rule set.

First Boot Overview

When a host that has not yet been provisioned with vSphere Auto Deploy boots (first boot), the host interacts with several vSphere Auto Deploy components.

1. When the administrator turns on a host, the host starts a PXE boot sequence.
The DHCP Server assigns an IP address to the host and instructs the host to contact the TFTP server.
2. The host contacts the TFTP server and downloads the iPXE file (executable boot loader) and an iPXE configuration file.
3. iPXE starts executing.
The configuration file instructs the host to make a HTTP boot request to the vSphere Auto Deploy server. The HTTP request includes hardware and network information.
4. In response, the vSphere Auto Deploy server performs these tasks:
 - a. Queries the rules engine for information about the host.
 - b. Streams the components specified in the image profile, the optional host profile, and optional vCenter Server location information.
5. The host boots using the image profile.
If the vSphere Auto Deploy server provided a host profile, the host profile is applied to the host.
6. vSphere Auto Deploy adds the host to the vCenter Server system that vSphere Auto Deploy is registered with.
 - a. If a rule specifies a target folder or cluster on the vCenter Server system, the host is placed in that folder or cluster. The target folder must be under a data center.
 - b. If no rule exists that specifies a vCenter Server inventory location, vSphere Auto Deploy adds the host to the first datacenter displayed in the vSphere Client UI.
7. If the host profile requires the user to specify certain information, such as a static IP address, the host is placed in maintenance mode when the host is added to the vCenter Server system.
You must reapply the host profile and update the host customization to have the host exit maintenance mode. When you update the host customization, answer any questions when prompted.
8. If the host is part of a DRS cluster, virtual machines from other hosts might be migrated to the host after the host has successfully been added to the vCenter Server system.

See [Provision a Host \(First Boot\)](#).

Figure 18: vSphere Auto Deploy Installation, First Boot**Subsequent Boots Without Updates**

For hosts that are provisioned with vSphere Auto Deploy and managed by vCenter Server, subsequent boots can become completely automatic.

1. The administrator reboots the host.
2. As the host boots up, vSphere Auto Deploy provisions the host with its image profile and host profile.
3. Virtual machines are brought up or migrated to the host based on the settings of the host.
 - Standalone host. Virtual machines are powered on according to autostart rules defined on the host.
 - DRS cluster host. Virtual machines that were successfully migrated to other hosts stay there. Virtual machines for which no host had enough resources are registered to the rebooted host.

If vCenter Server is unavailable, the host contacts the vSphere Auto Deploy server and is provisioned with an image profile. The host continues to contact the vSphere Auto Deploy server until vSphere Auto Deploy reconnects to the vCenter Server.

vSphere Auto Deploy cannot set up vSphere distributed switches if vCenter Server is unavailable, and virtual machines are assigned to hosts only if they participate in an vSphere HA cluster. Until the host is reconnected to vCenter Server and the host profile is applied, the switch cannot be created. Because the host is in maintenance mode, virtual machines cannot start. See [Reprovision Hosts with Simple Reboot Operations](#).

Any hosts that are set up to require user input are placed in maintenance mode. See [Update the Host Customization in the vSphere Client](#).

Subsequent Boots With Updates

You can change the image profile, host profile, vCenter Server location, or script bundle for hosts. The process includes changing rules and testing and repairing the host's rule compliance.


1. The administrator uses the `Copy-DeployRulePowerCLI` cmdlet to copy and edit one or more rules and updates the rule set. See [Overview of the vSphere Auto Deploy Process by Using PowerCLI](#) for an example.
2. The administrator runs the `Test-DeployRulesetCompliance` cmdlet to check whether each host is using the information that the current rule set specifies.

3. The host returns a PowerCLI object that encapsulates compliance information.
4. The administrator runs the `Repair-DeployRulesetCompliance` cmdlet to update the image profile, host profile, or vCenter Server location the vCenter Server system stores for each host.
5. When the host reboots, it uses the updated image profile, host profile, vCenter Server location, or script bundle for the host.

If the host profile is set up to request user input, the host is placed in maintenance mode. Follow the steps in [Update the Host Customization in the vSphere Client](#).

See [Test and Repair Rule Compliance](#).

Figure 19: vSphere Auto Deploy Installation, Subsequent Boots



Provisioning of Systems that Have Distributed Switches

You can configure the host profile of a vSphere Auto Deploy reference host with a distributed switch.

When you configure the distributed switch, the boot configuration parameters policy is automatically set to match the network parameters required for host connectivity after a reboot.

When vSphere Auto Deploy provisions the ESXi host with the host profile, the host goes through a two-step process.

1. The host creates a standard virtual switch with the properties specified in the boot configuration parameters field.
2. The host creates the VMkernel NICs. The VMkernel NICs allow the host to connect to vSphere Auto Deploy and to the vCenter Server system.

When the host is added to vCenter Server, vCenter Server removes the standard switch and reapplies the distributed switch to the host.

Note: Do not change the boot configuration parameters to avoid problems with your distributed switch.

Provision a Host (First Boot)

Provisioning a host that has never been provisioned with vSphere Auto Deploy (first boot) differs from subsequent boot processes. You must prepare the host and fulfill all other prerequisites before you can provision the host. You can optionally define a custom image profile with vSphere ESXi Image Builder by using the vSphere Client or PowerCLI cmdlets.

- Make sure your host meets the hardware requirements for ESXi hosts.

- Prepare the system for vSphere Auto Deploy. See [Install and Configure vSphere Auto Deploy](#).
- Write rules that assign an image profile to the host and optionally assign a host profile and a vCenter Server location to the host. See [Managing vSphere Auto Deploy with PowerCLI Cmdlets](#) or [Managing vSphere Auto Deploy with the vSphere Client](#).

When the setup is complete, the vSphere Auto Deploy service is enabled, DHCP setup is complete, and rules for the host that you want to provision are in the active rule set.

1. Turn on the host.

The host contacts the DHCP server and downloads iPXE from the location the server points it to. Next, the vSphere Auto Deploy server provisions the host with the image specified by the rule engine. The vSphere Auto Deploy server might also apply a host profile to the host if one is specified in the rule set. Finally, vSphere Auto Deploy adds the host to the vCenter Server system that is specified in the rule set.

2. Optional: If vSphere Auto Deploy applies a host profile that requires user input such as an IP address, the host is placed in maintenance mode. Reapply the host profile with the vSphere Client and provide the user input when prompted.

After the first boot process, the host is running and managed by a vCenter Server system. The vCenter Server stores the host's image profile, host profile, and location information.

You can now reboot the host as needed. Each time you reboot, the host is reprovisioned by the vCenter Server system.

Reprovision hosts as needed. See [Reprovisioning Hosts](#).

If you want to change the image profile, host profile, custom script, or location of the host, update the rules and activate them by using the vSphere Client or perform a test and repair compliance operation in a PowerCLI session. See [Rules and Rule Sets](#) or [Test and Repair Rule Compliance](#).

Reprovisioning Hosts

Use vSphere Auto Deploy to reprovision ESXi hosts with a different image profile or a different host profile.

vSphere Auto Deploy supports multiple reprovisioning options. You can perform a simple reboot or reprovision with a different image profile or a different host profile.

The following reprovisioning operations are available.

- Simple reboot.
- Reboot of hosts for which the user answered questions during the boot operation.
- Reprovision with a different image profile.
- Reprovision with a different host profile.

Reprovision Hosts with Simple Reboot Operations

You can reprovisions ESXi hosts with the image profile, host profile, custom script, and vCenter Server location assigned during first boot.

A simple reboot of a host that is provisioned with vSphere Auto Deploy requires only that all prerequisites are still met. The process uses the previously assigned image profile, host profile, custom script, and vCenter Server location.

- Verify that all associated items like are available. An item can be an image profile, host profile, custom script or vCenter Server inventory location.
- Verify that the host has the identifying information (asset tag, IP address) it had during previous boot operations.

1. Place the host in maintenance mode.

Host Type	Action
Host is part of a DRS cluster	VMware DRS migrates virtual machines to appropriate hosts when you place the host in maintenance mode.

Host Type	Action
Host is not part of a DRS cluster	You must migrate all virtual machines to different hosts and place each host in maintenance mode.

2. Reboot the host.

The host shuts down. When the host reboots, it uses the image profile that the vSphere Auto Deploy server provides. The vSphere Auto Deploy server also applies the host profile stored on the vCenter Server system.

Use PowerCLI To Reprovision a Host

You can use vSphere Auto Deploy to reprovision a host with a new image profile in a PowerCLI session.

- Verify that the setup you performed during the first boot operation is in place.

Several options for reprovisioning hosts exist.

- If the VIBs that you want to use support live update, you can use an `esxcli software vib update` command. In that case, you must also update the rule set to use an image profile that includes the new VIBs.
- During testing, you can apply an image profile to an individual host with the `Apply-EsxImageProfile` cmdlet and reboot the host so the change takes effect. The `Apply-EsxImageProfile` cmdlet updates the association between the host and the image profile but does not install VIBs on the host.
- In all other cases, use this procedure.

1. At the PowerShell prompt, run the `Connect-VIServerPowerCLI` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Determine the location of a public software depot that contains the image profile that you want to use, or define a custom image profile with vSphere ESXi Image Builder.
3. Run `Add-EsxSoftwareDepot` to add the software depot that contains the image profile to the PowerCLI session.

Depot Type	Cmdlet
Remote depot	Run <code>Add-EsxSoftwareDepot depot_url</code> .
ZIP file	1. Download the ZIP file to a local file path or create a mount point local to the PowerCLI machine. 2. Run <code>Add-EsxSoftwareDepot C:\file_path\my_offline_depot.zip</code>.

4. Run `Get-EsxImageProfile` to see a list of image profiles, and decide which profile you want to use.
5. Run `Copy-DeployRule` and specify the `ReplaceItem` parameter to change the rule that assigns an image profile to hosts.

The following cmdlet replaces the current image profile that the rule assigns to the host with the `my_new_imageprofile` profile. After the cmdlet completes, `myrule` assigns the new image profile to hosts. The old version of `myrule` is renamed and hidden.

```
Copy-DeployRule myrule -ReplaceItem my_new_imageprofile
```

6. Test the rule compliance for each host that you want to deploy the image to.

- a) Verify that you can access the host for which you want to test rule set compliance.

```
Get-VMHost -Name ESXi_hostname
```

- b) Run the cmdlet that tests rule set compliance for the host, and bind the return value to a variable for later use.

```
$str = Test-DeployRuleSetCompliance ESXi_hostname
```

- c) Examine the differences between the contents of the rule set and configuration of the host.

```
$str.itemlist
```

The system returns a table of current and expected items if the host for which you want to test the new rule set compliance is compliant with the active rule set.

CurrentItem	ExpectedItem
-----	-----
my_old_imageprofile	my_new_imageprofile

- d) Remediate the host to use the revised rule set the next time you boot the host.

```
Repair-DeployRuleSetCompliance $str
```

7. Reboot the host to provision it with the new image profile.

Reprovision a Host with a New Image Profile by Using the vSphere Client

You can use vSphere Auto Deploy to reprovision a host with a new image profile with the vSphere Client by changing the rule that the host corresponds to and activating the rule.

- Verify that the image profile you want to use to reprovision the host is available. See [Create an Image Profile](#).
- Verify that the setup you performed during the first boot operation is in place.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Deploy Rules** tab, from the list of rules in the inventory select the rule and click **Edit**.

The Edit Deploy Rule dialog box appears.

3. Optional: On the **Name and hosts** page of the wizard, enter a name for the new rule.

4. Optional: Select to either apply the rule to all hosts in the inventory or only to hosts that match a specific pattern.

You can select one or more patterns.

For example, the rule can apply only to hosts in a vCenter Single Sign-On domain, with a specific host name, or that match a specific IPv4 range.

5. On the **Configuration** page of the wizard, you can optionally include items in the rule.

Each enabled item adds a new page to the wizard.

Option	Action
Host Location	Add the hosts that match the criteria of the rule to a specific location.
Image Profile	Assign an image profile to the hosts that match the rule criteria.
Host Profile	Assign a host profile to the hosts that match the rule criteria.
Script Bundle	Assign a script bundle to the host that match the rule criteria.

6. Click **Next** to skip the Host Location selection.
7. On the **Select image profile** page of the wizard, assign an image profile to the hosts that match the rule criteria.

Option	Action
If you do not want to change the image profile	Select the Same image profile check box.
If you want to assign a new image profile to the selected hosts	1. Select the Browse for Image Profile check box. 2. Select a software depot from the drop-down menu. 3. Select an image profile from the list. 4. (Optional) If you want to bypass the acceptance level verification for the image profile, select the Skip image profile signature check check box.

8. Click **Next** to skip the Host profile selection.
9. On the **Ready to complete** page, review the summary information for the new image profile and click **Finish**.
10. Click **Activate/Deactivate rules**.
11. From the list of inactive rules, select the rule that you want to activate and click the **Activate** button.
12. Optional: If you want to reorder the rules in the active rule list, select a rule that you want to move up or down in the list and click **Move up** or **Move down** above the list of active rules.
The rules are listed by priority. For example, if two or more rules apply to the same host but are set to provision the host with different host locations, image profiles, and host profiles, the rule that is highest in the list takes effect on the host.
13. Optional: If you want to test an inactive rule before activation, click **Test rules before activation**.
 - a) Select a host from the list and click **Check Compliance** to view the current status of the host and the changes that are expected after the activation of the rule.
If the host is compliant with the rule, you do not need to remediate the host after you activate the rule.
 - b) Optional: If you want to remediate the selected hosts after the rule activation, enable the toggle button or select the **Remediate all host associations after rule activation** check box to remediate all hosts.
14. Review the list of active rules and click **OK**.
15. Reboot the host to provision it with the new image profile.

Update the Host Customization in the vSphere Client

If a host required user input during a previous boot, the answers are saved with the vCenter Server. If you want to prompt the user for new information, you must remediate the host.

Attach a host profile that prompts for user input to the host.

1. Migrate all virtual machines to different hosts, and place the host into maintenance mode.

Host Type	Action
Host is part of a DRS cluster	VMware DRS migrates virtual machines to appropriate hosts when you place the host in maintenance mode.
Host is not part of a DRS cluster	You must migrate all virtual machines to different hosts and place each host in maintenance mode.

2. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

3. On the **Deployed Hosts** tab, select an ESXi host.
4. Click **Remediate Host Associations**.
You can monitor the progress of the remediation process in the Recent Tasks pane.
5. When prompted, provide the user input.
6. Direct the host to exit maintenance mode.

The host customization is saved and takes effect the next time you boot the host.

Provision ESXi Host with an Image Profile Without VMware Tools

You can select to provision ESXi hosts with an image profile that does not contain VMware Tools binaries.

When you provision ESXi hosts with vSphere Auto Deploy, you can select to provision the hosts with an image profile that does not contain VMware Tools binaries. This image profile is smaller, has a lower memory overhead, and boots faster in a PXE-boot environment.

Download the `xxxxxx-no-tools` image profile from the VMware download site.

If the network boot time is too slow when using the standard image, or if you want to save space on the hosts, you can use an image profile that does not include VMware Tools and place the VMware Tools binaries on a shared storage.

1. Boot an ESXi host that was not provisioned with vSphere Auto Deploy.
2. Copy the `/productLocker` directory from the ESXi host to a shared storage.
You can connect to an ESXi host using an SSH client, see Knowledge Base article [1019852](#).
3. Change the `UserVars.ProductLockerLocation` variable to point to the new `/productLocker` directory location.
 - a) In the vSphere Client, select the reference host and click the **Configure** tab.
 - b) Under **System**, click **Advanced System Settings**.
 - c) Click **Edit**.
 - d) Filter the settings for `uservars`, and select **UserVars.ProductLockerLocation**.
 - e) Click the current value and edit the location so it points to the shared storage.
4. Create a host profile from the reference host.
5. Create a vSphere Auto Deploy rule that assigns the `xxxxxx-no-tools` image profile and host profile from the reference host to all other hosts.
6. Boot your target hosts with the rule so they pick up the product locker location from the reference host.

Use vSphere Auto Deploy for Stateless Caching and Stateful Installs

With stateless caching, you can cache the image of an ESXi host. With stateful installs, you can install hosts over the network.

The vSphere Auto Deploy stateless caching feature lets you cache the host's image. The vSphere Auto Deploy stateful installs feature lets you install hosts over the network. After the initial network boot, these hosts boot like other ESXi hosts. The stateless caching solution is primarily intended for situations when several hosts boot simultaneously. The locally cached image helps prevent a bottleneck that results if several hundreds of hosts connect to the vSphere Auto Deploy server simultaneously. After the boot operation is complete, hosts connect to vSphere Auto Deploy to complete the setup.

The stateful installs feature lets you provision hosts with the image profile over the network without having to set up the PXE boot infrastructure.

Introduction to Stateless Caching and Stateful Installs

You can use the System Cache Configuration host profile to provision hosts with vSphere Auto Deploy stateless caching and stateful installs.

Examples of Stateless Caching and Stateful Installs

Hosts provisioned with vSphere Auto Deploy cache the image (stateless caching)

Set up and apply a host profile for stateless caching. You can cache the image on a local disk, a remote disk, or a USB drive. Continue provisioning this host with vSphere Auto Deploy. If the vSphere Auto Deploy server becomes unavailable, for example because hundreds of hosts attempt to access it simultaneously, the host boots from the cache. The host attempts to reach the vSphere Auto Deploy server after the boot operation to complete configuration.

Hosts provisioned with vSphere Auto Deploy become stateful hosts

Set up and apply a host profile for stateful installs. When you provision a host with vSphere Auto Deploy, the image is installed on the local disk, a remote disk, or a USB drive. For subsequent boots, you boot from the disk. The host no longer uses vSphere Auto Deploy.

Preparation

To successfully use stateless caching or stateful installs, decide how to configure the system and set the boot order.

Table 20: Preparation for Stateless Caching or Stateful Installs

Requirement or Decision	Description
Decide on VMFS partition overwrite	When you install ESXi by using the interactive installer, you are prompted whether you want to overwrite an existing VMFS datastore. The System Cache Configuration host profile provides an option to overwrite existing VMFS partitions. The option is not available if you set up the host profile to use a USB drive.
Decide whether you need a highly available environment	If you use vSphere Auto Deploy with stateless caching, you can set up a highly available vSphere Auto Deploy environment to guarantee that virtual machines are migrated on newly provisioned hosts and that the environment supports vNetwork Distributed Switch even if the vCenter Server system becomes temporarily unavailable.
Set the boot order	The boot order you specify for your hosts depends on the feature you want to use. - To set up vSphere Auto Deploy with stateless caching, configure your host to first attempt to boot from the network, and to then attempt to boot from disk. If the vSphere Auto Deploy server is not available, the host boots using the cache. - To set up vSphere Auto Deploy for stateful installs on hosts that do not currently have a bootable disk, configure your hosts to first attempt to boot from disk, and to then attempt to boot from the network. Note: If you currently have a bootable image on the disk, configure the hosts for one-time PXE boot, and provision the host with vSphere Auto Deploy to use a host profile that specifies stateful installs.

Stateless Caching and Loss of Connectivity

If the ESXi hosts that run your virtual machines lose connectivity to the vSphere Auto Deploy server, the vCenter Server system, or both, some limitations apply the next time you reboot the host.

- If vCenter Server is available but the vSphere Auto Deploy server is unavailable, hosts do not connect to the vCenter Server system automatically. You can manually connect the hosts to the vCenter Server, or wait until the vSphere Auto Deploy server is available again.
- If both vCenter Server and vSphere Auto Deploy are unavailable, you can connect to each ESXi host by using the VMware Host Client, and add virtual machines to each host.
- If vCenter Server is not available, vSphere DRS does not work. The vSphere Auto Deploy server cannot add hosts to the vCenter Server. You can connect to each ESXi host by using the VMware Host Client, and add virtual machines to each host.
- If you make changes to your setup while connectivity is lost, the changes are lost when the connection to the vSphere Auto Deploy server is restored.

Understanding Stateless Caching and Stateful Installs

When you want to use vSphere Auto Deploy with stateless caching or stateful installs, you must set up a host profile, apply the host profile, and set the boot order.

When you apply a host profile that enables caching to a host, vSphere Auto Deploy partitions the specified disk. What happens next depends on how you set up the host profile and how you set the boot order on the host.

- vSphere Auto Deploy caches the image when you apply the host profile if **Enable stateless caching on the host** is selected in the System Cache Configuration host profile. No reboot is required. When you later reboot, the host continues to use the vSphere Auto Deploy infrastructure to retrieve its image. If the vSphere Auto Deploy server is not available, the host uses the cached image.
- vSphere Auto Deploy installs the image if **Enable stateful installs on the host** is selected in the System Cache Configuration host profile. When you reboot, the host initially boots using vSphere Auto Deploy to complete the installation. A reboot is then issued automatically, after which the host boots from disk, similar to a host that was provisioned with the installer. vSphere Auto Deploy no longer provisions the host.

You can apply the host profile from the vSphere Client, or write a vSphere Auto Deploy rule in a PowerCLI session that applies the host profile.

Using the vSphere Client to Set Up vSphere Auto Deploy for Stateless Caching or Stateful Installs

You can create a host profile on a reference host and apply that host profile to additional hosts or to a vCenter Server folder or cluster. The following workflow results.

1. You provision a host with vSphere Auto Deploy and edit that host's System Image Cache Configuration host profile.
2. You place one or more target hosts in maintenance mode, apply the host profile to each host, and instruct the host to exit maintenance mode.
3. What happens next depends on the host profile you selected.
 - If the host profile enabled stateless caching, the image is cached to disk. No reboot is required.
 - If the host profile enabled stateful installs, the image is installed. When you reboot, the host uses the installed image.

Using PowerCLI to Set Up vSphere Auto Deploy for Stateless Caching or Stateful Installs

You can create a host profile for a reference host and write a vSphere Auto Deploy rule that applies that host profile to other target hosts in a PowerCLI session. The following workflow results.

1. You provision a reference host with vSphere Auto Deploy and create a host profile to enable a form of caching.

2. You write a rule that provisions additional hosts with vSphere Auto Deploy and that applies the host profile of the reference host to those hosts.
3. vSphere Auto Deploy provisions each host with the image profile or by using the script bundle associated with the rule. The exact effect of applying the host profile depends on the host profile you selected.
 - For stateful installs, vSphere Auto Deploy proceeds as follows:
 - During first boot, vSphere Auto Deploy installs the image on the host.
 - During subsequent boots, the host boots from disk. The hosts do not need a connection to the vSphere Auto Deploy server.
 - For stateless caching, vSphere Auto Deploy proceeds as follows:
 - During first boot, vSphere Auto Deploy provisions the host and caches the image.
 - During subsequent boots, vSphere Auto Deploy provisions the host. If vSphere Auto Deploy is unavailable, the host boots from the cached image, however, setup can only be completed when the host can reach the vSphere Auto Deploy server.

Configure a Host Profile to Use Stateless Caching

If the vSphere Auto Deploy Server is not available, the host uses a cached image and to use stateless caching, you must configure a host profile.

When a host is set up to use stateless caching, if the vSphere Auto Deploy Server is not available, the host uses a cached image. To use stateless caching, you must configure a host profile. You can apply that host profile to other hosts that you want to set up for stateless caching.

- Decide which disk to use for caching and determine whether the caching process will overwrite an existing VMFS partition.
- In production environments, protect the vCenter Server system and the vSphere Auto Deploy server by including them in a highly available environment. Having the vCenter Server in a management cluster guarantees that VDS and virtual machine migration are available. If possible, also protect other elements of your infrastructure. See [Set Up Highly Available vSphere Auto Deploy Infrastructure](#).
- Set up your environment for vSphere Auto Deploy. See [Install and Configure vSphere Auto Deploy](#).
- Verify that a disk with at least 4GB of free space is available. If the disk is not yet partitioned, partitioning happens when you apply the host profile.
- Set up the host to attempt a network boot first and to boot from disk if network boot fails. See your hardware vendor's documentation.
- Create a host profile. See the *Host Profiles* documentation.

1. Navigate to **Home > Policies and Profiles > Host Profiles**.
2. Click the host profile you want to configure and select the **Configure** tab.
3. Click **Edit Host Profile**.
4. On the Edit host profile page of the wizard, select **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
5. In the **System Image Cache Profile Settings** drop-down menu, choose a policy option.

Option	Description
Enable stateless caching on the host	Caches the image to disk.
Enable stateless caching to a USB disk on the host	Caches the image to a USB disk attached to the host.

6. Depending on the policy option you select, you must do the following:
 - a) If you select **Enable stateless caching on the host**:
 1. Specify the information about the disk to use.

1. In the **System Disk Configuration** drop-down menu, select **User must explicitly choose the policy option**.
- b) If you select **Enable stateless caching to a USB disk on the host**:
 1. In the **System Disk Configuration** drop-down menu, select **System disk specified by user in host customization**
 2. Under **Hosts and Clusters**, right-click the host and select **Host Profiles > Edit Host Customizations**. Define a disk with persistent storage in the **Value** field for the **System disk** property.
7. Click **Save** to complete the host profile configuration.

Apply the host profile to individual hosts by using the Host Profiles feature in the vSphere Client. See the *Host Profiles* documentation. Alternatively, you can create a rule to assign the host profile to hosts with the vSphere Client or by using PowerCLI. See [Write a Rule and Assign a Host Profile to Hosts](#).

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Configure a Host Profile to Enable Stateful Installs

To set up an ESXi host provisioned with vSphere Auto Deploy to boot from disk, you must configure a host profile.

- Decide which disk to use for storing the image, and determine whether the new image will overwrite an existing VMFS partition.
- Set up your environment for vSphere Auto Deploy. See [Install and Configure vSphere Auto Deploy](#).
- Verify that a disk with at least 4GB of free space is available. If the disk is not yet partitioned, partitioning happens when you apply the host profile.
- Set up the host to boot from disk. See your hardware vendor's documentation.
- Create a host profile. See the *Host Profiles* documentation.

You can configure the host profile on a single host. You can also create a host profile on a reference host and apply that host profile to other hosts.

1. Navigate to **Home > Policies and Profiles > Host Profiles**.
2. Click the host profile you want to configure and select the **Configure** tab.
3. Click **Edit Host Profile**.
4. On the Edit host profile page of the wizard, select **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
5. In the **System Image Cache Profile Settings** drop-down menu, choose a policy option.

Option	Description
Enable stateful installs on the host	Caches the image to a disk.
Enable stateful installs to a USB disk on the host	Caches the image to a USB disk attached to the host.

6. Depending on the policy option you select, you must do the following:
 - a) If you select **Enable stateful installs on the host**:
 1. Specify the information about the disk to use.

1. In the **System Disk Configuration** drop-down menu, select **User must explicitly choose the policy option**.
- b) If you select **Enable stateful installs to a USB disk on the host**:
 1. In the **System Disk Configuration** drop-down menu, select **System disk specified by user in host customization**
 2. Under **Hosts and Clusters**, right-click the host and select **Host Profiles > Edit Host Customizations**. Define a disk with persistent storage in the **Value** field for the **System disk** property.
7. Click **Save** to complete the host profile configuration.

Apply the host profile to individual hosts by using the Host Profiles feature in the vSphere Client. See the *Host Profiles* documentation. Alternatively, you can create a rule to assign the host profile to hosts with the vSphere Client or by using PowerCLI. See [Write a Rule and Assign a Host Profile to Hosts](#).

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Setting Up a vSphere Auto Deploy Reference Host

In an environment where no state is stored on the ESXi host, a reference host helps you set up multiple hosts with the same configuration.

You configure the reference host with the logging, coredump, and other settings that you want, save the host profile, and write a rule that applies the host profile to other hosts as needed. You can configure the storage, networking, and security settings on the reference host and set up services such as syslog and NTP.

Understanding Reference Host Setup

A well-designed reference host connects to all services such as syslog, NTP, and so on. The reference host setup might also include security, storage, networking, and ESXi Dump Collector. You can apply such a host's setup to other hosts by using host profiles.

The exact setup of your reference host depends on your environment, but you might consider the following customization.

NTP Server Setup

When you collect logging information in large environments, you must make sure that log times are coordinated. Set up the reference host to use the NTP server in your environment that all hosts can share. You can specify an NTP server by running the `esxcli system ntp set` command. You can start and stop the NTP service for a host with the `esxcli system ntp set` command, or the vSphere Client.

Syslog Server Setup

All ESXi hosts run a syslog service (`vm syslogd`), which logs messages from the VMkernel and other system components to a file. You can specify the log host and manage the log location, rotation, size, and other attributes by running the `esxcli system syslog` command or by using the vSphere Client. Setting up logging on a remote host is especially important for hosts provisioned with vSphere Auto Deploy that have no local storage.

You can optionally install VMware vCenter Log Insight, which provides log aggregation and analytics.

Core Dump Setup

You can set up your reference host to send core dumps to a shared SAN LUN, or you can install ESXi Dump Collector in your environment and configure the reference host to use ESXi Dump Collector. See [Configure ESXi Dump Collector with ESXCLI](#). You

Security Setup

can either install ESXi Dump Collector by using the vCenter Server installation media or use the ESXi Dump Collector that is included in vCenter Server. After setup is complete, VMkernel memory is sent to the specified network server when the system encounters a critical failure.

In most deployments, all hosts that you provision with vSphere Auto Deploy must have the same security settings. You can, for example, set up the firewall to allow certain services to access the ESXi system, set up the security configuration, user configuration, and user group configuration for the reference host with the vSphere Client or with ESXCLI commands. Security setup includes shared user access settings for all hosts. You can achieve unified user access by setting up your reference host to use Active Directory. See the *vSphere Security* documentation.


Note: If you set up Active Directory by using host profiles, the passwords are not protected. Use the vSphere Authentication Service to set up Active Directory to avoid exposing the Active Directory password.

Networking and Storage Setup

If you reserve a set of networking and storage resources for use by hosts provisioned with vSphere Auto Deploy, you can set up your reference host to use those resources.

In large deployments, the reference host setup supports an Enterprise Network Manager, which collects all information coming from the different monitoring services that are running in the environment.

Figure 20: vSphere Auto Deploy Reference Host Setup



[Options for Configuration of a vSphere Auto Deploy Reference Host](#) explains how to perform this setup.

Watch the video "Auto Deploy Reference Hosts" for information about the reference host setup:

Options for Configuration of a vSphere Auto Deploy Reference Host

You can configure a reference host by using the vSphere Client, ESXCLI, or host profiles.

To set up a reference host, you can use the approach that suits you best.

vSphere Client	The vSphere Client supports setup of networking, storage, security, and most other aspects of an ESXi host. Set up your environment and create a host profile from the reference host for use by vSphere Auto Deploy.
ESXCLI	You can use ESXCLI for setup of many aspects of your host. ESXCLI is suitable for configuring many of the services in the vSphere environment. Commands include <code>esxcli system ntp</code> for setting up an NTP server, <code>esxcli system syslog</code> for setting up a syslog server, <code>esxcli network route</code> for adding routes and set up the default route, and <code>esxcli system coredump</code> for configuring ESXi Dump Collector.
Host Profiles Feature	Best practice is to set up a host with vSphere Client or ESXCLI and create a host profile from that host. You can instead use the Host Profiles feature in the vSphere Client and save that host profile.

vSphere Auto Deploy applies all common settings from the host profile to all target hosts. If you set up the host profile to prompt for user input, all hosts provisioned with that host profile come up in maintenance mode. You must reapply the host profile or reset host customizations to be prompted for the host-specific information.

Configure an ESXi Dump Collector

You can configure an ESXi Dump Collector to store core dumps either by using ESXCLI commands or by configuring a reference host.

Hosts provisioned with vSphere Auto Deploy do not have a local disk on which to store core dumps. You can configure an ESXi Dump Collector to store core dumps either by using ESXCLI commands or by configuring a reference host to use the ESXi Dump Collector by using the Host Profiles feature in the vSphere Client.

Configure ESXi Dump Collector with ESXCLI

Hosts provisioned with vSphere Auto Deploy do not have a local disk to store core dumps on. You can configure ESXi Dump Collector by using ESXCLI commands and keep core dumps on a network server for use during debugging.

Install ESXCLI if you want to configure the host to use ESXi Dump Collector. In troubleshooting situations, you can use ESXCLI in the ESXi Shell instead.

A core dump is the state of working memory if there is host failure. By default, a core dump is saved to the local disk. ESXi Dump Collector is especially useful for vSphere Auto Deploy, but is supported for any ESXi host. ESXi Dump Collector supports other customization, including sending core dumps to the local disk and is included with the vCenter Server management node.

Note: ESXi Dump Collector is not supported to be configured on a VMkernel interface that is running on a NSX-T N-VDS switch.

If you intend to use IPv6, and if both the ESXi host and ESXi Dump Collector are on the same local link, both can use either local link scope IPv6 addresses or global scope IPv6 addresses.

If you intend to use IPv6, and if ESXi and ESXi Dump Collector are on different hosts, both require global scope IPv6 addresses. The traffic routes through the default IPv6 gateway.

1. Set up an ESXi system to use ESXi Dump Collector by running `esxcli system coredump` in the local ESXi Shell or by using ESXCLI.

```
esxcli system coredump network set --interface-name vmk0 --server-ip 10xx.xx.xx.xx --server-port 6500
```


You must specify a VMkernel NIC and the IP address and optional port of the server to send the core dumps to. You can use an IPv4 address or an IPv6 address. If you configure an ESXi system that is running on a virtual machine that is using a vSphere standard switch, you must select a VMkernel port that is in promiscuous mode.

2. Enable ESXi Dump Collector.

```
esxcli system coredump network set --enable true
```

3. Optional: Verify that ESXi Dump Collector is configured correctly.

```
esxcli system coredump network check
```

The host on which you have set up ESXi Dump Collector is configured to send core dumps to the specified server by using the specified VMkernel NIC and optional port.

- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
- For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the new host profile.

Configure ESXi Dump Collector from the Host Profiles Feature in the vSphere Client

Hosts provisioned with vSphere Auto Deploy do not have a local disk to store core dumps on. You can configure a reference host to use ESXi Dump Collector by using the Host Profiles feature in the vSphere Client.

- Verify that you have created the host profile on which you want to configure a coredump policy. For more information on how to create a host profile, see the *vSphere Host Profiles* documentation.
- Verify that at least one partition has sufficient storage capability for core dumps from multiple hosts provisioned with vSphere Auto Deploy.

Best practice is to set up hosts to use ESXi Dump Collector with the `esxcli system coredump` command and save the host profile. For more information, see [Configure ESXi Dump Collector with ESXCLI](#).

1. Navigate to **Home > Policies and Profiles > Host Profiles**.
 2. Click the host profile you want to configure and select the **Configure** tab.
 3. Click **Edit Host Profile**.
 4. Select **Networking Configuration > Network Coredump Settings**.
 5. Select the **Enabled** check box.
 6. Specify the host NIC to use, the Network Coredump Server IP, and the Network Coredump Server Port.
 7. Click **Save** to complete the host profile configuration.
- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
 - For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
 - Power on unprovisioned hosts to provision them with the new host profile.

Configure Syslog from Host Profiles in the vSphere Client

Specify a remote syslog server by applying a host profile.

Hosts provisioned with vSphere Auto Deploy usually do not have sufficient local storage to save system logs. You can specify a remote syslog server for those hosts by setting up a reference host, saving the host profile, and applying that host profile to other hosts as needed.

- If you intend to use a remote syslog host, set up that host before you customize host profiles.

- Verify that you have access to the vSphere Client and the vCenter Server system.

Best practice is to set up the syslog server on the reference host with the vSphere Client or the `esxcli system syslog` command and to save the host profile. You can also set up syslog from the Host Profiles feature in the vSphere Client.

1. Navigate to **Home > Policies and Profiles > Host Profiles**.
 2. Optional: If no reference host exists in your environment, click **Extract Profile from Host** to create a host profile.
 3. Click the host profile you want to configure and select the **Configure** tab.
 4. Click **Edit Host Profile**.
 5. Select **Advanced Configuration Settings > Advanced Options > Advanced configuration options**.
You can select specific sub-profiles and edit the syslog settings.
 6. Optional: To create an advanced configuration option.
 - a) Click the **Add sub-profile** icon.
 - b) From the **Advanced option** drop-down list select **Configure a fixed option**.
 - c) Specify *Syslog.global.loghost* as the name of the option, and your host as the value of the option.
 7. Click **Save** to complete the host profile configuration.
- Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
 - For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
 - Power on unprovisioned hosts to provision them with the new host profile.

Enable NTP Client on a Reference Host in the vSphere Client

When you collect logging information in large environments, you must ensure that log times are coordinated.

You can set up the reference host to use the NTP server in your environment, extract the host profile and create a vSphere Auto Deploy rule to apply it to other hosts.

1. Navigate to **Home > Hosts and Clusters**, and select an ESXi host that you want to use as a reference host.
 2. Select the **Configure** tab.
 3. Under **System**, select **Time Configuration** and click **Edit**.
 4. Select the **Use Network Time Protocol (Enable NTP client)** radio button.
This option synchronizes the time and date of the host with an NTP server. The NTP service on the host periodically takes the time and date from the NTP server.
 5. From the **NTP Service Startup Policy** drop-down menu, select **Start and stop with host**.
 6. In the **NTP Servers** text box, enter the IP addresses or host names of the NTP servers that you want to use.
 7. Click **OK**.
- Extract a host profile from the reference host. See the *Host Profiles* documentation.
 - Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).

Configure Networking for Your vSphere Auto Deploy Host in the vSphere Client

You can set up networking for your vSphere Auto Deploy reference host and apply the host profile to all other hosts to guarantee a fully functional networking environment.

Provision the host you want to use as your reference host with an ESXi image by using vSphere Auto Deploy.

1. Navigate to **Home > Hosts and Clusters**, and select an ESXi host that you want to use as a reference host.
2. Select the **Configure** tab and navigate to **Networking**.
3. Perform the networking setup.
 - If you are using virtual switches and not vSphere Distributed Switch, do not add other VMkernel NICs to vSwitch0.
4. After the reference host is configured, reboot the system to verify that vmk0 is connected to the Management Network.
5. If no host profile exists for your reference host, create a host profile.
 - Create a rule that applies the host profile to all hosts that you want to provision with the settings specified in the reference host. For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).
 - For hosts that are already provisioned with vSphere Auto Deploy, perform the test and repair compliance operations in a PowerCLI session, see [Test and Repair Rule Compliance](#).
 - Power on unprovisioned hosts to provision them with the new host profile.

Configure a Reference Host for Auto-Partitioning

You can create a reference host to auto-partition all host you provision with vSphere Auto Deploy.

By default, vSphere Auto Deploy provisions hosts only if a partition is available on the host. The auto-partitioning option creates a VMFS datastore on your host's local storage. You can set up a reference host to auto-partition all hosts that you provision with vSphere Auto Deploy.

- Provision the host that you want to use as your reference host with an ESXi image by using vSphere Auto Deploy.
- Verify that you have access to a vSphere Client that can connect to the vCenter Server system.

CAUTION: If you change the default auto-partitioning behavior, vSphere Auto Deploy overwrites existing partitions regardless of their content. If you turn on this option, ensure that no unintended data loss results.

To ensure that local SSDs remain unpartitioned during auto-partitioning, you must set the parameter `skipPartitioningSsds=TRUE` on the reference host.

For more information about preventing SSD formatting during auto-partitioning, see the *vSphere Storage* documentation.

1. Navigate to **Home > Hosts and Clusters**, and select an ESXi host that you want to use as a reference host.
2. Select the **Configure** tab.
3. Under **System**, select **Advanced System Settings** and click **Edit**.
4. Search for the `VMkernel.Boot.autoPartition` key and set the value to **true**.
5. Optional: If you want the local SSDs to remain unpartitioned, search for the `VMkernel.Boot.skipPartitioningSsds` key and set the value to **true**.
6. Click **OK**.
7. If no host profile exists for your reference host, create a host profile.

Auto-partitioning is performed when the hosts boot.

- Use vSphere Auto Deploy to create a rule that applies the host profile of your reference host to all hosts immediately when they boot. To create a rule with the vSphere Client, see [Create a Deploy Rule](#). For writing a rule in a PowerCLI session, see [Write a Rule and Assign a Host Profile to Hosts](#).

Converting Stateless Hosts to Stateful Hosts

You can add physical storage to your stateless ESXi, convert them to stateful ESXi hosts, and add the hosts to a cluster that you manage by an image.

Your ESXi hosts must have physical storage attached, because a cluster managed by an image does not support stateless ESXi hosts, that use a Preboot Execution Environment (PXE) boot.

To ensure that image components are installed on a physical disk of an ESXi host during a boot process, Auto Deploy verifies that a host profile is part of the PXE boot image. The host profile must contain a "System Image Cache Profile Settings" policy with configuration set to "Enable stateful installs on the host" or "Enable stateful installs to a USB disk on the host". If the attached host profile does not contain this policy or the policy's configuration differs, then the policy is automatically configured to support a stateful installation. If a host profile is missing, a new host profile is attached to the cluster, containing only the "System Image Cache Profile Settings" policy with a configuration for a stateful installation.

Convert a Diskless ESXi Host

To add diskless ESXi hosts to a cluster managed by an image, add physical storage, convert them to stateful hosts and install an ESXi 8.0 image.

- Prepare your system and install the Auto Deploy Server. For more information, see *Prepare Your System for vSphere Auto Deploy* in the *VMware ESXi Installation and Setup* documentation.
- The host is booting from an Auto Deploy rule that contains an Image Profile.
- The Auto Deploy service is running.
- You have an empty USB stick or other supported storage.

For information on stateless caching and stateful installs, see [Use vSphere Auto Deploy for Stateless Caching and Stateful Installs](#).

1. From the home menu, click **Policies and Profiles**, and select **Host Profiles**.
2. Extract a host profile from a running host or duplicate an existing host profile with a specified configuration and a host location.
The existing host profile can be attached either to a host or to a cluster.
3. Right-click the new host profile, select **Edit Host Profile**, and browse to **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
4. In the **System Image Cache Profile Settings** drop-down menu, choose a policy option.

Option	Description
Enable stateless caching on the host	Caches the image to disk.
Enable stateless caching to a USB disk on the host	Caches the image to a USB disk attached to the host.

5. Depending on the policy option you select, you must do the following:
 - a) If you select **Enable stateless caching on the host**:
 1. Enter arguments for the first disk, and select the check boxes if needed.
 2. In the **System Disk Configuration** drop-down menu, select **User must explicitly choose the policy option**.
 - b) If you select **Enable stateless caching to a USB disk on the host**:
 1. In the **System Disk Configuration** drop-down menu, select **System disk specified by user in host customization**
 2. Under **Hosts and Clusters**, right-click the host and select **Host Profiles > Edit Host Customizations**. Define a disk with persistent storage in the **Value** field for the **System disk** property.

6. Click **Save** to complete the host profile configuration.
7. If you duplicated an existing host profile and the host profile was attached to a cluster, attach the new host profile to the cluster.
8. From the home menu, click **Auto Deploy**.
9. Deactivate the Auto Deploy rule that contains an Image Profile, and click **Edit**.
The Edit Deploy Rule dialog box appears.
10. On the **Select image profile** page of the wizard, select an 8.0 image profile.
11. On the **Select host profile** page of the wizard, select the new host profile.
12. Activate the rule and move the rule to the rule's initial position in the ordered list.
13. On the **Deployed Hosts** tab, select a single or multiple hosts.
14. Click **Remediate Host Associations** for the host.
You can monitor the progress of the remediation process in the Recent Tasks pane.
15. Shut down the remediated ESXi hosts.
16. After the ESXi hosts are powered off, install the selected storage as local boot disks.
17. Power on each ESXi host, enter the BIOS/UEFI setup, and change the boot order to boot first from the newly added storage and then from the network.
Since the newly added storage is empty, each ESXi host boots from the network and installs the ESXi 8.0 image that you specified earlier on the storage. After the installation, each ESXi host reboots and boots from the newly added storage.

The ESXi hosts boot by default from the new storage and operate as if the ESXi8.0 image is installed from a standard DVD.

Add the ESXi hosts to a cluster that you manage by a single image. For more information, see *Add a Host* from the *vCenter Server and Host Management* documentation.

Convert a Stateless ESXi Host with Enabled Stateless Caching

To add your stateless ESXi hosts to a cluster that you manage by an image, convert them to stateful hosts by installing a standard 8.0 8.0 image.

- Prepare your system and install the Auto Deploy Server. For more information, see *Prepare Your System for vSphere Auto Deploy* in the *VMware ESXi Installation and Setup* documentation.
- The host is booting from an Auto Deploy rule that contains an Image Profile.
- The Auto Deploy service is running.

For information on stateless caching and stateful installs, see [Use vSphere Auto Deploy for Stateless Caching and Stateful Installs](#).

1. From the home menu, click **Policies and Profiles**, and select **Host Profiles**.
2. Extract a host profile from a running host or duplicate an existing host profile with a specified configuration and a host location.
The existing host profile can be attached either to a host or to a cluster.

3. Right-click the new host profile, select **Edit Host Profile**, and browse to **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
4. Select **Enable stateful installs on the host** from the drop-down menu, enter `local` in the **Arguments for first disk** text box, and select the **Check to overwrite any VMFS volumes on the selected disk** check box.
5. If you duplicated an existing host profile and the host profile was attached to a cluster, attach the new host profile to the cluster.
6. From the home menu, click **Auto Deploy**.
7. Deactivate the Auto Deploy rule that contains an Image Profile, and click **Edit**.
The Edit Deploy Rule dialog box appears.
8. On the **Select image profile** page of the wizard, select an 8.0 image profile.
9. On the **Select host profile** page of the wizard, select the new host profile.
10. Activate the rule and move the rule to the rule's initial position in the ordered list.
11. On the **Deployed Hosts** tab, select a single or multiple hosts.
12. Click **Remediate Host Associations** for the host.
You can monitor the progress of the remediation process in the Recent Tasks pane.
13. Restart the hosts.

If an ESXi host has a legacy BIOS, you can change the boot order in the BIOS setup to first to boot from the local storage. For UEFI-based ESXi hosts the boot order changes automatically during the reboot.

Each ESXi host boots from the network and installs the ESXi8.0 image that you specified earlier on the storage used previously for caching. After the installation, each ESXi host reboots again and boots from the local storage.

The ESXi hosts boot by default from the new storage and operate as if the ESXi8.0 image is installed from a standard DVD.

Add the ESXi hosts to a cluster that you manage by a single image. For more information, see *Add a Host* from the *vCenter Server and Host Management* documentation.

Convert a Stateless ESXi Host with a Single VMFS Partition on the Local Disk

To add your stateless ESXi hosts to a cluster that you manage by an image, you must first convert the hosts to stateful hosts by repartitioning the VMFS partitions and installing a standard ESXi8.0 image.

- Prepare your system and install the Auto Deploy Server. For more information, see *Prepare Your System for vSphere Auto Deploy* in the *VMware ESXi Installation and Setup* documentation.
- The host is booting from an Auto Deploy rule that contains an Image Profile.
- The Auto Deploy service is running.

For information on stateless caching and stateful installs, see [Use vSphere Auto Deploy for Stateless Caching and Stateful Installs](#).

1. From the home menu, click **Policies and Profiles**, and select **Host Profiles**.
2. Extract a host profile from a running host or duplicate an existing host profile with a specified configuration and a host location.
The existing host profile can be attached either to a host or to a cluster.

3. Right-click the new host profile, select **Edit Host Profile**, and browse to **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
4. Select **Enable stateful installs on the host** from the drop-down menu, enter `localesx` in the **Arguments for first disk** text box.
5. If you duplicated an existing host profile and the host profile was attached to a cluster, attach the new host profile to the cluster.
6. From the home menu, click **Auto Deploy**.
7. Deactivate the Auto Deploy rule that contains an Image Profile, and click **Edit**.
The Edit Deploy Rule dialog box appears.
8. On the **Select image profile** page of the wizard, select an 8.0 image profile.
9. On the **Select host profile** page of the wizard, select the new host profile.
10. Activate the rule and move the rule to the rule's initial position in the ordered list.
11. On the **Deployed Hosts** tab, select a single or multiple hosts.
12. Click **Remediate Host Associations** for the host.
You can monitor the progress of the remediation process in the Recent Tasks pane.
13. Restart the hosts.

If an ESXi host has a legacy BIOS, you can change the boot order in the BIOS setup to first to boot from the local storage. For UEFI-based ESXi hosts the boot order changes automatically during the reboot.

Each ESXi host boots from the network, repartitions the VMFS partition to a standard ESXi installation partition, and installs the ESXi8.0 image that you specified earlier on the partition. After the installation, each ESXi host reboots again and boots from the standard ESXi partition.

The ESXi hosts boot by default from the new partition and operate as if the ESXi8.0 image is installed from a standard DVD.

Add the ESXi hosts to a cluster that you manage by a single image. For more information, see *Add a Host* from the *vCenter Server and Host Management* documentation.

Convert a Stateless ESXi Host with a Single VMFS Partition on a Remote Disk

To add your stateless ESXi hosts to a cluster that you manage by an image, first you must convert the hosts to stateful hosts by repartitioning their remote VMFS partitions and installing a standard ESXi8.0 image. For example, your ESXi host can boot from Fibre Channel SAN or from iSCSI SAN.

- Prepare your system and install the Auto Deploy Server. For more information, see *Prepare Your System for vSphere Auto Deploy* in the *VMware ESXi Installation and Setup* documentation.
- The host is booting from an Auto Deploy rule that contains an Image Profile.
- The Auto Deploy service is running.

For information on stateless caching and stateful installs, see [Use vSphere Auto Deploy for Stateless Caching and Stateful Installs](#).

For more information on booting from a remote disk, see the *Bootting from iSCSI SAN* in the *vSphere Storage* documentation.

1. From the home menu, click **Policies and Profiles**, and select **Host Profiles**.
2. Extract a host profile from a running host or duplicate an existing host profile with a specified configuration and a host location.
The existing host profile can be attached either to a host or to a cluster.

3. Right-click the new host profile, select **Edit Host Profile**, and browse to **Advanced Configuration Settings > System Image Cache Configuration > System Image Cache Configuration**.
4. Select **Enable stateful installs on the host** from the drop-down menu, enter `remoteesx`, `remote` or `sortedremoteesx`, `sortedremote` in the **Arguments for first disk** text box.
5. If you duplicated an existing host profile and the host profile was attached to a cluster, attach the new host profile to the cluster.
6. From the home menu, click **Auto Deploy**.
7. Deactivate the Auto Deploy rule that contains an Image Profile, and click **Edit**.
The Edit Deploy Rule dialog box appears.
8. On the **Select image profile** page of the wizard, select an 8.0 image profile.
9. On the **Select host profile** page of the wizard, select the new host profile.
10. Activate the rule and move the rule to the rule's initial position in the ordered list.
11. On the **Deployed Hosts** tab, select a single or multiple hosts.
12. Click **Remediate Host Associations** for the host.
You can monitor the progress of the remediation process in the Recent Tasks pane.
13. Restart the hosts.
Each ESXi host boots from the network, repartitions the VMFS partition to a standard ESXi installation partition, and installs the ESXi8.0 image that you specified earlier on the partition. After the installation, each ESXi host reboots again and boots from the remote standard ESXi partition.

The ESXi hosts boot by default from the new partition and operate as if the ESXi8.0 image is installed from a standard DVD.

Add the ESXi hosts to a cluster that you manage by a single image. For more information, see *Add a Host* from the *vCenter Server and Host Management* documentation.

vSphere Auto Deploy Best Practices

Set up a highly available vSphere Auto Deploy infrastructure in large production environments or when using stateless caching.

Follow best practices when installing vSphere Auto Deploy and when using vSphere Auto Deploy with other vSphere components. Set up a highly available vSphere Auto Deploy infrastructure in large production environments or when using stateless caching. Follow all security guidelines that you would follow in a PXE boot environment, and consider the recommendations in this chapter.

vSphere Auto Deploy Best Practices

Follow vSphere Auto Deploy best practices to set up networking, configure vSphere HA, and optimize your environment for vSphere Auto Deploy.

See the VMware Knowledge Base for additional best practice information.

vSphere Auto Deploy and vSphere HA Best Practices

You can improve the availability of the virtual machines running on hosts provisioned with vSphere Auto Deploy by following best practices.

Some environments configure the hosts provisioned with vSphere Auto Deploy with a distributed switch or configure virtual machines running on the hosts with Auto Start Manager. In such environments, deploy the vCenter Server system so that its availability matches the availability of the vSphere Auto Deploy server. Several approaches are possible.

- Deploy vCenter Server. The vSphere Auto Deploy server is included.
- Run vCenter Server in a vSphere HA enabled cluster and configure the virtual machine with a vSphere HA restart priority of high. Include two or more hosts in the cluster that are not managed by vSphere Auto Deploy and pin the vCenter Server virtual machine to these hosts by using a rule (vSphere HA DRS required VM to host rule). You can set up the rule and then deactivate DRS if you do not want to use DRS in the cluster. The greater the number of hosts that are not managed by vSphere Auto Deploy, the greater your resilience to host failures.

Note: This approach is not suitable if you use Auto Start Manager. Auto Start Manager is not supported in a cluster enabled for vSphere HA.

vSphere Auto Deploy Networking Best Practices

Prevent networking problems by following vSphere Auto Deploy networking best practices.

vSphere Auto Deploy and IPv6

Because vSphere Auto Deploy takes advantage of the iPXE infrastructure, if the hosts that you plan to provision with vSphere Auto Deploy are with legacy BIOS, the vSphere Auto Deploy server must have an IPv4 address. PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

IP Address Allocation

Use DHCP reservations for address allocation. Fixed IP addresses are supported by the host customization mechanism, but providing input for each host is not recommended.

VLAN Considerations

Use vSphere Auto Deploy in environments that do not use VLANs.

If you intend to use vSphere Auto Deploy in an environment that uses VLANs, make sure that the hosts that you want to provision can reach the DHCP server. How hosts are assigned to a VLAN depends on the setup at your site. The VLAN ID might be assigned by the switch or the router, or might be set in the host's BIOS or through the host profile. Contact your network administrator to determine the steps for allowing hosts to reach the DHCP server.

vSphere Auto Deploy and VMware Tools Best Practices

When you provision hosts with vSphere Auto Deploy, you can select an image profile that includes VMware Tools, or select the smaller image associated with the image profile that does not contain VMware Tools.

You can download two image profiles from the VMware download site.

- `xxxxx-standard`: An image profile that includes the VMware Tools binaries, required by the guest operating system running inside a virtual machine. The image is usually named `esxi-version-xxxxx-standard`.
- `xxxxx-no-tools`: An image profile that does not include the VMware Tools binaries. This image profile is usually smaller has a lower memory overhead, and boots faster in a PXE-boot environment. This image is usually named `esxi-version-xxxxx-no-tools`.

You can deploy ESXi using either image profile.

- If the network boot time is of no concern, and your environment has sufficient extra memory and storage overhead, use the image that includes VMware Tools.

- If you find the network boot time too slow when using the standard image, or if you want to save some space on the hosts, you can use the image profile that does not include VMware Tools, and place the VMware Tools binaries on shared storage. See, [Provision ESXi Host with an Image Profile Without VMware Tools](#).

vSphere Auto Deploy Load Management Best Practices

Simultaneously booting large numbers of hosts places a significant load on the vSphere Auto Deploy server. Because vSphere Auto Deploy is a Web server at its core, you can use existing Web server scaling technologies to help distribute the load. For example, one or more caching reverse proxy servers can be used with vSphere Auto Deploy. The reverse proxies serve up the static files that make up the majority of an ESXi boot image. Configure the reverse proxy to cache static content and pass all requests through to the vSphere Auto Deploy server. For more information, watch the video "Using Reverse Web Proxy Servers for vSphere Auto Deploy Scalability":

Use multiple TFTP servers to point to different proxy servers. Use one TFTP server for each reverse proxy server. After that, set up the DHCP server to send different hosts to different TFTP servers.

When you boot the hosts, the DHCP server redirects them to different TFTP servers. Each TFTP server redirects hosts to a different server, either the vSphere Auto Deploy server or a reverse proxy server, significantly reducing the load on the vSphere Auto Deploy server.

After a massive power outage, bring up the hosts on a per-cluster basis. If you bring multiple clusters online simultaneously, the vSphere Auto Deploy server might experience CPU bottlenecks. All hosts might come up after a delay. The bottleneck is less severe if you set up the reverse proxy.

vSphere Auto Deploy Logging and Troubleshooting Best Practices

To resolve problems that you encounter with vSphere Auto Deploy, use the vSphere Auto Deploy logging information from the vSphere Client and set up your environment to send logging information and core dumps to remote hosts.

vSphere Auto Deploy Logs

Download the vSphere Auto Deploy logs by going to the vSphere Auto Deploy page in the vSphere Client. See, [Download vSphere Auto Deploy Logs](#).

Setting Up Syslog

Set up a remote syslog server. See the *vCenter Server and Host Management* documentation for syslog server configuration information. Configure the first host you boot to use the remote syslog server and apply that host's host profile to all other target hosts. Optionally, install and use the vSphere Syslog Collector, a vCenter Server support tool that provides a unified architecture for system logging, enables network logging, and lets you combine logs from multiple hosts.

Setting Up ESXi Dump Collector

Hosts provisioned with vSphere Auto Deploy do not have a local disk to store core dumps on. Install ESXi Dump Collector and set up your first host so that all core dumps are directed to ESXi Dump Collector, and apply the host profile from that host to all other hosts. See [Configure ESXi Dump Collector with ESXCLI](#).

Using vSphere Auto Deploy in a Production Environment

When you move from a proof of concept setup to a production environment, take care to make the environment resilient.

- Protect the vSphere Auto Deploy server. See [vSphere Auto Deploy and vSphere HA Best Practices](#).
- Protect all other servers in your environment, including the DHCP server and the TFTP server.
- Follow VMware security guidelines, including those outlined in [vSphere Auto Deploy Security Considerations](#).

Set Up Highly Available vSphere Auto Deploy Infrastructure


Highly available vSphere Auto Deploy infrastructure prevents data loss and is a prerequisite for using vSphere Auto Deploy with stateless caching.

In many production situations, a highly available vSphere Auto Deploy infrastructure is required to prevent data loss. Such infrastructure is also a prerequisite for using vSphere Auto Deploy with stateless caching.

For the management cluster, install ESXi on three hosts. Do not provision the management cluster hosts with vSphere Auto Deploy.

Watch the video "Highly Available vSphere Auto Deploy Infrastructure" for information about the implementation of a highly available vSphere Auto Deploy infrastructure:

Figure 21: Highly Available vSphere Auto Deploy Infrastructure



1. Enable vSphere HA and vSphere DRS on the management cluster.
2. Set up the following virtual machines on the management cluster.

Infrastructure Component	Description
PXE boot infrastructure	TFTP and DHCP servers.
Infrastructure VM	Active Directory, DNS, vCenter Server.

Infrastructure Component	Description
vSphere Auto Deploy environment	PowerCLI, vSphere Auto Deploy server, vCenter Server. Set up this environment on a single virtual machine or on three separate virtual machines in production systems.

The vCenter Server on the infrastructure virtual machine differs from the vCenter Server in the vSphere Auto Deploy environment.

- Set up vSphere Auto Deploy to provision other hosts as needed.

Because the components on the management cluster are protected with vSphere HA, high availability is supported.

vSphere Auto Deploy Security Considerations

When you use vSphere Auto Deploy, pay careful attention to networking security, boot image security, and potential password exposure through host profiles to protect your environment.

Networking Security

Secure your network just as you secure the network for any other PXE-based deployment method. vSphere Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or of the Auto Deploy server is not checked during a PXE boot.

You can greatly reduce the security risk of Auto Deploy by completely isolating the network where Auto Deploy is used.

Boot Image and Host Profile Security

The boot image that the vSphere Auto Deploy server downloads to a machine can have the following components.

- The VIB packages that the image profile consists of are always included in the boot image.
- The host profile and host customization are included in the boot image if Auto Deploy rules are set up to provision the host with a host profile or host customization.
 - The administrator (root) password and user passwords that are included with host profile and host customization are hashed with SHA-512.
 - Any other passwords associated with profiles are in the clear. If you set up Active Directory by using host profiles, the passwords are not protected.

Use the vSphere Authentication Proxy to avoid exposing the Active Directory passwords. If you set up Active Directory using host profiles, the passwords are not protected.
- The host's public and private SSL key and certificate are included in the boot image.

Device Alias Configuration

Device aliases, also called device names, are short names associated with I/O adapters in an I/O subsystem.

For example, network uplinks have aliases such as vmnic0, vmnic1, and so on. SCSI adapter objects in the storage subsystem and graphics device objects also have aliases. A hardware device can be presented as multiple I/O adapters in the I/O subsystem. The I/O adapters can be of a different type from the underlying physical device. For example, an FCoE device is a storage I/O adapter that uses NIC hardware. Software iSCSI is a storage adapter using the network stack at the IP layer. Therefore, in the ESXi native driver model, aliases formally refer only to I/O adapters, and not to physical devices such as a PCI NIC or a PCI HBA.

Device Alias Assignment

A stateless ESXi deployment model is one where the ESXi host is not installed on hard disks, and is typically booted by using PXE. A stateful ESXi deployment model is one where the ESXi host is installed on local hard disks. Device alias assignment occurs during a stateless ESXi boot or a fresh installation of stateful ESXi. The ESXi host assigns aliases to I/O adapters in an order which is based on the underlying hardware enumeration order. The ESXi host assigns aliases first

to on-board devices and then to add-in cards based on slot order. The ESXi host cannot assign aliases to absent devices or devices without supported drivers.

An uplink that uses a NIC that is built into the motherboard receives a vmnicN alias with lower number compared to an uplink of a PCI add-in card. The NIC driver might register more than one uplink. If one of the uplinks does not correspond to an enumerable hardware device, the ESXi host assigns the next available alias to the uplink after the uplink is registered with the system.

Persistence of Device Alias Configuration

After the ESXi host assigns aliases, alias configuration is persisted. The ESXi host attempts to keep the alias of each device the same regardless of the ESXi version updates, or hardware changes, such as adding or removing devices from slots.

The persistence of the alias configuration depends on the deployment model.

- In stateful systems, the alias configuration is persisted locally on the host.
- In stateless systems, if you do not manage the stateless system by using host profiles, the alias configuration is not persisted locally on the host.
- In stateful and stateless systems that you manage by using host profiles, the alias configuration is persisted in the host profile. If you apply a host profile to a stateful host, the host profile overrides any locally persisted alias configuration.

Changes in the Device Alias Configuration

The persistence of alias configuration is based on the bus addresses of devices. If the bus address of a device is altered, the persisted alias configuration becomes inapplicable and the aliases assigned to the device might change.

Changes in the device alias configuration might occur in the following cases:

- A driver upgrade might enumerate or present an I/O adapter differently to the system compared to how the I/O adapter is presented before the driver upgrade.
- A stack upgrade might result in changes to parts of a multi-module driver setup, or to the ESXi I/O stack that supports a multi-module driver.
- BIOS or device firmware upgrades might lead to incomplete port or slot information.
- Changes in the slot position of a device.

Note: If you remove a device from the system, the alias configuration of the I/O adapters of the device is removed. If you add the same devices back to the system later, the I/O adapters of the device might not receive their previous aliases.

Device Alias Configuration in ESXi Clusters

Initial alias configuration is the same across a cluster of identical systems. However, even on a cluster that is considered homogenous, small differences in hardware or firmware might result in differences in the alias configuration between hosts.

Differences in the processing order during driver binding can also result in differences in alias configuration. For example, a NIC driver registers two uplinks, uplink-1 and uplink-2, for two ports of the same PCI device, where one of the ports is not hardware enumerable by the system. Timing changes in the order of registration of the uplinks might result in differences in how the ESXi hosts assign aliases to the uplinks. One ESXi host might assign the hardware-based alias to uplink-1, and another ESXi host might assign the hardware-based alias to uplink-2.

To match alias configuration across homogeneous hosts, you can use host profiles. The Device Alias Configuration host profile applies alias configuration to an ESXi host by mapping devices in the alias configuration to the ESXi host devices. The mapping operation is based on the hardware information sources that are used as the basis for initial alias assignment. For more information about information sources used for alias assignment, see the Knowledge Base article [KB 2091560](#).

The Device Alias Configuration host profile also flags errors, for example, when a device is present in the host profile but not present on the host.

A heterogeneous cluster does not have the same default alias configuration across its hosts. Due to the differences between the devices, a host profile cannot be applied cleanly.

Using ESXi Shell Commands to View Device Alias Information

On a running ESXi system, you can view information about I/O adapter aliases by running commands in an ESXi Shell.

Using ESXi Shell Commands to View Device Alias Information

Command	Description
<code>device alias list</code>	Lists all current I/O adapter aliases.
<code>device alias get -n<alias></code>	Displays the physical device that an I/O adapter alias maps to.
<code>network nic list</code>	Lists aliases and general information about network devices.
<code>storage core adapter list</code>	Lists all storage I/O adapters.

Note: ESXCLI commands are supported commands. Using alternative sources for displaying alias information is not recommended.

Change Device Aliases on an ESXi Host by Using Host Profiles

You can modify a device alias on ESXi hosts by editing the device alias section of the host profile attached to the hosts.

A device alias change can occur, for example, when you apply a BIOS or device firmware update. For more information about changes in the device alias configuration, see [Device Alias Configuration](#).

I/O adapters that are based on PCI hardware usually have a logical and a PCI alias entry. Both aliases must have the same value. Some I/O adapters usually have a logical alias entry only. PCI hardware devices without I/O adapters usually have a PCI alias entry only. Modify a PCI alias entry only when you need the alias for another device.

Note: Two different I/O adapters must not have the same alias, except I/O adapters that are based on PCI hardware and have a logical and a PCI alias entry.

Procedure

For information about exporting, extracting, and editing a host profile, see the *vSphere Host Profiles* documentation.

1. Export the current host profile attached to the ESXi host.
2. Extract a new host profile from the ESXi host, but do not apply the host profile.
3. Remove all entries from the device alias section of the current host profile.
4. Transfer all entries from the device alias section of the new host profile to the same section of the current host profile.
You can modify the aliases in the entries before adding them to the current host profile.
5. Reapply the current host profile to the ESXi host.

Change Device Aliases on a Stateful ESXi Host by Using ESXCLI Commands

You can modify a device alias on a stateful ESXi host that does not have an attached host profile by running ESXCLI commands on the host.

A device alias change can occur, for example, when the slot position of a device changes. For more information about changes in the device alias configuration, see [Device Alias Configuration](#).

Note: The following ESXCLI commands might not be compatible with ESXi version 6.7 and earlier.

1. To list the current assignment of aliases to device addresses, in the ESXi Shell, run `localcli --plugin-dir /usr/lib/vmware/esxcli/int/ deviceInternal alias list`.

For I/O adapters that are based on PCI hardware, you usually see a logical and a PCI alias entry.

For some I/O adapters, you usually see a logical alias entry only.

For PCI hardware devices without I/O adapters, you usually see a PCI alias entry only.

2. To change an alias, replace the `ALIAS`, `PCI_ADDRESS`, and `LOGICAL_ADDRESS` placeholder values with actual values, and follow one of the substeps.

- a) If an I/O adapter has a logical and a PCI alias, run `localcli --plugin-dir /usr/lib/vmware/esxcli/int/ deviceInternal alias store --bus-type logical --alias ALIAS --bus-address LOGICAL_ADDRESS` and `localcli --plugin-dir /usr/lib/vmware/esxcli/int/ deviceInternal alias store --bus-type pci --alias ALIAS --bus-address PCI_ADDRESS`.

The logical and PCI aliases must have the same value.

- b) If an I/O adapter has a logical alias entry only, run `localcli --plugin-dir /usr/lib/vmware/esxcli/int/ deviceInternal alias store --bus-type logical --alias ALIAS --bus-address LOGICAL_ADDRESS`.
- c) If a PCI hardware device has a PCI alias entry only, run `localcli --plugin-dir /usr/lib/vmware/esxcli/int/ deviceInternal alias store --bus-type pci --alias ALIAS --bus-address PCI_ADDRESS`.

Modify a PCI alias only when you need the alias for another device.

Note: Two different I/O adapters must not have the same alias, except I/O adapters that are based on PCI hardware and have a logical and a PCI alias entry.

3. Reboot the system.

vSphere PowerCLI Scenario for vSphere Auto Deploy


Scenario to set up and configure a working vSphere Auto Deploy environment by using PowerCLI.

In this scenario, you are going to set up and configure a working vSphere Auto Deploy environment that includes four hosts. You will create rules and provision two of the hosts with an image profile and the other two with the same image profile and a host profile that is set up to prompt for user input. This scenario can provide you with the basis for a production environment. The task descriptions assume that you are using a flat network with no VLAN tagging between the physical hosts and the rest of your environment.

To perform the tasks in this scenario, you should have the following background knowledge and privileges.

- Experience with vSphere (vCenter Server and ESXi).
- Basic knowledge of Microsoft PowerShell and PowerCLI.
- Administrator rights to a Windows system and a vCenter Server system.

Follow the tasks in the order presented in this scenario. Some steps can be performed in a different order, but the order used here limits repeated manipulation of some components. For details on the preinstallation checklist and other prerequisites to configure vSphere Auto Deploy, see [Install and Configure vSphere Auto Deploy](#).

Figure 22: vSphere Auto Deploy Setup and Hosts Provisioning Workflow

vSphere Auto Deploy takes advantage of the iPXE infrastructure and PXE booting with legacy BIOS firmware is possible only over IPv4. If the hosts that you want to provision with vSphere Auto Deploy are with legacy BIOS, the vSphere Auto Deploy server must have an IPv4 address. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

Install the TFTP Server

To set up a vSphere Auto Deploy infrastructure, you must install a TFTP server in your environment. vSphere Auto Deploy relies on a TFTP server for sending a boot image to the hosts that it provisions.

This task only installs the TFTP server. You later download a configuration file to the server. See [Configure the vSphere Auto Deploy and TFTP Environment in the vSphere Client](#).

1. Download your preferred TFTP server to a location that has network access to your vCenter Server and install the server.
2. Configure the TFTP root directory, for example `D:\TFTP_Root\`.

Install PowerCLI, to manage vSphere Auto Deploy with PowerCLI cmdlets.

Install PowerCLI

Before you can manage vSphere Auto Deploy with rules that you create with PowerCLI cmdlets, you must install PowerCLI.

Use the [VMware PowerCLI User's Guide](#) to learn about PowerShell basics, PowerCLI concepts, and how to install and configure PowerCLI.

Configure the settings of your target hosts to prepare them for provisioning with vSphere Auto Deploy.

Prepare the vSphere Auto Deploy Target Hosts

You must configure the BIOS settings of the four hosts and reconfirm the MAC address of the primary network device to prepare the target hosts for provisioning with vSphere Auto Deploy.

Hosts that you want to provision with vSphere Auto Deploy must meet the requirements for ESXi.

1. Change the BIOS settings of each of the four physical hosts to force the hosts to boot from the primary network device.
2. Reconfirm the MAC address of the primary network device.

Set up the DHCP server to serve each target host with an iPXE binary.


Prepare the DHCP Server for vSphere Auto Deploy Provisioning

When you prepare the vSphere Auto Deploy target hosts, you must set up the DHCP server in this scenario to serve each target host with an iPXE binary.

The environment in this scenario uses Active Directory with DNS and DHCP. The DHCP server is included in the vSphere supported Windows Server versions.

1. Log in to your DHCP Server with administrator privileges.
2. Create a DHCP scope for your IP address range.
 - a) Click **Start > Settings > Control Panel > Administrative Tools** and click **DHCP**.
 - b) Navigate to **DHCP > *hostname* > IPv4**.
 - c) Right-click **IPv4** and select **New Scope**.
 - d) On the Welcome screen, click **Next**, and specify a name and description for the scope.
 - e) Specify an IP address range and click **Next**.
 - f) Click **Next** until you reach the Configure DHCP Options screen and select **No, I will configure this option later**.
3. Create a DHCP reservation for each target ESXi host.
 - a) In the DHCP window, navigate to **DHCP > *hostname* > IPv4 > Autodeploy Scope > Reservations**.
 - b) Right-click **Reservations** and select **New Reservation**.
 - c) In the New Reservation window, specify a name, IP address, and the MAC address for one of the hosts. Do not include the colon (:) in the MAC address.

- d) Repeat the process for each of the other hosts.
4. Set up the DHCP Server to point the hosts to the TFTP Server.
 - a) In the DHCP window, navigate to **DHCP > *hostname* > IPv4 > Autodeploy Scope > Scope Options**.
 - b) Right click **Scope Options** and choose **Configure Options**.
 - c) In the Scope Options window, click the **General** tab.
 - d) Click **066 Boot Server Host Name** and enter the address of the TFTP server that you installed in the String value field below the Available Options.



e) Click **067 Bootfile Name** and enter `undionly.kpxe.vmw-hardwired`.

The `undionly.kpxe.vmw-hardwired` iPXE binary will be used to boot the ESXi hosts.

f) Click **Apply** and click **OK** to close the window.

5. In the DHCP window, right-click **DHCP > hostname > IPv4 > Scope > Activate** and click **Activate**.

6. Do not log out from the DHCP Server if you are using Active Directory for DHCP and DNS, or log out otherwise.

start the vCenter Server service of vSphere Auto Deploy and configure the TFTP server.

Configure the vSphere Auto Deploy and TFTP Environment in the vSphere Client

After you prepare the DHCP server, you must start the vSphere Auto Deploy vCenter Server service and configure the TFTP server. You must download a TFTP ZIP file from your vSphere Auto Deploy server. The customized FTP server serves the boot images that vSphere Auto Deploy provides.

1. Use the vSphere Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.

2. Start the vSphere Auto Deploy service.

a) Navigate to **Home > Auto Deploy**.

b) On the **Auto Deploy** page, select your vCenter Server from the drop-down menu at the top.

c) Click **Enable Auto Deploy and Image Builder** to activate the service.

If the **Image Builder** service is already active, select the **Configure** tab and click **Enable Auto Deploy Service**.

3. In the Auto Deploy inventory, click the **Configure** tab.

4. Click the **Download TFTP Zip File**.

5. Save the file `deploy-tftp.zip` to the `TFTP_Root` directory that you created when you installed the TFTP Server, and unzip the file.

Add a software depot to your inventory and use an image profile from the depot to create a rule for host provisioning.

Prepare the ESXi Software Depot and Write a Rule

After you configure the vSphere Auto Deploy infrastructure, you must add an ESXi software depot, specify an image profile, write a rule, and add it to the active rule set.

Verify that you can access the ESXi hosts that you want to provision from the system on which you run PowerCLI.

vSphere Auto Deploy provisions hosts with image profiles that define the set of VIBs that an ESXi installation process uses. Image profiles are stored in software depots. You must make sure that the correct image profile is available before you start provisioning hosts. When you add a software depot to a PowerCLI session, it is available only during the current session. It does not persist across sessions.

The steps in this task instruct you to run PowerCLI cmdlets. For additional information about the vSphere Auto Deploy cmdlets that you can run in a PowerCLI session, see [vSphere Auto Deploy PowerCLI Cmdlet Overview](#).

1. Log in as an administrator to the console of the Windows system, either directly or by using RDP.

This task assumes that you installed PowerCLI on the system on which the vCenter Server system is running.

2. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate issues occur. In a development environment, you can ignore the warning.

3. Enter the vCenter Server credentials.

4. Run `Add-EsxSoftwareDepot` to add the online depot to the PowerCLI session.

```
Add-EsxSoftwareDepot https://hostupdate.vmware.com/software/VUM/PRODUCTION/main/vmw-depot-index.xml
```

Adding the software depot is required each time you start a new PowerCLI session.

5. Validate that you successfully added the software depot by checking the contents of the depot with the `Get-EsxImageProfile` cmdlet.

The cmdlet returns information about all image profiles in the depot.

6. Create a new rule by running the `New-DeployRule` cmdlet.

```
New-DeployRule -Name "InitialBootRule" -Item ESX-9.0.0-xxxxxxx-standard -AllHosts
```

The cmdlet creates a rule that assigns the specified image profile to all hosts in the inventory.

7. Add the new rule to the active rule set to make the rule available to the vSphere Auto Deploy server.

```
Add-DeployRule -DeployRule "InitialBootRule"
```

Provision your first host with vSphere Auto Deploy and verify its image provisioning.

Provision the First Host with vSphere Auto Deploy

After creating a rule and adding it to the active rule set, you can provision the first host and check its vCenter Server location to complete verification of the image provisioning of your setup.

1. Open a console session to the physical host that you want to use as the first ESXi target host, boot the host, and look for messages that indicate a successful iPXE boot.

During the boot process, DHCP assigns an IP address to the host. The IP address matches the name you specified earlier in the DNS server. The host contacts the vSphere Auto Deploy server and downloads the ESXi binaries from the HTTP URL indicated in the iPXE tramp file that you downloaded earlier to the TFTP_Root directory. Each instance of vSphere Auto Deploy produces a custom set of files for the TFTP Server.

2. Use the vSphere Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.
3. On the vSphere Client Home page, click **Hosts and Clusters**.
4. Verify that the newly provisioned host is now in the vCenter Server inventory at the datacenter level.

By default, vSphere Auto Deploy adds hosts at the datacenter level when the boot process completes.

Extract a host profile from the host and configure it to require user input.

Extract and Configure a Host Profile from the Reference Host

After provisioning the first host, you can extract and configure a host profile that can be used to apply the same configuration to other target hosts. Configuration that differs for different hosts, such as a static IP address, can be managed through the host customization mechanism.

vSphere Auto Deploy can provision each host with the same host profile. vSphere Auto Deploy can also use host customization that allows you to specify different information for different hosts. For example, if you set up a VMkernel port for vMotion or for storage, you can specify a static IP address for the port by using the host customization mechanism.

1. Use the vSphere Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.
2. Click **Policies and Profiles** and select **Host Profiles**.
3. Click **Extract Host Profile**.
4. On the **Select host** page of the wizard, select the reference host that you configured earlier and click **Next**.
5. On the **Name and Description** page of the wizard, enter a name and description for the new profile and click **Finish**.
6. Select the host profile that you want to edit and click the **Configure** tab.
7. Click **Edit Host Profile**.
8. Select **Security and Services > Security Settings > Security > User Configuration > root**.
9. From the **Password** drop-down menu, select **User Input Password Configuration**.
10. Click **Save** to configure the host profile settings.

Create a vSphere Auto Deploy rule to apply the host profile to other ESXi hosts.

Create a Rule that Provisions Hosts from a Specific IP Range

After creating a host profile from a reference host, you can create a rule that applies the previously verified image profile and the host profile that you extracted to target hosts from a specific IP range.

1. Log in with administrator privileges to the console of the Windows system, either directly or by using RDP.
2. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate issues occur. In a development environment, you can ignore the warning.

3. Run `Add-ESXSoftwareDepot` to add the online depot to the PowerCLI session.

```
Add-ESXSoftwareDepot https://hostupdate.vmware.com/software/VUM/PRODUCTION/main/vmw-depot-index.xml
```

Adding the software depot is required each time you start a new PowerCLI session.

4. Required: Display the rules in the active rule set by running the `Get-DeployRuleset` cmdlet.
5. Create a rule that instructs vSphere Auto Deploy to provision the set of hosts from a specified IP range with the image profile that you previously selected and the host profile that you created from the reference host.

```
New-DeployRule -name "Production01Rule" -item "image_profile",ESXiGold -Pattern "ipv4=IP_range"
```

6. Add the new rule to the active rule set.

```
Add-DeployRule -DeployRule "Production01Rule"
```

7. Check the active rule set by running the `Get-DeployRuleset` command.

PowerCLI displays information similar to the following example.

```
Name: Production01Rule
```



```
PatternList:      {ipv4=address_range}
ItemList:         {ESXi-version-XXXXXX-standard, Compute01, ESXiGold}
```

Provision the hosts and set up the host customizations.

Provision Hosts and Set Up Host Customizations

With the rule in place that provisions hosts using an image profile and a host profile, you can provision specific target hosts. If any host profile items are set to prompt the user for input, the host comes up in maintenance mode. You apply the host profile or check host compliance to be prompted for the information. The system associates the host customization with the host.

1. Boot the remaining hosts you want to provision.
vSphere Auto Deploy boots the hosts, applies the host profile, and adds the hosts to the vCenter Server inventory. The hosts remain in maintenance mode because the host profile from the reference host is set up to require user input for each host.
2. Use the vSphere Client to connect to the vCenter Server system that manages the vSphere Auto Deploy server.
3. Click **Policies and Profiles** and select **Host Profiles**.
4. Right-click the newly created host profile for Auto Deploy and click **Edit Host Customizations**.
5. Select the hosts, enter the required host customizations and click **Finish**.
Alternatively, you can also **Import Host Customizations** file.
6. Apply the host profile to each of the hosts and get the hosts out of maintenance mode. Alternatively, you can reboot each host.
When the reboot progress completes, all hosts are running with the image you specify and use the configuration in the reference host profile. The cluster shows that all hosts are fully compliant.

All hosts are now configured with the shared information through the reference host profile and with the host-specific information through the host customization mechanism. The next time you boot the hosts, they receive the complete Host Profile information, including the host-specific information, and boot up completely configured and out of Maintenance Mode.

Troubleshooting vSphere Auto Deploy

The vSphere Auto Deploy troubleshooting topics cover some situations when provisioning hosts with vSphere Auto Deploy does not work as expected.

vSphere Auto Deploy Rule Takes Long to Complete

After you run an Auto Deploy rule, it takes long to complete, and you see no progress for the tasks.

Auto Deploy automatically detects if a new version of the image that a cluster uses is available, or when the cache lifetime of files expires, and refreshes the cache by re-downloading the files from the latest version of the software depot. As a result, you might see a delay in the implementation of Auto Deploy rules.

When the cache lifetime of files cached by Auto Deploy expires, Auto Deploy automatically refreshes the cache from the software depot. Since a default software depot is usually more than 300 MB, depending on the network, the download might take long.

You can also see a delay in deploying Auto Deploy rules when converting a cluster that you manage with a single image to one that you manage with a configuration on a cluster level. If a host attempts to boot during the time Auto Deploy caches the configuration file, you might see a delay, because of the time Auto Deploy needs to create the cache.

A general vSphere infrastructure problem might also prevent Auto Deploy rules from running in a timely manner.

Regardless of any connectivity issues or breakage, Auto Deploy keeps active sessions and persists in booting hosts until it succeeds.

vSphere Auto Deploy TFTP Timeout Error at Boot Time

A TFTP Timeout error message appears when a host provisioned with vSphere Auto Deploy boots. The text of the message depends on the BIOS.

A TFTP Timeout error message appears when a host provisioned with vSphere Auto Deploy boots. The text of the message depends on the BIOS.

The TFTP server is down or unreachable.

- Ensure that your TFTP service is running and reachable by the host that you are trying to boot.
- To view the diagnostic logs for details on the present error, see your TFTP service documentation.

vSphere Auto Deploy Host Boots with Wrong Configuration

A host is booting with a different ESXi image, host profile, or folder location than the one specified in the rules.

A host is booting with a different ESXi image profile or configuration than the image profile or configuration that the rules specify. For example, you change the rules to assign a different image profile, but the host still uses the old image profile.

After the host has been added to a vCenter Server system, the boot configuration is determined by the vCenter Server system. The vCenter Server system associates an image profile, host profile, or folder location with the host.

Use the `Test-DeployRuleSetCompliance` and `Repair-DeployRuleSetCompliance` vSphere PowerCLI cmdlets to reevaluate the rules and to associate the correct image profile, host profile, or folder location with the host.

Host Is Not Redirected to vSphere Auto Deploy Server

During boot, a host that you want to provision with vSphere Auto Deploy loads iPXE. The host is not redirected to the vSphere Auto Deploy server.

During boot, a host that you want to provision with vSphere Auto Deploy loads iPXE. The host is not redirected to the vSphere Auto Deploy server.

The `tramp` file that is included in the TFTP ZIP file has the wrong IP address for the vSphere Auto Deploy server.

Correct the IP address of the vSphere Auto Deploy server in the `tramp` file, as explained in the *vSphere Installation and Setup* documentation.

Package Warning Message When You Assign an Image Profile to a vSphere Auto Deploy Host

When you run a vSphere PowerCLI cmdlet that assigns an image profile that is not vSphere Auto Deploy ready, a warning message appears.

When you write or modify rules to assign an image profile to one or more hosts, the following error results:

```
Warning: Image Profile <name-here> contains one or more software packages that are not
stateless-ready. You may experience problems when using this profile with Auto Deploy.
```

Each VIB in an image profile has a `stateless-ready` flag that indicates that the VIB is meant for use with vSphere Auto Deploy. You get the error if you attempt to write a vSphere Auto Deploy rule that uses an image profile in which one or more VIBs have that flag set to `FALSE`.

Note: You can use hosts provisioned with vSphere Auto Deploy that include VIBs that are not stateless ready without problems. However booting with an image profile that includes VIBs that are not stateless ready is treated like a fresh

install. Each time you boot the host, you lose any configuration data that would otherwise be available across reboots for hosts provisioned with vSphere Auto Deploy.

1. Use vSphere ESXi Image Builder cmdlets in a vSphere PowerCLI session to view the VIBs in the image profile.
2. Remove any VIBs that are not stateless-ready.
3. Rerun the vSphere Auto Deploy cmdlet.

vSphere Auto Deploy Host with a Built-In USB Flash Drive Does Not Send Coredumps to Local Disk

If your vSphere Auto Deploy host has a built-in USB flash drive, and an error results in a coredump, the coredump is lost. Set up your system to use ESXi Dump Collector to store coredumps on a networked host.

If your vSphere Auto Deploy host has a built-in USB Flash, and if it encounters an error that results in a coredump, the coredump is not sent to the local disk.

1. Install ESXi Dump Collector on a system of your choice.
ESXi Dump Collector is included with the vCenter Server installer.

2. Use ESXCLI to configure the host to use ESXi Dump Collector.

```
esxcli conn_options system coredump network set IP-addr,port
esxcli system coredump network set -e true
```

3. Use ESXCLI to deactivate local coredump partitions.

```
esxcli conn_options system coredump partition set -e false
```

vSphere Auto Deploy Host Reboots After Five Minutes

A vSphere Auto Deploy host boots and displays iPXE information, but reboots after five minutes.

A host to be provisioned with vSphere Auto Deploy boots from iPXE and displays iPXE information on the console. However, after five minutes, the host displays the following message to the console and reboots.

```
This host is attempting to network-boot using VMware
AutoDeploy. However, there is no ESXi image associated with this host.
Details: No rules containing an Image Profile match this
host. You can create a rule with the New-DeployRule PowerCLI cmdlet
and add it to the rule set with Add-DeployRule or Set-DeployRuleSet.
The rule should have a pattern that matches one or more of the attributes
listed below.
```

The host might also display the following details:

```
Details: This host has been added to VC, but no Image Profile
is associated with it. You can use Apply-ESXiImageProfile in the
PowerCLI to associate an Image Profile with this host.
Alternatively, you can reevaluate the rules for this host with the
Test-DeployRuleSetCompliance and Repair-DeployRuleSetCompliance cmdlets.
```

The console then displays the host's machine attributes including vendor, serial number, IP address, and so on.

No image profile is currently associated with this host.

You can assign an image profile to the host by running the `Apply-ESXiImageProfile` cmdlet, or by creating the following rule:

1. Run the `New-DeployRule` cmdlet to create a rule that includes a pattern that matches the host with an image profile.
2. Run the `Add-DeployRule` cmdlet to add the rule to a ruleset.

3. Run the `Test-DeployRuleSetCompliance` cmdlet and use the output of that cmdlet as the input to the `Repair-DeployRuleSetCompliance` cmdlet.

vSphere Auto Deploy Host Cannot Contact TFTP Server

The host that you provision with vSphere Auto Deploy cannot contact the TFTP server.

When you attempt to boot a host provisioned with vSphere Auto Deploy, the host performs a network boot and is assigned a DHCP address by the DHCP server, but the host cannot contact the TFTP server.

The TFTP server might have stopped running, or a firewall might block the TFTP port.

- If you installed the WinAgents TFTP server, open the WinAgents TFTP management console and verify that the service is running. If the service is running, check the Windows firewall's inbound rules to make sure the TFTP port is not blocked. Turn off the firewall temporarily to see whether the firewall is the problem.
- For all other TFTP servers, see the server documentation for debugging procedures.

vSphere Auto Deploy Host Cannot Retrieve ESXi Image from vSphere Auto Deploy Server

The host that you provision with vSphere Auto Deploy stops at the iPXE boot screen.

When you attempt to boot a host provisioned with vSphere Auto Deploy, the boot process stops at the iPXE boot screen and the status message indicates that the host is attempting to get the ESXi image from the vSphere Auto Deploy server.

The vSphere Auto Deploy service might be stopped or the vSphere Auto Deploy server might be inaccessible.

1. Log in to the system on which you installed the vSphere Auto Deploy server.
2. Check that the vSphere Auto Deploy server is running.
 - a) Click **Start > Settings > Control Panel > Administrative Tools**.
 - b) Double-click **Services** to open the Services Management panel.
 - c) In the Services field, look for the VMware vSphere Auto Deploy Waiter service and restart the service if it is not running.
3. Open a Web browser, enter the following URL, and check whether the vSphere Auto Deploy server is accessible.
`https://Auto_Deploy_Server_IP_Address:Auto_Deploy_Server_Port/vmw/rdb`

Note: Use this address only to check whether the server is accessible.
4. If the server is not accessible, a firewall problem is likely.
 - a) Try setting up permissive TCP Inbound rules for the vSphere Auto Deploy server port.
 The port is 6501 unless you specified a different port during installation.
 - b) As a last resort, deactivate the firewall temporarily and enable it again after you verified whether it blocked the traffic. Do not deactivate the firewall on production environments.
 To deactivate the firewall, run `netsh firewall set opmode disable`. To enable the firewall, run `netsh firewall set opmode enable`.

vSphere Auto Deploy Host Does Not Get a DHCP Assigned Address

The host you provision with vSphere Auto Deploy fails to get a DHCP Address.

When you attempt to boot a host provisioned with vSphere Auto Deploy, the host performs a network boot but is not assigned a DHCP address. The vSphere Auto Deploy server cannot provision the host with the image profile.

You might have a problem with the DHCP service or with the firewall setup.

1. Check that the DHCP server service is running on the Windows system on which the DHCP server is set up to provision hosts.
 - a) Click **Start > Settings > Control Panel > Administrative Tools**.
 - b) Double-click **Services** to open the Services Management panel.
 - c) In the Services field, look for the DHCP server service and restart the service if it is not running.
2. If the DHCP server is running, recheck the DHCP scope and the DHCP reservations that you configured for your target hosts.
If the DHCP scope and reservations are configured correctly, the problem most likely involves the firewall.
3. As a temporary workaround, turn off the firewall to see whether that resolves the problem.
 - a) Open the command prompt by clicking **Start > Program > Accessories > Command prompt**.
 - b) Type the following command to temporarily turn off the firewall. Do not turn off the firewall in a production environment.
`netsh firewall set opmode disable`
 - c) Attempt to provision the host with vSphere Auto Deploy.
 - d) Type the following command to turn the firewall back on.
`netsh firewall set opmode enable`
4. Set up rules to allow DHCP network traffic to the target hosts.
See the firewall documentation for DHCP and for the Windows system on which the DHCP server is running for details.

vSphere Auto Deploy Host Does Not Network Boot

The host you provision with vSphere Auto Deploy comes up but does not network boot.

When you attempt to boot a host provisioned with vSphere Auto Deploy, the host does not start the network boot process.

You did not enable your host for network boot.

1. Reboot the host and follow the on-screen instructions to access the BIOS configuration.
2. In the BIOS configuration, enable Network Boot in the Boot Device configuration.

Recovering from Database Corruption on the vSphere Auto Deploy Server

In some situations, you might have a problem with the vSphere Auto Deploy database. The most efficient recovery option is to replace the existing database file with the most recent backup.

When you use vSphere Auto Deploy to provision the ESXi hosts in your environment, you might encounter a problem with the vSphere Auto Deploy database.

Important: This is a rare problem. Follow all other vSphere Auto Deploy troubleshooting strategies before you replace the current database file. Rules or associations that you created since the backup you choose are lost.

This problem happens only with hosts that are provisioned with vSphere Auto Deploy.

1. Stop the vSphere Auto Deploy server service.
2. Find the vSphere Auto Deploy log by going to the vSphere Auto Deploy page in the vSphere Client.
3. Check the logs for the following message:

```
DatabaseError: database disk image is malformed.
```

If you see the message, replace the existing database with the most recent backup.

4. Go to the vSphere Auto Deploy data directory at `/var/lib/abd`.
The directory contains a file named `db`, and backup files named `db-yyy-mm-dd`.
5. Rename the current `db` file.
VMware Support might ask for that file if you call for assistance.
6. Rename the most recent backup to `db`.
7. Restart the vSphere Auto Deploy server service.
8. If the message still appears in the log, repeat the steps to use the next recent backup until vSphere Auto Deploy works without database errors.

Setting Up ESX

These topics provide information about using the direct console user interface and configuring defaults for ESXi.

- [Initial ESXi Configuration](#)
- [Enable ESXi Shell and SSH Access with the Direct Console User Interface](#)
- [Set the Password for the Administrator Account](#)
- [Configuring the BIOS Boot Settings](#)
- [Configuring Network Settings](#)
- [Test the Management Network](#)
- [Restart the Management Agents](#)
- [Restart the Management Network](#)
- [Test Connectivity to Devices and Networks](#)
- [Restoring the Standard Switch](#)
- [Configuring System Logging](#)
- [Set the Host Image Profile Acceptance Level](#)
- [Remove All Custom Packages on ESXi](#)
- [Modify ESXi Configuration Files](#)
- [Deactivate Support for Non-ASCII Characters on ESXi](#)
- [Reset the System Configuration](#)

Initial ESXi Configuration

ESXi Autoconfiguration

When you turn on the ESXi host for the first time or after resetting the configuration defaults, the host enters an autoconfiguration phase. This phase configures system network and storage devices with default settings.

By default, Dynamic Host Configuration Protocol (DHCP) configures IP, and all visible blank internal disks are formatted with the virtual machine file system (VMFS) so that virtual machines can be stored on the disks.

Managing ESXi Remotely

You can use the VMware Host Client, the vSphere Client and vCenter Server to manage your ESXi hosts.

For instructions about downloading and installing vCenter Server and the vCenter Server components, see *vCenter Server Installation and Setup*. For information about installing the VMware Host Client, see *vSphere Single Host Management*.

About the Direct Console ESXi Interface

Use the direct console interface for initial ESXi configuration and troubleshooting.

Connect a keyboard and monitor to the host to use the direct console. After the host completes the autoconfiguration phase, the direct console appears on the monitor. You can examine the default network configuration and change any settings that are not compatible with your network environment.

Key operations available to you in the direct console include:

- Configuring hosts
- Setting up administrative access
- Troubleshooting

You can also use vSphere Client to manage the host by using vCenter Server.

Table 21: Navigating in the Direct Console

Action	Key
View and change the configuration	F2
Change the user interface to high-contrast mode	F4
Shut down or restart the host	F12
View the VMkernel log	Alt+F12
Switch to the shell console	Alt+F1
Switch to the direct console user interface	Alt+F2
Move the selection between fields	Arrow keys
Select a menu item	Enter
Toggle a value	Spacebar
Confirm sensitive commands, such as resetting configuration defaults	F11
Save and exit	Enter
Exit without saving	Esc
Exit system logs	q

Configure the Keyboard Layout for the Direct Console

You can configure the layout for the keyboard that you use with the direct console.

1. From the direct console, select **Configure Keyboard** and press Enter.
2. Select the layout to use.
3. Press the spacebar to toggle selections on and off.
4. Press Enter.

Create a Security Banner for the Direct Console

A security banner is a message that is displayed on the direct console **Welcome** screen.

1. From the vSphere Client, connect to the vCenter Server.
2. Select the host in the inventory.
3. Click the **Configure** tab.
4. Under System, select **Advanced System Settings**.
5. Select `Annotations.WelcomeMessage`.
6. Click the **Edit** icon.
7. Enter a security message.

The message is displayed on the direct console **Welcome** screen.

Redirecting the Direct Console to a Serial Port

To manage your ESXi host remotely from a serial console, you can redirect the direct console to a serial port.

vSphere supports the VT100 terminal type and the PuTTY terminal emulator to view the direct console over the serial port.

You can redirect the direct console to a serial port in several ways.

Redirect the Direct Console to a Serial Port by Setting the Boot Options Manually

When you redirect the direct console to a serial port by setting the boot options, the change does not persist for subsequent boots.

Verify that the serial port is not in use for serial logging and debugging.

1. Start the host.
2. When the Loading VMware Hypervisor window appears, press Shift+O to edit boot options.
3. Deactivate the logPort and gdbPort on com1 and set tty2Port to com1 by entering the following boot options:

```
"gdbPort=none logPort=none tty2Port=com1";
```

To use com2 instead, replace `com1` with `com2`.

The direct console is redirected to the serial port until you reboot the host. To redirect the direct console for subsequent boots, see [Redirect the Direct Console to a Serial Port from the vSphere Client](#)

Redirect the Direct Console to a Serial Port from the vSphere Client

You can manage the ESXi host remotely from a console that is connected to the serial port by redirecting the direct console to either of the serial ports com1 or com2. When you use the vSphere Client to redirect the direct console to a serial port, the boot option that you set persists after subsequent reboots.

- Verify that you can access the host from the vSphere Client.
 - Verify that the serial port is not in use for serial logging and debugging, or for ESX Shell (tty1Port).
1. From the vSphere Client, connect to the vCenter Server.
 2. Select the host in the inventory.
 3. Click the **Configure** tab.
 4. Under System, select **Advanced System Settings**.
 5. Make sure that the **VMkernel.Boot.logPort** and **VMkernel.Boot.gdbPort** fields are not set to use the com port that you want to redirect the direct console to.
 6. Set **VMkernel.Boot.tty2Port** to the serial port to redirect the direct console to: `com1` or `com2`.
 7. Reboot the host.

You can now manage the ESXi host remotely from a console that is connected to the serial port.

Redirect the Direct Console to a Serial Port in a Host Deployed with Auto Deploy

After you redirect the direct console to a serial port, you can make that setting part of the host profile that persists when you reprovision the host with Auto Deploy.

The serial port must not already be in use for serial logging and debugging.

1. From the , connect to the vCenter Server.
2. Select the host in the inventory.
3. Under System, select **Advanced System Settings**.
4. Make sure that the **VMkernel.Boot.logPort** and **VMkernel.Boot.gdbPort** fields are not set to use the com port that you want to redirect the direct console to.
5. Set **VMkernel.Boot.tty2Port** to the serial port to redirect the direct console to: `com1` or `com2`.
6. Click **OK**.
7. Save the host profile and attach the host to the profile. See the *vSphere Host Profiles* documentation.

The setting to redirect the direct console to a serial port is stored by vCenter Server and persists when you reprovision the host with Auto Deploy.

Enable ESXi Shell and SSH Access with the Direct Console User Interface

Use the direct console user interface to enable the ESXi Shell.

1. From the Direct Console User Interface, press F2 to access the System Customization menu.
2. Select **Troubleshooting Options** and press Enter.
3. From the Troubleshooting Mode Options menu, select a service to enable.
 - Enable ESXi Shell
 - Enable SSH

4. Press Enter to enable the service.

5. Optional: Set the timeout for the ESXi Shell.

By default, timeouts for the ESXi Shell is 0 (not active).

The availability timeout setting is the number of minutes that can elapse before you must log in after the ESXi Shell is enabled. After the timeout period, if you have not logged in, the shell is deactivated.

Note: If you are logged in when the timeout period elapses, your session will persist. However, the ESXi Shell is deactivated, preventing other users from logging in.

a) From the Troubleshooting Mode Options menu, select **Modify ESXi Shell and SSH timeouts** and press Enter.

b) Enter the availability timeout in minutes.

The availability timeout is the number of minutes that can elapse before you must log in after the ESXi Shell is enabled.

c) Press Enter.

d) Enter the idle timeout.

The idle timeout is the number of minutes that can elapse before the user is logged out of an idle interactive sessions. Changes to the idle timeout apply the next time a user logs in to the ESXi Shell and do not affect existing sessions.

6. Press Esc until you return to the main menu of the Direct Console User Interface.

Set the Password for the Administrator Account

You can use the direct console to set the password for the administrator account (root).

The administrative user name for the ESXi host is root. By default, the administrative password is not set.

1. From the direct console, select **Configure Password**.

2. Optional: If a password is already set up, type the password in the **Old Password** line and press Enter.

3. In the **New Password** line, type a new password and press Enter.

4. Retype the new password and press Enter.

Configuring the BIOS Boot Settings

If your server has multiple drives, you might need to configure the BIOS settings.

The BIOS boot configuration determines how your server boots. Generally, the CD-ROM device is listed first.

Note: If you are using ESXi Embedded, the BIOS boot configuration determines whether your server boots into the ESXi boot device or another boot device. Generally, the USB flash device is listed first in the BIOS boot settings on the machine that hosts ESXi.

When you install or upgrade ESXi in UEFI mode, the installer creates an UEFI boot option named VMware ESXi and makes it the default boot option, so you do not need to change the boot order.

You can change the boot setting by configuring the boot order in the BIOS during startup or by selecting a boot device from the boot device selection menu. When you change the boot order in the BIOS, the new setting affects all subsequent reboots. When you select a boot device from the boot device selection menu, the selection affects the current boot only.

Some servers do not have a boot device selection menu, in which case you must change the boot order in the BIOS even for one-time boots, and then change it back again during a subsequent reboot.

Change the BIOS Boot Setting for ESXi

Configure the BIOS boot setting for ESXi if you want the server to boot into ESXi by default.

ESXi Installable and ESXi Embedded cannot exist on the same host.

1. While the ESXi host is powering on, press the key required to enter your host's BIOS setup.
Depending on your server hardware, the key might be a function key or Delete. The option to enter the BIOS setup might be different for your server.
2. Select the BIOS boot setting.

Option	Description
If you are using the installable version of ESXi	Select the disk on which you installed the ESXi software and move it to the first position in the list. The host boots into ESXi.
If you are using ESXi Embedded	Select the USB flash device and move it to the first position in the list. The host starts in ESXi mode.

Configure the Boot Setting for Virtual Media

If you are using remote management software to set up ESXi, you might need to configure the boot setting for virtual media.

ESXi Installable and ESXi Embedded cannot exist on the same host.

Virtual media is a method of connecting a remote storage media such as CD-ROM, USB mass storage, ISO image, and floppy disk to a target server that can be anywhere on the network. The target server has access to the remote media, and can read from and write to it as if it were physically connected to the server's USB port.

1. Connect the media to the virtual device.
For example, if you are using a Dell server, log in to the Dell Remote Access Controller (DRAC) or a similar remote management interface and select a physical floppy or CD-ROM drive, or provide a path to a floppy image or CD-ROM image.
2. Reboot the server.
3. While the server is powering on, enter the device selection menu.
Depending on your server hardware, the key might be a function key or Delete.
4. Follow the instructions to select the virtual device.

The server boots from the configured device once and goes back to the default boot order for subsequent boots. When you install or upgrade ESXi in UEFI mode, you do not need to change the boot order, because the system default boot order is set to VMware ESXi.

Configuring Network Settings

ESXi requires one IP address for the management network. To configure basic network settings, use the vSphere Client or the direct console.

Use the vSphere Client if you are satisfied with the IP address assigned by the DHCP server.

Use the direct console for network configuration in the following cases:

- You are not satisfied with the IP address assigned by the DHCP server.
- You are not allowed to use the IP address assigned by the DHCP server.
- ESXi does not have an IP address. This situation might occur if the autoconfiguration phase did not succeed in configuring DHCP.
- The wrong network adapter was selected during the autoconfiguration phase.

Use ESXCLI commands to configure your network settings. See [esxcli network Commands](#).

Network Access to Your ESXi Host

The default behavior is to configure the ESXi management network using DHCP. You can override the default behavior and use static IP settings for the management network after the installation is completed.

Table 22: Network Configuration Scenarios Supported by ESXi

Scenario	Approach
You want to accept the DHCP-configured IP settings.	In the ESXi direct console, you can find the IP address assigned through DHCP to the ESXi management interface. You can use that IP address to connect to the host from the vSphere Client and customize settings, including changing the management IP address.
One of the following is true: - You do not have a DHCP server. - The ESXi host is not connected to a DHCP server. - Your connected DHCP server is not functioning properly.	During the autoconfiguration phase, the software assigns the link local IP address, which is in the subnet 169.254.x.x/16. The assigned IP address appears on the direct console. You can override the link local IP address by configuring a static IP address using the direct console.
The ESXi host is connected to a functioning DHCP server, but you do not want to use the DHCP-configured IP address.	During the autoconfiguration phase, the software assigns a DHCP-configured IP address. You can make the initial connection by using the DHCP-configured IP address. Then you can configure a static IP address. If you have physical access to the ESXi host, you can override the DHCP-configured IP address by configuring a static IP address using the direct console.
Your security deployment policies do not permit unconfigured hosts to be powered on the network.	Follow the setup procedure in Configure the Network Settings on a Host That Is Not Attached to the Network .

ESXi Networking Security Recommendations

Isolation of network traffic is essential to a secure ESXi environment. Different networks require a different access and level of isolation.

Your ESXi host uses several networks. Use appropriate security measures for each network, and isolate traffic for specific applications and functions. For example, ensure that VMware vSphere® vMotion® traffic does not travel over networks where virtual machines are located. Isolation prevents snooping. Having separate networks is also recommended for performance reasons.

- vSphere infrastructure networks are used for features such as vSphere vMotion, VMware vSphere Fault Tolerance, VMware vSAN, and storage. Isolate these networks for their specific functions. It is often not necessary to route these networks outside a single physical server rack.
- A management network isolates client traffic, command-line interface (CLI) or API traffic, and third-party software traffic from other traffic. In general, the management network is accessible only by system, network, and security administrators. To secure access to the management network, use a bastion host or a virtual private network (VPN). Strictly control access within this network.
- Virtual machine traffic can flow over one or many networks. You can enhance the isolation of virtual machines by using virtual firewall solutions that set firewall rules at the virtual network controller. These settings travel with a virtual machine as it migrates from host to host within your vSphere environment.

Choose Network Adapters for the Management Network

Traffic between an ESXi host and any external management software is transmitted through an Ethernet network adapter on the host. You can use the direct console to choose the network adapters that are used by the management network.

Examples of external management software include the vCenter Server and SNMP client. Network adapters on the host are named `vmnicN`, where N is a unique number identifying the network adapter, for example, `vmnic0`, `vmnic1`, and so forth.

During the autoconfiguration phase, the ESXi host chooses `vmnic0` for management traffic. You can override the default choice by manually choosing the network adapter that carries management traffic for the host. In some cases, you might want to use a Gigabit Ethernet network adapter for your management traffic. Another way to help ensure availability is to select multiple network adapters. Using multiple network adapters enables load balancing and failover capabilities.

1. From the direct console, select **Configure Management Network** and press Enter.
2. Select **Network Adapters** and press Enter.
3. Select a network adapter and press Enter.

After the network is functional, you can use the vSphere Client to connect to the ESXi host through vCenter Server.

Set the VLAN ID

You can set the virtual LAN (VLAN) ID number of the ESXi host.

1. From the direct console, select **Configure Management Network** and press Enter.
2. Select **VLAN** and press Enter.
3. Enter a VLAN ID number from 1 through 4094.

Configuring IP Settings for ESXi

By default, DHCP sets the IP address, subnet mask, and default gateway.

For future reference, write down the IP address.

When you have access to the direct console, you can optionally configure a static network address. The default subnet mask is 255.255.0.0.

Configure IP Settings from the Direct Console

If you have physical access to the host or remote access to the direct console, you can use the direct console to configure the IP address, subnet mask, and default gateway.

1. Select **Configure Management Network** and press Enter.
2. Select **IP Configuration** and press Enter.
3. Select **Set static IP address and network configuration**.
4. Enter the IP address, subnet mask, and default gateway and press Enter.

Configure IP Settings from the vSphere Client

If you do not have physical access to the host, you can use the vSphere Client to configure static IP settings.

1. Log in to the vCenter Server from the vSphere Client.
2. Select the host in the inventory.
3. On the **Configure** tab, expand **Networking**.
4. Select **VMkernel adapters**.
5. Select **vmk0 Management Network** and click the edit icon.
6. Select **IPv4 settings**.
7. Select **Use static IPv4 settings**.
8. Enter or change the static IPv4 address settings.
9. Optional: Set static IPv6 addresses.
 - a) Select **IPv6 settings**.
 - b) Select **Static IPv6 addresses**.
 - c) Click the add icon.
 - d) Type the IPv6 address and click **OK**.
10. Click **OK**.

Configuring DNS for ESXi

You can select either manual or automatic DNS configuration of the ESXi host.

The default is automatic. For automatic DNS to work, your network environment must have a DHCP server and a DNS server.

In network environments where automatic DNS is not available or not desirable, you can configure static DNS information, including a host name, a primary name server, a secondary name server, and DNS suffixes.

Configure DNS Settings from the Direct Console

If you have physical access to the host or remote access to the direct console, you can use the direct console to configure DNS information.

1. Select **Configure Management Network** and press Enter.
2. Select **DNS Configuration** and press Enter.
3. Select **Use the following DNS server addresses and hostname**.
4. Enter the primary server, an alternative server (optional), and the host name.

Configure DNS Suffixes

If you have physical access to the host, you can use the direct console to configure DNS information. By default, DHCP acquires the DNS suffixes.

1. From the direct console, select **Configure Management Network**.
2. Select **Custom DNS Suffixes** and press Enter.
3. Enter new DNS suffixes.

Configure the Network Settings on a Host That Is Not Attached to the Network

Some highly secure environments do not permit unconfigured hosts on the network to be powered on. You can configure the host before you attach the host to the network.

Verify that no network cables are connected to the host.

1. Power on the host.
2. Use the direct console user interface to configure the password for the administrator account (root).
3. Use the direct console user interface to configure a static IP address.
4. Connect a network cable to the host.
5. Optional: Use the vSphere Client to connect to a vCenter Server system.
6. Optional: Add the host to the vCenter Server inventory.

Test the Management Network

You can use the direct console to do simple network connectivity tests.

The direct console performs the following tests.

- Pings the default gateway
- Pings the primary DNS name server
- Pings the secondary DNS nameserver

- Resolves the configured host name

1. From the direct console, select **Test Management Network** and press Enter.
2. Press Enter to start the test.

Restart the Management Agents

The management agents synchronize VMware components and let you access the ESXi host by using the vSphere Client and vCenter Server.

The vSphere Client and vCenter Server are installed with the vSphere software. You might need to restart the management agents if remote access is interrupted. Restarting the management agents restarts all management agents and services that are installed and running in `/etc/init.d` on the ESXi host. Typically, these agents include `hostd`, `ntpd`, `sfcbd`, `slpd`, `wsman`, and `vobd`. The software also restarts the Fault Domain Manager (FDM) if installed.

Users accessing this host by using the vSphere Client and vCenter Server lose connectivity when you restart management agents.

1. From the direct console, select **Troubleshooting Options** and press Enter.
2. Select **Restart Management Agents** and press Enter.
3. Press F11 to confirm the restart.

The ESXi host restarts the management agents and services.

Restart the Management Network

Restarting the management network interface might be required to restore networking or to renew a DHCP lease.

Restarting the management network will result in a brief network outage that might temporarily affect running virtual machines.

If a renewed DHCP lease results in a new network identity (IP address or host name), remote management software will be disconnected.

1. From the direct console, select **Restart Management Network** and press Enter.
2. Press F11 to confirm the restart.

Test Connectivity to Devices and Networks

You can use the direct console to perform some simple network connectivity tests, and specify other devices and networks.

1. From the direct console, select **Test Management Network** and press Enter.
2. Type addresses to ping or another DNS host name to resolve.
3. Press Enter to start the test.

Restoring the Standard Switch

A vSphere Distributed Switch functions as a single virtual switch across all associated hosts.

Verify that your management network is connected to a distributed switch.

Virtual machines can maintain a consistent network configuration as they migrate across multiple hosts. If you migrate an existing standard switch, or virtual adapter, to a Distributed Switch and the Distributed Switch becomes unnecessary or stops functioning, you can restore the standard switch to ensure that the host remains accessible. When you restore

the standard switch, a new virtual adapter is created and the management network uplink that is currently connected to Distributed Switch is migrated to the new virtual switch.

You might need to restore the standard switch for the following reasons:

- The Distributed Switch is not needed or is not functioning.
 - The Distributed Switch needs to be repaired to restore connectivity to vCenter Server and the hosts need to remain accessible.
 - You do not want vCenter Server to manage the host. When the host is not connected to vCenter Server, most Distributed Switch features are unavailable to the host.
1. From the direct console, select **Restore Standard Switch** and press Enter.
If the host is on a standard switch, this selection is dimmed, and you cannot select it.
 2. Press F11 to confirm.

Configuring System Logging

ESXi hosts run the syslog service (`vm syslogd`) that writes messages from system components to log files and can forward messages to syslog collectors.

You can configure the amount and location of the logs. You can also create and apply log filters to modify the logging policy of an ESXi host.

When configuring remote hosts for syslog, you also need to open some specified ports in the ESXi host firewall to allow the transmission of log messages. For more information, see [Opening the firewall for syslog emission to remote hosts](#).

Configure Syslog on ESXi Hosts

You can use the vSphere Client, the VMware Host Client, or the `esxcli system syslog` command to configure the syslog service.

The syslog service receives, categorizes, and stores log messages for analyses that help you take preventive action in your environment.

Set ESXi Syslog by using the vSphere Client

You can use the vSphere Client command to configure the syslog service globally and edit various advanced settings.

1. Browse to the ESXi host in the vSphere Client inventory.
2. Click **Configure**.
3. Under **System**, click **Advanced System Settings**.
4. Click **Edit**.
5. Filter for `syslog`.
6. To set up logging globally and configure various advanced settings, see [ESXi Syslog Options](#).
7. Optional: To overwrite the default log size and log rotation for any of the logs:
 - a) Click the name of the log that you want to customize.
 - b) Enter the number of rotations and the log size you want.
8. Click **OK**.

Changes to the syslog options take effect.

Note: Syslog parameter settings that you define by using the vSphere Client or VMware Host Client are effective immediately. However, most settings you define by using ESXCLI require an additional command to take effect. For more details, see [ESXi Syslog Options](#).

Set ESXi Syslog by using the vSphere Client

You can use the vSphere Client to configure and edit syslog service parameters on ESXi hosts.

1. In the vSphere Client, under **Host**, click **Manage > System > Advanced settings**.
2. In the **Search** panel, type a syslog setting you want to define. See [ESXi Syslog Options](#).
3. Select the setting and click **Edit option**.
4. Set the value as described in the table of parameters in [ESXi Syslog Options](#).
5. Click **Save**.

Set ESXi Syslog by using ESXCLI

You can configure the syslog service on ESXi hosts by using the ESXCLI command: `esxcli system syslog config set <syslog option>`.

For information about using the `esxcli system syslog` command and other ESXCLI commands, see [Getting Started with ESXCLI](#). For details how to open the ESXi firewall for the port specified in each remote host specification, see [Configuring the ESXi Firewall](#).

Note:

Using ESXCLI requires ESXi to open SSH logins, which is a security risk, and is not recommended. If you chose to use ESXCLI, make sure you use the `esxcli system syslog reload` command after setting each parameter to make sure that it takes effect.

Use the ESXCLI command `esxcli system syslog config set <syslog option>` to set a syslog option that you decide to enable. For example, to set the `Syslog.global.logHost` option, use the command `esxcli system syslog config set --loghost=<str>`

After setting `Syslog.global.logHost`, ESXi hosts open and maintain connections to the syslog collectors, and the transmission of messages begins immediately. When ESXi generates a syslog message, it writes it to the appropriate log file on the ESXi host and also forwards it to all configured syslog collectors.

Fine-tune Syslog on ESXi Hosts

By using the right syslog settings, you can achieve proactive monitoring of your environment, reduce downtime and take preventive action on servers.

While setting up syslog, you need to consider several parameters that affect log file retention, syslog transmission, transmission length, error handling, and the set up of SSL certificates for secure syslog message transmission. What follows are recommendations for fine-tuning your syslog parameters. You can see a description of all available parameters at [ESXi Syslog Options](#).

How to Specify Log File Retention

By default, log files cannot expand past a configured size. Once a log file reaches the configured size, logging is routed to a new log file and the oldest log file is deleted.

Note: Best practice is to balance the rotate and size settings. Increasing the rotate setting ensures that syslog files are generated often enough to prevent any potential corruption or destruction from the other log files. Increasing the size setting reduces the time for switching to another log file. Optimal size settings are a multiple of 1024 KiB.

Use the `Syslog.global.defaultSize` setting to specify the log file maximum size in KiB, and `Syslog.global.defaultRotate` to set the maximum number of old log files to keep before rotating to a new log file. To change the log file retention parameters associated with a specific program, use the `Syslog.loggers.<progName>.rotate` and

`Syslog.loggers.<progName>.size` settings, where `<progName>` is the name of the program whose parameters you want to adjust.

Manage Settings that Affect the Virtual Machine Log File

You can configure some settings that affect the virtual machine log file, `vmware.log`, either in the `vmx` file or in the `/etc/vmware/config` file. You must power off a virtual machine to edit the `vmx` file and edits take effect only on that virtual machine. If you use the `/etc/vmware/config` file, you must add the prefix `"vmx"` to the setting, for example `vmx.log.keepOld = "20"`, and edits affect all virtual machines on the ESXi host.

Table 23: Configurable settings for the `vmware.log` file

Parameter	Description	Example	Notes
<code>logging</code>	Disables all virtual machine logging.	The default value is <code>logging = "TRUE"</code> To disable virtual machine logging: <code>logging = "FALSE"</code>	Do not use this setting, because disabling virtual machine logging makes it difficult or impossible to get support for virtual machine problems. If you need to use this setting for some reason, you can only place it in the <code>vmx</code> file of a virtual machine.
<code>log.throttleBytesPerSec</code>	Controls when a log file throttles. Log file throttling occurs when writes to the <code>vmware.log</code> exceed the specified rate for a significant amount of time. This occurs when code within the VMX process, which controls a virtual machine, creates excessive log messages. The default value for this setting is 1 KB/sec. In case of log throttling, you see <code><<< Log Throttled >>></code> in the <code>vmware.log</code> file.	<code>log.throttleBytesPerSec = "1500"</code> To disable log throttling, use <code>log.throttleBytesPerSec = "0xFFFFFFFF"</code>	Log file throttling might obscure information necessary to diagnose problems with the affected virtual machine. If you need to disable log throttling, place the line in the example in the <code>vmx</code> file of the affected virtual machine. Remove the line after the debugging session ends.
<code>log.keepOld</code>	Controls the number of older <code>vmware.log</code> file to retain.	<code>log.keepOld = "20"</code>	Do not put the value of this setting below the default value (10). If virtual machines are frequently modified or moved, consider raising this setting to 20 or more.
<code>log.rotateSize</code>	Controls the maximum size of a <code>vmware.log</code> file in bytes.	<code>log.rotateSize = "2500000"</code> To disable limiting the maximum size of a <code>vmware.log</code> file, use <code>log.rotateSize = "0"</code>	A value of this setting below 100,000 can cause a loss of critical log messages and affect virtual machine performance. In ESXi 7.x and earlier, the default value of this setting places no limit on the size of a <code>vmware.log</code> file. In ESXi 8.x and later, the default value of this setting is 2,048,000.

Parameter	Description	Example	Notes
<code>log.fileName</code>	Controls the name and location of virtual machine log files.	<pre>log.fileName = "myVMLog"</pre> This setting changes the name of the virtual machine log files from <code>vmware.log</code> to <code>myVMlog</code>. <pre>log.fileName = "/vmfs/volumes/vol1/myVM/myVM.log"</pre> This setting directs virtual machine log files to a directory on a different VMFS volume (vol1) by using <code>myVM</code> for a file name.	Do not place a log file outside the virtual machine directory to make sure that the collection of host support bundles picks up the log file, which can be critical to debug virtual machine problems.
<code>log.fileLevel</code>	Controls the minimum level at which messages are written to <code>vmware.log</code>. Every log message has a level associated with it. Levels below the specified setting are not added to a virtual machine log file. The virtual machine message log levels (from most to least restricted) are: • error • warning • notice • info (default) • trivia • debug • debug1 • debug2 • debug3 • debug4 • debug5 • debug6 • debug7 • debug8 • debug9 • debug10	<pre>log.fileLevel = "debug1"</pre>	Do not set a more restrictive level than "info" to avoid filtering out messages that are necessary for debugging virtual machine problems. Lower the level below "info" only upon request by licensed support. Restore the setting to "info" after debugging ends.
<code>log.filter.minLogLevel.<groupName></code>	Controls the output of specialized debugging messages.	<pre>log.filter.minLogLevel = "debug5"</pre>	Use this setting only upon request by licensed support, who should provide one or more <code><groupName></code> parameters. Remove the setting after debugging ends.
<code>log.syslogID</code>	Activates the sending of virtual machine log messages to the system logger of an ESXi host, such as the syslog.	<pre>log.syslogID = "vmx"</pre>	Use "vmx" as value for this setting to allow the ESXi syslog daemon, <code>vmssyslogd</code> , to send these messages to a separate log file.

Parameter	Description	Example	Notes
<code>log.syslogLevel</code>	Controls the minimum level at which messages are output to the system logger of an ESXi host, such as the syslog.	<code>log.syslogLevel = "debug"</code>	The levels and functioning of this setting are identical to those for the <code>log.fileLevel</code> setting.

How to Specify Message Transmission to Remote Hosts

Optionally, you can configure ESXi to send syslog messages to one or more remote hosts, called syslog collectors, such as VMware Aria Operations for Logs (formerly VMware vRealize Log Insight and vCenter Log Insight), to collect syslog messages.

Note: Best practice is that you configure each ESXi host to send syslog messages to at least one syslog collector. This helps ensure that the messages are preserved in case of a catastrophic system event and that you can process syslog messages in various ways, such as real-time categorization and analysis (for example, by type, time span, or machine), or archive messages.

Use the `Syslog.global.logHost` setting to define remote host specifications. Separate multiple remote host specifications with a comma (,). After setting `Syslog.global.logHost`, ESXi hosts open and maintain connections to the syslog collectors, and the transmission of messages begins immediately. When ESXi generates a syslog message, it writes it to the appropriate log file on the ESXi host and also forwards it to all configured syslog collectors.

In addition to syslog messages, audit messages can also be transmitted to syslog collectors for security purposes. Audit records track security-related activity on the ESXi host. For more information about audit records, see [Audit Records](#).

Note: Consult with your company security response team if and how to set audit records. Certified configurations usually require audit records to be activated.

What follows is the syntax for `Syslog.global.logHost` remote host specifications:

```
protocol://target[:port][?formatter=value[&framing=value]]
```

Parameter	Description	Notes
<code>protocol</code>	Specifies the networking protocol. Valid values are <code>udp</code> , <code>tcp</code> , and <code>ssl</code> .	The <code>ssl</code> protocol specifies that transmission of syslog messages is encrypted. The <code>tcp</code> and <code>udp</code> protocols do not encrypt the transmission. Note: If capturing syslog messages or audit messages is critical to your system, avoid using the <code>udp</code> protocol because the networking infrastructure external to ESXi might drop UDP messages.
<code>target</code>	Specifies the remote host. You can use either an IPV4 or IPV6 address, or a host name.	When you use an IPV6 address, you must embed it in square brackets <code>[xxx]</code> , where <code>xxx</code> is the IPV6 address.
<code>port</code>	(Optional) Specifies the remote host port to use. If you use UDP or TCP, the default port is 514. If you use SSL, the default port is 1514. If you choose to use different ports from 514 or 1514, you must adjust the ESXi firewall to open the port.	For details how to open the ESXi firewall for the port specified in each remote host specification, see Configuring the ESXi Firewall .
<code>formatter</code>	Specifies how transmissions are formatted. The formatter must be RFC 3164 or RFC 5424.	RFC 3164 is the default.

Parameter	Description	Notes
framing	Specifies if transmissions are framed. Framing must be <code>non_transparent</code> or <code>octet_counting</code> .	The default is <code>non_transparent</code> . Transmissions in RFC 5424 format must specify <code>octet_counting</code> framing. For more information, see Protocols, Formats and Framing of ESXi Syslog Messages .

Examples of remote machine specifications:

Syslog.global.logHost string example	Notes
<code>tcp://10.176.130.7:12345?formatter=RFC_3164</code>	Transmits syslog messages to 10.176.130.7 using TCP/IP and port 12345. Transmission format is RFC 3164 with no framing.
<code>tcp://10.176.130.7:12345?formatter=RFC_3164&framing=octet_counting</code>	Transmits syslog messages to 10.176.130.7 using TCP/IP and port 12345. Transmission format is RFC 3164 and framing is <code>octet_counting</code> .
<code>tcp://10.176.130.7:12345?formatter=RFC_5424&framing=octet_counting</code>	Transmits syslog messages to 10.176.130.7 using TCP/IP and port 12345. Transmission format is RFC 5424 and framing is <code>octet_counting</code> .
<code>tcp://[2001:db8:85a3:8d3:1319:8a2e:370:7348]</code>	Transmits syslog messages to an IPV6 address using port 1514.
<code>tcp://[2001:db8:85a3:8d3:1319:8a2e:370:7348]:5432</code>	Transmit syslog messages to an IPV6 address using port 54321.
<code>udp://company.com</code>	Transmits syslog messages to <code>company.com</code> using UDP and port 514.
<code>udp://company.com,tcp://10.20.30.40:1050</code>	Transmits syslog messages to two remote hosts. The first remote host uses UDP to communicate with <code>company.com</code> using port 514. The second remote host uses TCP to communicate with the IPV4 address <code>10.20.30.40</code> using port 1050.
<code>ssl://company.com</code>	Transmits syslog messages to <code>company.com</code> using SSL (TLS) and port 514.

Maximum Message Transmission Length

If you use UDP, the maximum syslog message transmission length is 480 bytes for IPV4 and 1180 bytes for IPV6.

For TCP or SSL, the default maximum syslog message transmission length is 1 kibibyte (KiB). You can increase this length by using the `Syslog.global.remoteHost.maxMsgLen` parameter. The maximum value is 16 KiB. Messages longer than 16 KiB are truncated.

Note: If increasing the maximum transmission length is necessary, best practice is to increase the length only as much as specifically necessary.

Increasing the maximum syslog message length can cause problems if the networking and syslog infrastructure external to ESXi is unable to handle messages longer than 1 KiB.

Note: Best practice is that you do not use UDP to transmit syslog messages due to the packet length constraints and the possibility that the external networking infrastructure might drop the messages.

Considerations for Certificates When Configuring SSL Transmissions to Remote Hosts

When configuring ESXi to transmit syslog messages to remote hosts using SSL, you must add an SSL certificate for each remote host to the ESXi host CA store. For more information, see [Managing Certificates for ESXi Hosts](#).

Note: Consult with your syslog collector documentation on how to configure the collector for secure receipt of syslog messages using SSL and a private key.

Additional SSL Transmission Parameters

An ESXi system complying with security certification requirements might require activating X509 CRL checks. You turn on the advanced settings `Syslog.global.certificate.strictX509Compliance` and `Syslog.global.certificate.checkCRL` by changing the default value of `false` to `true`. Due to implementation limitations, if you activate CRL checks by using the setting `Syslog.global.certificate.checkSSLCerts`, then all certificates in a certificate chain must provide a CRL link. By default, the setting is active. You can deactivate SSL certificate checks by changing the setting to `false`, but this is not a best practice. You might need to turn off SSL certificate checks when troubleshooting communications with a remote host, but do this only for a limited time.

Where to Find Syslog Daemon Error and Status Information

The ESXi syslog daemon uses the log file `/var/run/log/vmsyslogd.log` to store status and error information, including dropped messages. If audit record transmission is active, the syslog daemon also emits audit records related to its operation, such as daemon start, stop, and error conditions, which allows you to verify that the syslog daemon runs properly.

How to Change the Default Syslog Log File Storage Area

The default syslog log file storage area is `/var/run/log`, local to each ESXi host. Use the `Syslog.global.logDir` syslog configuration variable to change the default syslog log file storage area, as long as the location resides on persistent storage. If `Syslog.global.logDir` is configured to a persistent store shared by multiple ESXi hosts to store their syslog log files, change the `Syslog.global.logDirUnique` setting to `true` to prevent mixing logs. The `Syslog.global.logDirUnique` setting makes sure that each ESXi machine gets a unique name added to the `Syslog.global.logDir` path, separating the log files from other hosts.

Syslog Message Queueing for Remote Hosts and Message Drops

Once syslog emissions start, they never stop except for ESXi reboots and failures, or a syslog reconfiguration to stop.

To avoid dropping messages, ESXi uses an in-memory queue that allows the `vmsyslogd` service to handle the following conditions for a short time:

- ESXi generates log messages at a faster rate than the `vmsyslogd` service can process and transmit
- Network connectivity between ESXi and the remote host fails

If either of these conditions continues for a long period, the capacity of the in-memory queue might not be sufficient and `vmsyslog` might stop transmitting messages and audit logs to remote hosts. No data is lost, because dropped messages are written to the `/var/run/log/vmsyslogd-dropped.log` file.

To minimize the chance that message dropping occurs, place the `vmsyslogd` log files on the fastest available storage and configure `vmsyslogd` and its syslog collectors on a network with end-to-end bandwidth greater than 1 GigE.

Note: Consider an end-to-end bandwidth of 2.5 GigE or more, preferably 10 GigE, to optimize performance and prevent message drops.

If you see excessive logs from an application, file a support service request for analysis and correction.

You can see statistics about message drops in the syslog daemon log file.

You can see dropped messages at `/var/run/log/vmsyslogd-dropped.log`. This log file has retention settings specific to it, similar to those for the program-specific retention parameters. The dropped messages log file retention parameters are: `Syslog.global.droppedMsgs.fileRotate` and `Syslog.global.droppedMsgs.fileSize`.

ESXi Syslog Options

You can define the behavior of ESXi syslog files and transmissions by using a set of syslog options.

Apart from the base settings, such as `Syslog.global.logHost`, starting from ESXi 7.0 Update 1, a list of advanced options is available for customizations, and NIAP compliance.

Note: Always configure persistent storage before you set any of the audit record parameters or the `Syslog.global.logDir` parameter.

Note: All audit record settings, beginning with `Syslog.global.auditRecord`, take effect immediately. However, for other settings that you define by using ESXCLI, make sure to run the `esxcli system syslog reload` command to enable the changes.

Table 24: Legacy Syslog Options

Option	ESXCLI command	Description
<code>Syslog.global.logHost</code>	<code>esxcli system syslog config set --loghost=<str></code>	Defines a comma-delimited list of remote hosts and specifications for message transmissions. If the <code>loghost=<str></code> field is blank, no logs are forwarded. While no hard limit to the number of remote hosts to receive syslog messages exists, good practice is to keep the number of remote hosts to five or less. The format of a remote host specification is: <code>protocol://hostname ipv4 '['ipv6']':port[?formatter=value[&framing=value]]</code> . The protocol must be one of <code>tcp</code> , <code>udp</code> , or <code>ssl</code> . The value of a port can be any decimal number between 1 and 65535. If a port is not provided, UDP and TCP use 514; SSL uses 1514. For example: <code>ssl://hostName1:1514</code> . The formatter must be <code>RFC_3164</code> or <code>RFC_5424</code> ; <code>RFC_3164</code> is the default. The framing must be <code>non_transparent</code> or <code>octet_counting</code> ; the default is <code>non_transparent</code> . For more information, see Fine-tune Syslog on ESXi Hosts .
<code>Syslog.global.defaultRotate</code>	<code>esxcli system syslog config set --default-rotate=<long></code>	Maximum number of old log files to keep. You can set this number globally and for individual subloggers (see <code>Syslog.global.defaultSize</code>).
<code>Syslog.global.defaultSize</code>	<code>esxcli system syslog config set --default-size=<long></code>	Default size of log files, in KiB. After a file reaches the default size, the syslog service creates a new file. You can set this number globally and for individual subloggers.

Option	ESXCLI command	Description
Syslog.global.logDir	esxcli system syslog config set --logdir=<str>	Directory where logs reside. The directory can be on mounted NFS or VMFS volumes. Only the /scratch directory on the local file system is persistent across reboots. Specify the directory as [datastore name] path_to_file, where the path is relative to the root of the volume backing the datastore. For example, the path [storage1] /systemlogs maps to the path /vmfs/volumes/storage1/systemlogs.
Syslog.global.logDirUnique	esxcli system syslog config set --logdir-unique=<bool>	Specifies the ESXi host name to be concatenated to the value of Syslog.global.logDir. It is critical that you enable this setting when multiple ESXi hosts log to a shared file system. Selecting this option creates a subdirectory with the name of the ESXi host under the directory specified by Syslog.global.LogDir. A unique directory is useful if the same NFS directory is used by multiple ESXi hosts.
Syslog.global.certificate.checkSSL	esxcli system syslog config set --check-ssl-certs=<bool>	Enforces checking of SSL certificates when transmitting messages to remote hosts.

Table 25: Syslog Options Available Starting from ESXi 7.0 Update 1

Option	ESXCLI command	Description
Syslog.global.auditRecord.storageCapacity	esxcli system auditrecords local set --size=<long>	Specifies the capacity of the audit record storage directory located on the ESXi host, in MiB. You cannot decrease the capacity of the audit record storage. You can increase the capacity before or after the audit record storage is enabled (see Syslog.global.auditRecord.storageEnable).
Syslog.global.auditRecord.remoteEnable	esxcli system auditrecords remote enable	Enables sending audit records to remote hosts. Remote hosts are specified by using the Syslog.global.logHost parameter.
Syslog.global.auditRecord.storageDirectory	esxcli system auditrecords local set --directory=<dir>	Creates an audit record storage directory and unless specified, sets /scratch/auditLog as the default location. You must not manually create an audit record storage directory and you cannot change the audit record storage directory while audit record storage is enabled (see Syslog.global.auditRecord.storageEnable).

Option	ESXCLI command	Description
<code>Syslog.global.auditRecord.storageEnable</code>	<code>esxcli system auditrecords local enable</code>	Enables the storage of audit records on an ESXi host. If the audit record storage directory does not exist, it is created with the capacity specified by <code>Syslog.global.auditRecord.storageCapacity</code> .
<code>Syslog.global.certificate.checkCRL</code>	<code>esxcli system syslog config set --crl-check=<bool></code>	Enables checking the revocation status of all the certificates in an SSL certificate chain. Enables verification of X.509 CRLs, which are not checked by default in compliance with industry conventions. A NIAP-validated configuration requires CRL checks. Due to implementation limitations, if CRL checks are enabled, then all certificates in a certificate chain must provide a CRL link. Do not enable the <code>crl-check</code> option for installations not related to certification, because of the difficulty in properly configuring an environment that uses CRL checks.
<code>Syslog.global.certificate.strictX509Compliance</code>	<code>esxcli system syslog config set --x509-strict=<bool></code>	Enables strict compliance with X.509. Performs additional validity checks on CA root certificates during verification. These checks are generally not performed, as CA roots are inherently trusted, and might cause incompatibilities with existing, misconfigured CA roots. A NIAP-validated configuration requires even CA roots to pass validations. Do not enable the <code>x509-strict</code> option for installations not related to certification, because of the difficulty in properly configuring an environment that uses CRL checks.
<code>Syslog.global.droppedMsgs.fileRotate</code>	<code>esxcli system syslog config set --drop-log-rotate=<long></code>	Specifies the number of old dropped message log files to keep.
<code>Syslog.global.droppedMsgs.fileSize</code>	<code>esxcli system syslog config set --drop-log-size=<long></code>	Specifies the size of each dropped message log file before switching to a new one, in KiB.
<code>Syslog.global.logCheckSSLCerts</code>	<code>esxcli system syslog config set --check-ssl-certs=<bool></code>	Enforces checking of SSL certificates when transmitting messages to remote hosts. Note: Deprecated. Use <code>Syslog.global.certificate.checkSSLCerts</code> in ESXi 7.0 Update 1 and later.
<code>Syslog.global.logFilters</code>	<code>esxcli system syslog config logfilter [add remove set] ...</code>	Specifies one or more log filtering specifications. Each log filter must be separated by a double vertical bar " ". The format of a log filter is: <code>numLogs ident logRegexp.numLogs</code> sets the maximum number of log entries for the

Option	ESXCLI command	Description
		specified log messages. After reaching this number, the specified log messages are filtered and ignored. <code>ident</code> specifies one or more system components to apply the filter to the log messages that these components generate. <code>logRegexp</code> specifies a case-sensitive phrase with Python regular expression syntax to filter the log messages by their content.
<code>Syslog.global.logFiltersEnable</code>		Enables the use of log filters.
<code>Syslog.global.logLevel</code>	<code>esxcli system syslog config set --log-level=<str></code>	Specifies the log filtering level. You must change this parameter only when troubleshooting an issue with the syslog daemon. You can use the values <code>debug</code> for the most detailed level, <code>info</code> for the default detail level, <code>warning</code> for only warnings or errors, or <code>error</code> , only for errors.
<code>Syslog.global.msgQueueDropMark</code>	<code>esxcli system syslog config --queue-drop-mark=<long>)</code>	Specifies the percent of the message queue capacity at which messages are dropped.
<code>Syslog.global.remoteHost.connectRetryDelay</code>	<code>esxcli system syslog config set --default-timeout=<long></code>	Specifies the delay before retrying to connect to a remote host after a connection attempt fails, in seconds.
<code>Syslog.global.remoteHost.maxMsgLen</code>	<code>esxcli system syslog config set --remote-host-max-msg-len=<long></code>	For the TCP and SSL protocols, this parameter specifies the maximum length of a syslog transmission before truncation occurs, in bytes. The default maximum length for remote host messages is 1 KiB. You can increase the maximum message length to up to 16 KiB. However, raising this value above 1 KiB does not ensure that long transmissions arrive untruncated to a syslog collector. For example, when the syslog infrastructure that issues a message is external to ESXi. This setting does not affect the UDP protocol. RFC 5426 specifies the UDP message lengths that can be safely accepted at 480 bytes for IPV4 and 1180 bytes for IPV6. Because of these restrictions, and because the networking infrastructure can arbitrary drop UDP packets, the use of UDP for transmitting critical syslog messages is not recommended.
<code>Syslog.global.vsanBacking</code>	<code>esxcli system syslog config set --vsan-backing=<bool></code>	Allows log files and the audit record storage directory to be placed on a vSAN cluster. However, enabling this parameter might cause the ESXi host to become unresponsive.

Protocols, Formats and Framing of ESXi Syslog Messages

Starting with ESXi 8.0, the ESXi syslog daemon, `vm syslogd`, uses three parameters to define how to send messages and audit records - protocol, formatting, and framing.

The supported protocols are UDP, TCP, and TLS (SSL). Formatting of syslog messages is defined by either RFC 3164 or RFC 5424. Framing specifies how a message is encapsulated. Framing of encapsulated messages is defined as transparent, also called `octet_counting`, or non-transparent, if a message is not encapsulated. Transparent framing ensures that new lines embedded in a message do not confuse a syslog collector. Syslog messages sent by using the UDP protocol are considered transparently framed; a syslog collector is expected to understand this and accept the transmission as a single message.

RFC 3164 sets the maximum total length of a syslog message at 1024 bytes, while RFC 5424 extends this maximum to 2048 bytes.

The default maximum length for remote host messages in ESXi is 1 KiB. You can increase the maximum message length to up to 16 KiB. However, raising this value above 1 KiB does not ensure that long transmissions arrive untruncated to a syslog collector. For example, when the syslog infrastructure that issues a message is external to ESXi.

Syslog messages that the `vm syslogd` transmits consist of structured data, a property list formatted in compliance with RFC 5424, and free format, or unstructured, data.

When a message is longer than the maximum length, ESXi 8.0 mitigates the message, trying to preserve as much of the structured data as possible.

When a message is mitigated, three parameters are either added to existing structured data or structured data is created to contain these parameters: `msgModified`, `remoteHostMaxMsgLen`, and `originalLen`.

The `msgModified` parameter indicates how the mitigation impacts the message: only structured data, only unstructured data, or both.

The `remoteHostMaxMsgLen` parameter specifies the maximum message length that ESXi can handle.

The `originalLen` parameter specifies the message length before it is mitigated.

Supported options for protocols, formatting and framing of ESXi syslog messages:

Formatting	Framing	UDP	TCP	SSL	Comments
Unspecified	Unspecified	Supported RFC 5426	Supported	Supported	Formatting of messages complies to RFC 3164, only timestamps are in RFC 3339 format. Structured data is prepended to each message. Framing defaults to non-transparent with TCP or SSL (TLS) and embedded newlines in structured data might corrupt messages. With UDP, packets are framed.
Unspecified	Non_transparent	Forbidden	Supported	Supported	Formatting of messages complies to RFC 3164, only timestamps are in RFC 3339 format. Structured data is prepended to each message. Framing defaults to non-transparent with TCP or SSL (TLS) and embedded newlines in structured data might corrupt messages.

Formatting	Framing	UDP	TCP	SSL	Comments
Unspecified	Octet_counting	Forbidden	Supported RFC 6587	Supported RFC 6587	Formatting of messages complies to RFC 3164, only timestamps are in RFC 3339 format. Structured data is prepended to each message.
RFC 5424	Unspecified	Supported RFC 5426	Supported RFC 5425	Supported RFC 5424	Formatting of messages complies to RFC 5424. Framing defaults to octet-counting with TCP or SSL (TLS). With UDP, framing might not be explicitly specified.
RFC 5424	Non_transparent	Forbidden	Not Supported	Not Supported	Not supported because embedded newlines in structured data might create corrupted messages.
RFC 5424	Octet_counting	Forbidden	Supported RFC 5425	Supported RFC 5425	Formatting of messages complies to RFC 5424.
RFC 3164	Unspecified	Supported RFC 5426	Supported	Supported	Formatting of messages complies to RFC 3164, only timestamps are in RFC 3339 format. Structured data is prepended to each message. Framing defaults to non-transparent with TCP or SSL (TLS) and embedded newlines in structured data might corrupt messages. With UDP, packets are framed.
RFC 3164	Non_transparent	Forbidden	Supported	Supported	Formatting of messages complies to RFC 3164, only timestamps are in RFC 3339 format. Structured data is prepended to each message. Framing defaults to non-transparent with TCP or SSL (TLS) and embedded newlines in structured data might corrupt messages.
RFC 3164	Octet_counting	Forbidden	Supported RFC 6587	Supported RFC 6587	Formatting of messages complies to RFC 3164, only timestamps are in RFC 3339 format. Structured data is prepended to each message.

Log File Formats

Starting with ESXi 8.0, the format of log files is standardized and is expressed in Augmented Backus-Naur Form (ABNF).

In ESXi 8.0, log files are written either directly, from a single service such as VMX, or indirectly, when logs from a service are submitted to a syslog. For example, VMX always writes log messages in the `vmware.log` file of each virtual machine. To spare system resources, VMX does not submit log messages to the syslog. On the other hand, in some log files that `vm syslogd` generates, you see messages from multiple programs, because the ESXi syslog daemon creates and manages all log files, and messages to these files, from multiple services.

Format for direct log messages:

Parameter	Value
LOG-MSG	HEADER SP MSG
HEADER	TIMESTAMP SP SEVERITY SP THREAD-NAME SP OPID
TIMESTAMP	FULL-DATE T FULL-TIME (Complies to RFC 5424 with the requirement for UTC/GMT formatting and resolution in milliseconds, or more granular where possible.)
FULL-DATE	DATE-FULLYEAR - DATE-MONTH - DATE-MDAY
DATE-FULLYEAR	4DIGIT
DATE-MONTH	2DIGIT ; 01-12
DATE-MDAY	2DIGIT ; 01-28, 01-29, 01-30, 01-31 based on month/year
FULL-TIME	TIME-HOUR : TIME-MINUTE : TIME-SECOND[TIME-SECFRAC] Z
TIME-HOUR	2DIGIT ; 00-23
TIME-MINUTE	2DIGIT ; 00-59
TIME-SECOND	2DIGIT ; 00-59
TIME-SECFRAC	'.' 1*6DIGIT
SEVERITY	SEVERITY-STRING SEVERITY-VALUE [LINE-MARKER]
SEVERITY-STRING	Em / Al / Cr / Er / Wa / No / In / Db (The 8 severity levels specified in RFC 5424 are abbreviated as follows: • Em - Emergency • Al - Alert • Cr - Critical • Er - Error • Wa - Warning • No - Notice • In - Informational • Db - Debug
SEVERITY-VALUE	(*DIGIT) (The SEVERITY-VALUE is an optional expression of the numeric value associated with the SEVERITY-STRING. This allows levels supported by a logger to be collapsed into the 8 required strings with no loss of information (e.g. Db(5) - debug, level 5).)
LINE-MARKER	+ (The LINE-MARKER is added to each subsequent line generated from a multi-line submission. It identifies multiline submissions and prevents a log injection security attack.)
NILVALUE	- (A single threaded program might not have a thread name and NILVALUE is acceptable.)
THREAD-NAME	NILVALUE / 1*32PRINTUSASCII (The component (APP-NAME) is implied when a single program writes the file and no component field is necessary, only the thread name.)
OPID	NILVALUE / 1*128UTF-8-STRING
STRUCTURED-DATA	1*SD-ELEMENT
SD-ELEMENT	[SD-ID *(SP SD-PARAM)]
SD-PARAM	PARAM-NAME %d34 PARAM-VALUE %d34

SD-ID	SD-NAME
PARAM-NAME	SD-NAME
PARAM-VALUE	UTF-8-STRING ; characters ", '\ and ']' MUST be escaped.
SD-NAME	1*32PRINTUSASCII ; except ", SP, ']', %d34 ()
MSG	[STRUCTURED-DATA SP] UTF-8-STRING

Format for log files that the `vmssyslogd` service manages:

Parameter	Value
LOG-MSG	HEADER SP MSG
HEADER	TIMESTAMP SP SEVERITY SP APP-NAME [PROC-IDENTIFIER] :
APP-NAME	1*32PRINTUSASCII
PROC-IDENTIFIER	[*DIGITS] ; the PID associated with APP-NAME
TIMESTAMP	FULL-DATE T FULL TIME (resolution in milliseconds or more granular where possible.)
FULL-DATE	DATE-FULLYEAR - DATE-MONTH - DATE-MDAY
DATE-FULLYEAR	4DIGIT
DATE-MONTH	2DIGIT ; 01-12
DATE-MDAY	2DIGIT ; 01-28, 01-29, 01-30, 01-31 based on month/year
FULL-TIME	TIME-HOUR : TIME-MINUTE : TIME-SECOND[TIME-SECFRAC] Z
TIME-HOUR	2DIGIT ; 00-23
TIME-MINUTE	2DIGIT ; 00-59
TIME-SECOND	2DIGIT ; 00-59
TIME-SECFRAC	'.' 1*6DIGIT
SEVERITY-STRING	Em/ Al / Cr / Er / Wa / No / In / Db (The 8 severity levels specified in RFC 5424 are abbreviated as follows: • Em - Emergency • Al - Alert • Cr - Critical • Er - Error • Wa - Warning • No - Notice • In - Informational • Db - Debug
SEVERITY	SEVERITY-STRING PRI-STRING [LINE-MARKER]
PRIVAL	1*3DIGIT ; range 0 .. 191 (the MSG PRI; contains facility and severity values, ORed together)
PRI-STRING	(PRIVAL) (The PRIVAL contains the bits from the message PRI. This allows one to see the Facility of the message, as well as the severity bits themselves._

LINE-MARKER	+ (The LINE-MARKER is added to each subsequent line generated from a multi-line submission. It identifies multiline submissions and prevents a log injection security attack.)
STRUCTURED-DATA	1*SD-ELEMENT
SD-ELEMENT	[SD-ID *(SP SD-PARAM)]
SD-PARAM	PARAM-NAME %d34 PARAM-VALUE %d34
SD-ID	SD-NAME
PARAM-NAME	SD-NAME
PARAM-VALUE	UTF-8-STRING ; characters ", '\ and ']' MUST be escaped.
SD-NAME	1*32PRINTUSASCII ; except ", SP, ']', %d34 ()
MSG	[STRUCTURED-DATA SP] UTF-8-STRING

Audit Records

ESXi audit records, with facility code 13, are compliant to both RFC 3164 and 5424 formats and you find them in the structured data section. In the audit record, you also find event-based traceability information when such data is available. Audit records are stored in a special format, not a regular log file. You can access audit records locally by using the `viewAudit` program and the Virtual Infrastructure Management functionality `FetchAuditRecords`. Do not read, use, or edit an audit record storage file directly. Locally stored audit records comply with RFC 5424 transmission format where the `HOSTNAME` and `MSGID` are always `NILVALUE`.

ESXi Syslog Message Transmission Formats

ESXi 8.0 formats syslog messages in compliance with either RFC 3164 or RFC 5424.

Note: The timestamps associated with RFC 3164 messages are in RFC 3339 format, an exception to the RFC 3164 specification.

The definition of the ESXi transmission formats for RFC 3164 and RFC 5424 is in Augmented Backus-Naur Form (ABNF).

RFC 3164 Transmission Message Format

Parameter	Value
SYSLOG-MSG	HEADER SP MSG
HEADER	PRI TIMESTAMP SP HOSTNAME SP APP-NAME [PROC-IDENTIFIER] ":"
PRI	"<" PRIVAL ">"
PRIVAL	1*3DIGIT ; range 0 .. 191 (the MSG PRI; contains facility and severity values, ORed together)
APP-NAME	1*32PRINTUSASCII
HOSTNAME	1*255PRINTUSASCII
PROC-IDENTIFIER	"[" *DIGITS "]" ; the PID associated with APP-NAME
TIMESTAMP	FULL-DATE "T" UTC-TIME (Never set a TIME-OFFSET)
FULL-DATE	DATE-FULLYEAR "-" DATE-MONTH "-" DATE-MDAY
DATE-FULLYEAR	4DIGIT
DATE-MONTH	2DIGIT ; 01-12

DATE-MDAY	2DIGIT ; 01-28, 01-29, 01-30, 01-31 based on month/year
UTC-TIME	TIME-HOUR ":" TIME-MINUTE ":" TIME-SECOND [TIME-SECFRAC] "Z"
TIME-HOUR	2DIGIT ; 00-23
TIME-MINUTE	2DIGIT ; 00-59
TIME-SECOND	2DIGIT ; 00-59
TIME-SECFRAC	"." 1*6DIGIT
STRUCTURED-DATA	1*SD-ELEMENT
SD-ELEMENT	"[" SD-ID *(SP SD-PARAM) "]"
SD-PARAM	PARAM-NAME "%d34 PARAM-VALUE %d34
SD-ID	SD-NAME
PARAM-NAME	SD-NAME
PARAM-VALUE	UTF-8-STRING ; characters '"', '\', and ']' MUST be escaped.
SD-NAME	1*32PRINTUSASCII ; except '"', SP, ']', %d34 ('')
MSG	[STRUCTURED-DATA SP] UTF-8-STRING

RFC 5424 Transmission Message Format

Parameter	Value
SYSLOG-MSG	HEADER SP STRUCTURED-DATA [SP MSG]
HEADER	PRI VERSION SP TIMESTAMP SP HOSTNAME SP APP-NAME SP PROCID SP MSGID
PRI	"<" PRIVAL ">"
PRIVAL	1*3DIGIT ; range 0 .. 191; contains facility and severity data
VERSION	NONZERO-DIGIT 0*2DIGIT
HOSTNAME	1*255PRINTUSASCII
APP-NAME	1*48PRINTUSASCII
NILVALUE	'.'
PROCID	NILVALUE *DIGITS ; the PID associated with APP-NAME
MSGID	NILVALUE
TIMESTAMP	FULL-DATE "T" UTC-TIME
FULL-DATE	DATE-FULLYEAR "-" DATE-MONTH "-" DATE-MDAY
DATE-FULLYEAR	4DIGIT
DATE-MONTH	2DIGIT ; 01-12
DATE-MDAY	2DIGIT ; 01-28, 01-29, 01-30, 01-31 based on month/year
UTC-TIME	TIME-HOUR ":" TIME-MINUTE ":" TIME-SECOND [TIME-SECFRAC] "Z"
TIME-HOUR	2DIGIT ; 00-23
TIME-MINUTE	2DIGIT ; 00-59
TIME-SECOND	2DIGIT ; 00-59

TIME-SECFRAC	"." 1*6DIGIT
STRUCTURED-DATA	NILVALUE / 1*SD-ELEMENT
SD-ELEMENT	"[" SD-ID *(SP SD-PARAM) "]"
SD-PARAM	PARAM-NAME " %d34 PARAM-VALUE %d34
SD-ID	SD-NAME
PARAM-NAME	SD-NAME
PARAM-VALUE	UTF-8-STRING ; characters '"', '\', and ']' MUST be escaped.
SD-NAME	1*32PRINTUSASCII ; except ", SP, ']', %d34 (")
MSG	MSG-UTF8
MSG-UTF8	BOM UTF-8-STRING
BOM	%xEF.BB.BF

Configure Log Filtering on ESXi Hosts

The log filtering capability lets you modify the logging policy of the syslog service that is running on an ESXi host.

You can create log filters to reduce the number of repetitive entries in the ESXi logs and to denylist specific log events entirely.

Install ESXCLI. See *Getting Started with ESXCLI*. For troubleshooting, run `esxcli` commands in the ESXi Shell.

Starting with vSphere 7.0 Update 2, you can add log filters and enable log filtering by using ESXCLI. A log filter, once established, remains in place until it is removed, even across ESXi reboots.

Log filters affect all log events that are processed by the ESXi host `vm syslogd` service, whether they are recorded to a log directory or to a remote syslog server.

You must enable the log filtering capability and reload the syslog daemon to activate the log filters on the ESXi host.

ESXCLI commands to configure log filters follow this pattern: `esxcli system syslog config logfilter {cmd} [cmd options]`.

For example, to get the list of available log filters, run the following command: `[root@xxx-xx-dhcp-xx-xx:~] esxcli system syslog config logfilter list`.

Use the `set` command to activate or deactivate log filtering: `[root@xxx-xx-dhcp-xx-xx:~] esxcli system syslog config logfilter set`.

Use the `add` command to add a log filter and the `remove` command to remove a log filter.

Use the `get` command to determine if log filtering is enabled.

A log filter is specified by three components and uses the following syntax: `numLogs | ident | logRegexp`.

Parameter	Description
<code>numLogs</code>	Specifies the number of matches of the <code>logRegexp</code> Python regular expression that will be allowed before filtering begins.
<code>ident</code>	The <code>ident</code> string is how an application identifies itself to the syslog facility. The <code>logRegexp</code> filter must be associated with the same application. You can find the <code>ident</code> string of an application by inspecting the log files in <code>/var/run/log</code> . The third field of each log file begins with the <code>ident</code> string and ends with <code>[</code> .

Parameter	Description
logRegexp	Python regular expression that identifies the messages which you want to filter out.

For example, to filter out all messages from the `hostd` daemon that contain the word "mark" after the tenth occurrence, use the following command: `esxcli system syslog config logfilter add --filter="10|Hostd|mark"`.

To remove the log filter, use the command `esxcli system syslog config logfilter remove --filter="10|Hostd|mark"`.

For more information, see [ESXi Syslog Options](#).

Set the Host Image Profile Acceptance Level

The Host Image Profile acceptance level determines which vSphere installation bundles (VIBs) are accepted for installation.

Required privileges: **Host > Configuration > SecurityProfile** and **Host > Configuration > Firewall**

VIB signatures are checked and accepted for installation based on a combination of the VIB acceptance level and the host image profile acceptance level. VIBs are tagged with an acceptance level that depends on their signature status.

1. From the vSphere Client, connect to the vCenter Server.
2. Select the host in the inventory.
3. Click the **Configure** tab.
4. Under System, select **Security Profile**.
5. Scroll down to Host Image Profile Acceptance Level, and click **Edit**.
6. Select the acceptance level and click **OK**.

Table 26: Host Image Profile Acceptance Levels

Host Image Profile Acceptance Level	Accepted Levels of VIBs
VMware Certified	VMware Certified
VMware Accepted	VMware Certified, VMware Accepted
Partner Supported	VMware Certified, VMware Accepted, Partner Supported
Community Supported	VMware Certified, VMware Accepted, Partner Supported, Community Supported

Remove All Custom Packages on ESXi

After adding custom packages, you might decide to remove them.

Before you remove custom packages, shut down or migrate running virtual machines off of the ESXi host.

1. Reboot the ESXi host.
2. In the direct console, select **Remove Custom Extensions** and press F11 to confirm.
3. Reboot the host.

All custom packages are removed.

Modify ESXi Configuration Files

You can modify ESXi configuration files in the ESXi Configuration Store (ConfigStore) by using the `/bin/configstorecli` tool.

The goal of the ESXi Shell tool `configstorecli`, introduced in ESXi 7.0 Update 1, is to manage all configurations for an ESXi host centrally, instead of using different methods and a variety of configuration files.

For more information on how to manage ConfigStore, see VMware knowledge base articles [82227](#) and [93720](#).

Deactivate Support for Non-ASCII Characters on ESXi

You can deactivate support for non-ASCII characters for virtual machine files and directory names by using two methods, depending on the ESXi version.

By default, ESXi supports the use of non-ASCII characters for virtual machine file and directory names.

If you need to deactivate support for non-ASCII characters, for ESXi 7.0 Update 2 and later, see VMware knowledge base articles [82227](#) and [93720](#).

Prior to ESXi 7.0 Update 2, you can deactivate this support by modifying the `/etc/vmware/hostd/config.xml` file with the following steps:

1. Using a text editor, open the `/etc/vmware/hostd/config.xml` file for the ESXi host.
2. Within the `<config></config>` tag, add the following code.

```
<g11nSupport>false</g11nSupport>
```
3. Save and close the file.
4. Reboot the host.

After you deactivate this support, you can still enter non-ASCII characters for virtual machine names. vSphere user interfaces display the virtual machine names in the non-ASCII characters, but ESXi converts the actual file and directory names to ASCII strings.

Reset the System Configuration

If you are having trouble determining the source of a problem with your ESXi host, you can reset the system configuration.

Before resetting the configuration, back up your ESXi configuration in case you want to restore your configuration.

Changes in the system configuration can be related to various problems, including problems with connectivity to the network and devices. Resetting the system configuration might solve such problems. If resetting the system configuration does not solve the problem, it can still rule out configuration changes made since the initial setup as the source of the problem.

When you reset the configuration, the software overrides all your configuration changes, deletes the password for the administrator account (root), and reboots the host. Configuration changes made by your hardware vendor, such as IP address settings and license configuration, might also be deleted.

Resetting the configuration does not remove virtual machines on the ESXi host. After you reset the configuration defaults, the virtual machines are not visible, but you make them visible again by reconfiguring storage and reregistering the virtual machines.

CAUTION: When you reset the configuration defaults, users accessing the host lose connectivity.

1. Back up the configuration by using the `Get-VMHostFirmwarePowerCLI` cmdlet.
2. From the direct console, select **Reset System Configuration** and press Enter.
3. Press F11 to confirm.

The system reboots after all settings are reset to the default values.

After You Install and Set Up ESXi

After you install and set up ESXi, you can use manage hosts by various interfaces, license the hosts, and back up your configuration.

After ESXi is installed and set up, you can manage the host by using the vSphere Client and vCenter Server, license the host, and back up your ESXi configuration. You can also use the VMware Host Client to connect directly to the ESXi host and to manage it. For information about installing and using the VMware Host Client, see *vSphere Single Host Management*.

Note: When you install a standalone ESXi host with TMP enabled that is not connected to a vCenter Server instance, create a backup of the ESXi configuration recovery key. To get the recovery key, run the command `esxcli system settings encryption recovery list` on the ESXi host and note it down. It is possible that the host might not be able to complete booting due to host configuration encryption related problems and in such cases, you can restore the host configuration by using the recovery key and running the configuration recovery process.

For best practices and VMware recommendations, see [Best Practices for Secure ESXi Configuration](#).

Licensing ESXi Hosts

After you install ESXi, it has a 60-day evaluation period during which you can explore the full set of vSphere features provided with a vSphere Enterprise Plus license. You must assign the host an appropriate license before the evaluation period expires.

ESXi hosts are licensed with vSphere licenses that have per-CPU capacity. To license hosts correctly, you must assign them a vSphere license that has enough CPU capacity to cover all CPUs in the hosts. The license must support all features that the hosts are using. For example, if the hosts are connected to a vSphere Distributed Switch, you must assign a license that has the vSphere Distributed Switch feature.

You can use one of following methods to license ESXi hosts:

- License multiple hosts at a time by using the license management function in the vSphere Client. The hosts must be connected to a vCenter Server system. For more information, see *vCenter Server and Host Management*.
- Set up bulk licensing by using PowerCLI commands. Bulk licensing works for all ESXi hosts, but is especially useful for hosts provisioned with Auto Deploy. See [Set Up Bulk Licensing](#)
- License individual ESXi hosts by using a direct connection with the VMware Host Client. For information about assigning a license key to an ESXi host, see *vSphere Single Host Management*.

For more information, see [Licensing and Subscription in vSphere](#).

Recording the License Key of an ESXi Host

If a host becomes inaccessible or unbootable, you should have a record of its license key. You can write down the license key and tape it to the server, or put the license key in a secure location. You can access the license key from the direct console user interface or the vSphere Client.

View the License Keys of ESXi Hosts from the vSphere Client

You can view the license keys of the hosts that are connected to a vCenter Server system through the vSphere Client.

1. In the vSphere Client, select **Administration**.
2. Under Licensing, select **Licenses**.
3. On the **Assets** tab, select **Hosts**.
4. In the License column, click a license.

You view information about the license, such as its usage and license key.

Access the ESXi License Key from the Direct Console

If you have physical access to the host or remote access to the direct console, you can use the direct console to access the ESXi license key.

- From the direct console, select **View Support Information**.

The license key appears in the form XXXXX-XXXXX-XXXXX-XXXXX-XXXXX, labeled License Serial Number.

Note: The physical machine serial number also appears, labeled Serial Number. Do not confuse the license key with the physical machine serial number.

View System Logs

System logs provide detailed information about system operational events.

1. From the direct console, select **View System Logs**.
2. Press a corresponding number key to view a log.
vCenter Server Agent (vpxa) logs appear if you add the host to vCenter Server.
3. Press Enter or the spacebar to scroll through the messages.
4. Perform a regular expression search.
 - a) Press the slash key (/).
 - b) Type the text to find.
 - c) Press Enter.

The found text is highlighted on the screen.
5. Press q to return to the direct console.

See also [Configure Syslog on ESXi Hosts](#).

Troubleshooting ESXi Booting

The ESXi booting troubleshooting topics provide solutions to problems that you might encounter during the ESXi booting.

- [Host Stops Unexpectedly at Bootup When Sharing a Boot Disk with Another Host](#)
- [Host Fails to Boot After You Install ESXi in UEFI Mode](#)

Host Stops Unexpectedly at Bootup When Sharing a Boot Disk with Another Host

When more than one host, either physical or virtual, boots from the same shared physical disk or LUN, they cannot use the same scratch partition.

The host stops at bootup when sharing a boot disk with another host.

More than one ESXi host can share the same physical disk or LUN. When two such hosts also have the same scratch partition configured, either of the hosts can fail at bootup.

1. Set the hosts to boot sequentially, and boot the hosts.

This setting lets you start the hosts so that you can change the scratch partition for one of them.

2. From the vSphere Client, connect to the vCenter Server.
3. Select the host in the inventory.
4. Click the **Configure** tab.
5. Under System, select **Advanced System Settings**.

6. Select **ScratchConfig**.

The **ScratchConfig.CurrentScratchLocation** text box shows the current location of the scratch partition.

7. In the **ScratchConfig.ConfiguredScratchLocation** text box, enter a directory path that is unique for this host.
For example, `/vmfs/volumes/DatastoreUUID/DatastoreFolder`.
8. Reboot the host for the changes to take effect.

Host Fails to Boot After You Install ESXi in UEFI Mode

After you install ESXi on a host machine in UEFI mode, the machine might fail to boot.

When you install or upgrade ESXi, the installer tries to create a UEFI boot option named VMware ESXi and make it the default boot option. When you reboot after installing ESXi the reboot might fail. This problem is accompanied by an error message similar to `No boot device available`.

- When the installer creates the UEFI boot option, a problem occurs while writing to the NVRAM on the host motherboard.
- The host firmware does not recognize the attempt to set the UEFI boot option as the first boot option, or the firmware overrides the boot order.
- Either or both BIOS and firmware are outdated.
- The boot disk has an MBR or MSDOS partition table. Due to a technical limitation, the UEFI boot option is only created for a GUID Partition Table (GPT) partition table.

Note: UEFI firmware attempts to load the boot image from the EFI system partition, which is FAT based, on the disk. Booting from the EFI system partition only works if the disk is laid out by using a GPT. If the boot disk has an MBR or MSDOS partition table, a UEFI boot fails. You cannot add a boot entry for MBR. If the disk is fully consumed by ESXi, it cannot be converted to GPT, and you must boot in legacy BIOS mode.

1. While the error message is displayed on screen, open the boot options menu. Depending on your system, the boot options menu might open with a keyboard shortcut, in the BIOS menu, or in a BMC, iLO or iDRAC interface.
2. Check if a boot option VMware ESXi exists and try to boot from it. If the boot is successful, change the boot order and set VMware ESXi as the first boot option.
3. If the problem is not resolved, select an option similar to **Add boot option**.

The wording and location of the option might vary, depending on your system.

4. Select the file `\EFI\BOOT\BOOTx64.EFI` on the disk that you installed ESXi on.
5. Change the boot order so that the host boots from the option that you added.

Decommission an ESXi Host

If you do not want your server to be an ESXi host, you can decommission the ESXi host machine.

1. Remove VMFS datastores on the internal disks so that the internal disks are no longer set up to store virtual machines.
2. Change the boot setting in the BIOS so that the host no longer boots into ESXi.
If you installed ESXi in UEFI mode, delete the boot option VMware ESXi or any other boot option created manually.
3. Install another operating system in its place.

VMware ESXi Upgrade

VMware ESXi Upgrade describes how to upgrade VMware ESXi™ to the current version.

Intended Audience

VMware ESXi Upgrade is for anyone who needs to upgrade from earlier versions of ESXi. These topics are for experienced Microsoft Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

vSphere Upgrade Options


vSphere 8.0 provides many options for upgrading your vSphere deployment. For a successful vSphere upgrade, you must understand the upgrade options, the configuration details that impact the upgrade process, and the sequence of tasks.

The two core components of vSphere are VMware ESXi™ and VMware vCenter Server™. ESXi is the virtualization platform on which you can create and run virtual machines and virtual appliances. vCenter Server is a management plane that acts as a central administrator for ESXi hosts connected in a network. You use the vCenter Server system to pool and manage the resources of multiple hosts. vCenter Server appliance is a preconfigured virtual machine optimized to run vCenter Server.

You can upgrade existing vCenter Server deployments that include either an embedded or an external Platform Services Controller to a deployment consisting of a vCenter Server appliance.

Overview of the vSphere Upgrade Process

vSphere is a sophisticated product with multiple components to upgrade. Understanding the required sequence of tasks is vital for a successful vSphere upgrade.

Figure 23: Overview of vSphere Upgrade Tasks

Upgrading vSphere includes the following tasks:

1. Read the vSphere release notes.
2. Verify that you have backed up your configuration.
3. If your vSphere system includes VMware solutions or plug-ins, verify that they are compatible with the vCenter Server appliance version to which you are upgrading. See *VMware Product Interoperability Matrix* at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php.
4. Upgrade vCenter Server.
For detailed instructions, see [vCenter Upgrade](#).
5. Upgrade your ESXi hosts. See [Overview of the ESXi Host Upgrade Process](#).


6. To ensure sufficient disk storage for log files, consider setting up a syslog server for remote logging. Setting up logging on a remote host is especially important for hosts with a limited amount of local storage.
See [Required Free Space for System Logging](#) and [Configure Syslog on ESXi Hosts](#).
7. Upgrade your VMs manually or by using vSphere Lifecycle Manager to perform an orchestrated upgrade.
See [Upgrading Virtual Machines and VMware Tools](#)

Overview of the ESXi Host Upgrade Process

VMware provides several ways to upgrade ESXi hosts with earlier versions to ESXi version 8.0.

The details and level of support for an upgrade to ESXi8.0 depend on the host to be upgraded and the upgrade method that you use. Verify that the upgrade path from your current version of ESXi to the version to which you are upgrading, is supported. For more information, see the VMware Product Interoperability Matrices at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php.

You can upgrade an ESXi host to version 8.0 by using an interactive upgrade from a CD, DVD, or USB, a scripted upgrade, ESXCLI, or vSphere Lifecycle Manager. When you upgrade an ESXi host that has custom VIBs to version 8.0, all supported custom VIBs are migrated. For more information, see [Upgrading Hosts That Have Third-Party Custom VIBs](#).

Figure 24: Overview of the ESXi Host Upgrade Process

The following high-level steps are for upgrading ESXi.

1. Verify that your system meets the upgrade requirements. See [ESXi Requirements](#).
2. Prepare your environment before upgrading. See [Before Upgrading ESXi Hosts](#).
3. Determine where you want to locate and boot the ESXi installer. See [Media Options for Booting the ESXi Installer](#). If you are network booting the installer, verify that your network boot infrastructure is properly set up. See [Network Booting the ESXi Installer](#).
4. Upgrade ESXi. See [Upgrading ESXi Hosts](#).
5. After upgrading ESXi hosts, you must reconnect the hosts to the vCenter Server and reapply the licenses. See [After You Upgrade ESXi Hosts](#).

The following methods are supported for a direct upgrade to ESXi8.0.

- Use the interactive graphical user interface (GUI) installer from a CD, DVD, or USB drive.
- Perform a scripted upgrade.
- Use ESXCLI.
- Use vSphere Auto Deploy. If the ESXi host is deployed by using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with a 8.0 image.
- Use the vSphere Lifecycle Manager.

Graphical User Interface (GUI) Installer

You can upgrade interactively by using an ESXi installer ISO image on a CD, DVD, or USB flash drive or by network booting the installer. This method is appropriate for deployments with a few hosts. If during the installation process, you select a target disk that contains an ESXi installation, the installer upgrades the host to ESXi version 8.0. The installer also gives you the option to migrate some existing host settings and configuration files and to preserve the existing VMFS datastore. See [Upgrade Hosts Interactively](#).

Perform a Scripted Upgrade

To do a scripted upgrade, you can use the ESXi8.0 installer from a CD, DVD, or USB flash drive or by network booting the installer. This method is an efficient way to deploy multiple hosts. For more information, see [Installing or Upgrading Hosts by Using a Script](#). You can use ESXCLI to upgrade ESXi6.7 hosts or ESXi7.0 hosts to ESXi8.0 hosts.

ESXCLI

vSphere 8.0 introduces configuration files, components, base images, and add-ons as new software deliverables that you can use to update or patch ESXi8.0 hosts. For information about managing components, base images, and add-ons on ESXi, see [ESXCLI Concepts and Examples](#).

To use ESXCLI commands, you must install Standalone ESXCLI. For more information about installing and using ESXCLI, see the following documents.

- [Getting Started with ESXCLI](#)
- [ESXCLI Reference](#)

See [Upgrading Hosts by Using ESXCLI Commands](#).

vSphere Auto Deploy

If an ESXi host is deployed with vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host and reboot it with a new image profile or a configuration that you manage on a cluster level. An image profile contains an ESXi upgrade or patch, a host configuration profile, and optionally, third-party drivers or management agents that are provided by VMware partners. To add ESXi hosts to a cluster that manages ESXi configuration at a cluster level, you create a rule in Auto Deploy that assigns such a cluster as the host location for newly added hosts, which inherit

vSphere Lifecycle Manager

the same settings and do not require manual configuration. You can build custom images by using vSphere ESXi Image Builder CLI. For more information, see [Using vSphere Auto Deploy to Reprovision Hosts](#).

vSphere Lifecycle Manager is a vCenter Server service for installing, upgrading, and updating ESXi hosts. By using images and baselines, vSphere Lifecycle Manager enables centralized and simplified lifecycle management for multiple ESXi hosts at a cluster level. For more information about performing orchestrated installations, upgrades, and updates, see the *Managing Host and Cluster Lifecycle* documentation.

Upgrading Virtual Machines and VMware Tools

After you upgrade an ESXi host, you can upgrade the virtual machines on the host to take advantage of new features.

You have the following tools for upgrading virtual machines.

vSphere Client

You can use the vSphere Client to upgrade a virtual machine step by step. For more information about upgrading virtual machines, see the *vSphere Virtual Machine Administration* documentation.

vSphere Lifecycle Manager

You can use the vSphere Lifecycle Manager to upgrade the virtual machine hardware and VMware Tools versions of the virtual machines in your environment. The vSphere Lifecycle Manager automates the upgrade process and verifies that the steps occur in the correct order. For more information, see the *Managing Host and Cluster Lifecycle* documentation.

Upgrading ESXi Hosts

After you upgrade vCenter Server, upgrade your ESXi hosts. You can upgrade ESXi6.7 and 7.0 hosts directly to ESXi8.0.

To upgrade hosts, you can use the tools and methods that are described in [Overview of the ESXi Host Upgrade Process](#).

CAUTION: If you upgrade hosts managed by vCenter Server, you must upgrade vCenter Server before you upgrade the ESXi hosts. If you do not upgrade your environment in the correct order, you can lose data and lose access to servers.

ESXi Requirements

To install or upgrade ESXi, your system must meet specific hardware and software requirements.

- [ESXi System Storage Overview](#)
- [ESXi Hardware Requirements](#)
- [Using Remote Management Applications](#)
- [Recommendations for Enhanced ESXi Performance](#)
- [Incoming and Outgoing Firewall Ports for ESXi Hosts](#)
- [Required Free Space for System Logging](#)
- [VMware Host Client System Requirements](#)
- [ESXi Passwords and Account Lockout](#)

ESXi System Storage Overview

ESXi8.0 has a system storage layout that allows flexible management of partitions and third-party components, while facilitating debugging.

ESXi System Storage

The ESXi8.0 system storage layout consists of four partitions:

Table 27: ESXi system storage partitions:


Partition	Use	Type
System Boot	Stores boot loader and EFI modules.	FAT16
Boot-bank 0	System space to store ESXi boot modules.	FAT16
Boot-bank 1	System space to store ESXi boot modules.	FAT16
ESX-OSData	Acts as the unified location to store additional modules. Not used for booting and virtual machines. Consolidates the legacy <code>/scratch</code> partition, locker partition for VMware Tools, and core dump destinations. CAUTION: In case the installation media is a USB or an SD card device, best practice is to create ESX-OSData partitions on persistent storage device that is not shared between ESXi hosts.	VMFS-L

The ESX-OSData volume is divided into two high-level categories of data, persistent and non-persistent data. Persistent data contains of data written infrequently, for example, VMware Tools ISOs, configurations, and core dumps.

Non-persistent data contains of frequently written data, for example, logs, VMFS global traces, vSAN Entry Persistence Daemon (EPD) data, vSAN traces, and real-time databases.

Figure 25: Consolidated system storage in ESXi 8.0

ESXi 6.x
system storage layout



ESXi System Storage Sizes

Partition sizes, except for the system boot partition, can vary depending on the size of the boot media used. If the boot media is a high-endurance one with capacity larger than 142 GB, a VMFS datastore is created automatically to store virtual machine data.

You can review the boot media capacity and the automatic sizing as configured by the ESXi installer by using the vSphere Client and navigating to the **Partition Details** view. Alternatively, you can use ESXCLI, for example the `esxcli storage filesystem list` command.

Table 28: ESXi System Storage Sizes, Depending on the Used Boot Media and Its Capacity.

Boot Media Size	8-10 GB	10-32 GB	32-128 GB	>128 GB
System Boot	100 MB	100 MB	100 MB	100 MB
Boot-bank 0	500 MB	1 GB	4 GB	4 GB
Boot-bank 1	500 MB	1 GB	4 GB	4 GB
ESX-OSData	remaining space	remaining space	remaining space	up to 128 GB
VMFS datastore				remaining space for media size > 142 GB

You can use the ESXi installer boot option `systemMediaSize` to limit the size of system storage partitions on the boot media. If your system has a small footprint that does not require the maximum of 128 GB of system storage size, you can limit it to the minimum of 32 GB. The `systemMediaSize` parameter accepts the following values:

- min (32 GB, for single disk or embedded servers)
- small (64 GB, for servers with at least 512 GB of RAM)
- default (128 GB)
- max (consume all available space, for multi-terabyte servers)

The selected value must fit the purpose of your system. For example, a system with 1 TB of memory must use the minimum of 64 GB for system storage. To set the boot option at install time, for example `systemMediaSize=small`, refer to [Enter Boot Options to Run an Installation or Upgrade Script](#) [Boot Option](#). For more information, see Knowledge Base article [81166](#).

ESXi System Storage Links

The sub-systems that require access to the ESXi partitions, access these partitions by using the following symbolic links:

Table 29: ESXi system storage symbolic links.

System Storage Volume	Symbolic Link
Boot-bank 0	/bootbank
Boot-bank 1	/altbootbank
Persistent data	/productLocker /locker /var/core /usr/lib/vmware/isoimages /usr/lib/vmware/floppies
Non-persistent data	/var/run /var/log

System Storage Volume	Symbolic Link
	/var/vmware
	/var/tmp
	/scratch

Storage Behavior

When you start ESXi, the host enters an autoconfiguration phase during which system storage devices are configured with defaults.

When you reboot the ESXi host after installing the ESXi image, the host configures the system storage devices with default settings. Starting with ESXi 7.0, you can activate the option `autoPartition`, which automatically formats all available empty devices with VMFS, except for legacy SD and USB devices. The default is `autoPartition=FALSE`, which formats with VMFS only boot devices with size larger than 128 GB. For more information, see VMware knowledge base article [77009](#).

CAUTION: ESXi overwrites any disks that appear to be blank. Disks are considered to be blank if they do not have a valid partition table or partitions. If you are using software that uses such disks, in particular if you are using logical volume manager (LVM) instead of, or in addition to, conventional partitioning schemes, ESXi might cause local LVM to be reformatted. Back up your system data before you power on ESXi for the first time.

On the hard drive or USB device that the ESXi host is booting from, the disk-formatting software retains existing diagnostic partitions that the hardware vendor creates. In the remaining space, the software creates the partitions described below.

Partitions Created by ESXi on the Host Drive

For fresh installations, several new partitions are created for the system boot, boot banks, and ESX-OSData. Fresh ESXi installations use GUID Partition Tables (GPT) instead of MSDOS-based partitioning. The installer creates boot banks of varying size depending on the size of the disk. For more information on the scratch partition see [About the Scratch Partition](#).

The installer affects only the installation disk. The installer does not affect other disks of the server. When you install on a disk, the installer overwrites the entire disk. When the installer autoconfigures storage, the installer does not overwrite hardware vendor partitions.

To create the VMFS datastore, the ESXi installer requires a minimum of 128 GB of free space on the installation disk.

You might want to override this default behavior if, for example, you use shared storage devices instead of local storage. To prevent automatic disk formatting, detach the local storage devices from the host under the following circumstances:

- Before you start the host for the first time.
- Before you start the host after you reset the host to the configuration defaults.

To override the VMFS formatting if automatic disk formatting already occurred, you can remove the datastore. See the *vCenter Server and Host Management* documentation.

About the Scratch Partition

For new installations of ESXi, during the autoconfiguration phase, a scratch partition is created as part of the ESX-OSDATA partition.

Note: Partitioning for hosts that are upgraded to ESXi 7.0 and later from earlier versions differs significantly from partitioning for new installations of ESXi. The upgrade process to ESXi 7.0 and later repartitions the boot device and consolidates the original core dump, locker, and scratch partitions into the ESX-OSData volume.

The scratch partition serves to store system logs, which you need when you create a support bundle. If the scratch partition is not present, system logs are stored in a ramdisk. If no scratch partition is created, you can configure one. You can also override the default configuration.

You can create the scratch partition on a remote SAN or NFS-mounted directory.

Set the Scratch Partition from the vSphere Client

If a scratch partition is not set up, you might want to configure one, especially if the host is low on memory. When a scratch partition is not present, system logs are stored in a ramdisk.

The directory to use for the scratch partition must exist on the host.

1. From the vSphere Client, connect to the vCenter Server.
2. Select the host in the inventory.
3. Click the **Configure** tab.
4. Select **System**.
5. Select **Advanced System Settings**.

The setting **ScratchConfig.CurrentScratchLocation** shows the current location of the scratch partition.

6. In the **ScratchConfig.ConfiguredScratchLocation** text box, enter a directory path that is unique for this host. For example, `/vmfs/volumes/DatastoreUUID/DatastoreFolder`.
7. Reboot the host for the changes to take effect.

ESXi Hardware Requirements

Make sure that the host meets the minimum hardware configurations supported by ESXi8.0.

Hardware and System Resources

To install or upgrade ESXi, your hardware and system resources must meet the following requirements:

- Supported server platform. For a list of supported platforms, see the [VMware Compatibility Guide](#).
- ESXi8.0 requires a host with at least two CPU cores.
- ESXi8.0 supports a broad range of multi-core of 64-bit x86 processors. For a complete list of supported processors, see the [VMware Compatibility Guide](#).
- ESXi8.0 requires the NX/XD bit to be enabled for the CPU in the BIOS.
- ESXi8.0 requires a minimum of 8 GB of physical RAM. Provide at least 12 GB of RAM to run virtual machines in typical production environments.
- To support 64-bit virtual machines, support for hardware virtualization (Intel VT-x or AMD RVI) must be enabled on x64 CPUs.
- One or more Gigabit or faster Ethernet controllers. For a list of supported network adapter models, see the [VMware Compatibility Guide](#).
- ESXi8.0 requires a boot disk of at least 32 GB of persistent storage such as HDD, SSD, or NVMe. A boot device must not be shared between ESXi hosts.
- SCSI disk or a local, non-network, RAID LUN with unpartitioned space for the virtual machines.
- For Serial ATA (SATA), a disk connected through supported SAS controllers or supported on-board SATA controllers. SATA disks are considered remote, not local. These disks are not used as a scratch partition by default because they are seen as remote.

Note: You cannot connect a SATA CD-ROM device to a virtual machine on an ESXi host. To use the SATA CD-ROM device, you must use IDE emulation mode.

Storage Systems

For a list of supported storage systems, see the [VMware Compatibility Guide](#). Starting with ESXi 8.0, you cannot use software adapters for Fibre Channel over Ethernet (FCoE), only hardware FCoE adapters.

ESXi Booting Requirements

In vSphere 8.0, support for legacy BIOS is limited and booting ESXi hosts from the Unified Extensible Firmware Interface (UEFI) is recommended. With UEFI, you can boot systems from hard drives, CD-ROM drives, or USB media. vSphere Auto Deploy supports network booting and provisioning of ESXi hosts with UEFI. If your system has supported data processing units (DPU), you can only use UEFI to install and boot ESXi on the DPUs. For more information on VMware plans to deprecate support for legacy BIOS in server platforms, see [Deprecation of legacy BIOS support in vSphere](#).

ESXi can boot from a disk larger than 2 TB if the system firmware and the firmware on any add-in card that you are using support it. See the vendor documentation.

Storage Requirements for ESXi8.0 Installation or Upgrade

For best performance of an ESXi8.0 installation, use a persistent storage device that is a minimum of 32 GB for boot devices. Upgrading to ESXi8.0 requires a boot device that is a minimum of 8 GB. When booting from a local disk, SAN or iSCSI LUN, at least a 32 GB disk is required to allow for the creation of system storage volumes, which include a boot partition, boot banks, and a VMFS-L based ESX-OSData volume. The ESX-OSData volume takes on the role of the legacy `/scratch` partition, locker partition for VMware Tools, and core dump destination.

Note: In ESXi8.0, the ESX-OSData volume is considered a unified partition and the separate components, such as `/scratch` and VMware Tools, are consolidated into a single persistent OSData partition.

Other options for best performance of an ESXi8.0 installation are the following:

- A local disk of 128 GB or larger for optimal support of ESX-OSData. The disk contains the boot partition, ESX-OSData volume and a VMFS datastore.
- A device that supports the minimum of 128 terabytes written (TBW).
- A device that delivers at least 100 MB/s of sequential write speed.
- To provide resiliency in case of device failure, a RAID 1 mirrored device is recommended.

Legacy SD and USB devices are supported with the following limitations:

- SD and USB devices are supported for boot bank partitions. The use of SD and USB devices for storing ESX-OSData partitions is being deprecated and best practice is to provide a separate persistent local device with a minimum of 32 GB to store the ESX-OSData volume. The persistent local boot device can be an industrial grade M.2 flash (SLC and MLC), SAS, SATA, HDD, SSD, or a NVMe device. The optimal capacity for persistent local devices is 128 GB.
- If you do not provide persistent storage, you see an alarm such as `Secondary persistent device not found`. Please move installation to persistent storage as support for SD-Card/USB only configuration is being deprecated.
- You must use an SD flash device that is approved by the server vendor for the particular server model on which you want to install ESXi on an SD flash storage device. You can find a list of validated devices on [partnerweb.vmware.com](#).
- See [SD card/USB boot device revised guidance](#) on updated guidance for SD card or USB-based environments.
- To choose a proper SD or USB boot device, see [Boot device guidance for low endurance media \(vSphere and vSAN\)](#).

The upgrade process to ESXi8.0 from versions earlier than 7.x repartitions the boot device and consolidates the original core dump, locker, and scratch partitions into the ESX-OSData volume.

The following events occur during the repartitioning process:

- If a custom core dump destination is not configured, then the default core dump location is a file in the ESX-OSData volume.

- If the syslog service is configured to store log files on the 4 GB VFAT scratch partition, the log files in `var/run/log` are migrated to the ESX-OSData volume.
- VMware Tools are migrated from the locker partition and the partition is wiped.
- The core dump partition is wiped. The application core dump files that are stored on the scratch partition are deleted.

Note: Rollback from ESXi 8.x to a version of ESXi earlier than 7.x is not possible due to the repartitioning process of the boot device. To use a version of ESXi earlier than 7.x after upgrading to version 8.0, you must create a backup of the boot device before the upgrade, and restore the ESXi boot device from the backup. Rollback from ESXi 8.x to 7.x is possible as long as no changes to the bootbank partitions have been made and no corrupt partition is detected.

If you use USB or SD devices to perform an upgrade, best practice is to allocate an ESX-OSData region on an available persistent disk or a SAN LUN. If persistent storage or a SAN LUN are not available, ESX-OSData is automatically created on a RAM disk. VMFS can also be used for ESX-OSData partition.

After upgrade, if ESX-OSData resides on a RAM disk and a new persistent device is found on subsequent boots, and this device has the setting `autoPartition=True`, ESX-OSData is automatically created on the new persistent device. ESX-OSData does not move between persistent storage automatically, but you can manually change the ESX-OSData location on a supported storage.

To configure the size of ESXi system partitions, you can use the `systemMediaSize` option. For more information, see [Boot option to configure the size of ESXi system partitions](#).

In Auto Deploy installations, the installer attempts to allocate a scratch region on an available local disk or datastore. If no local disk or datastore is found, installation fails.

For environments that boot from a SAN or use Auto Deploy, the ESX-OSData volume for each ESXi host must be set up on a separate SAN LUN.

Using Remote Management Applications

Remote management applications allow you to install ESXi on servers that are in remote locations.

Remote management applications supported for installation include HP Integrated Lights-Out (iLO), Dell Remote Access Card (DRAC), IBM management module (MM), and Remote Supervisor Adapter II (RSA II). For support on remote management applications, contact the vendor.

You can use remote management applications to do both interactive and scripted installations of ESXi remotely.

If you use remote management applications to install ESXi, the virtual CD might encounter corruption problems with systems or networks operating at peak capacity. If a remote installation from an ISO image fails, complete the installation from the physical CD media.

Supported Remote Management Server Models and Firmware Versions

You can use remote management applications to install or upgrade ESXi, or to manage hosts remotely.

Table 30: Supported Remote Management Server Models and Minimum Firmware Versions

Remote Management Server Model	Firmware Version	Java
Dell DRAC 9	6.0.30.00	N/A
Dell DRAC 7	1.30.30 (Build 43)	1.7.0_60-b19
Dell DRAC 6	1.54 (Build 15), 1.70 (Build 21)	1.6.0_24
Dell DRAC 5	1.0, 1.45, 1.51	1.6.0_20, 1.6.0_203
Dell DRAC 4	1.75	1.6.0_23

Remote Management Server Model	Firmware Version	Java
Fujitsu iRMC S5	1.10P	1.7.0_60-b19
Fujitsu iRMC S6	1.06S	N/A
HP ILO	1.81, 1.92	1.6.0_22, 1.6.0_23
HP ILO 2	1.8, 1.81	1.6.0_20, 1.6.0_23
HP ILO 3	1.28	1.7.0_60-b19
HP ILO 4	1.13	1.7.0_60-b19
HP ILO 5	2.72	N/A
IBM RSA 2	1.03, 1.2	1.6.0_22

Recommendations for Enhanced ESXi Performance

To enhance performance, install or upgrade ESXi on a robust system with more RAM than the minimum required and with multiple physical disks.

For ESXi system requirements, see [ESXi Hardware Requirements](#).

Table 31: Recommendations for Enhanced Performance

System Element	Recommendation
RAM	ESXi hosts require more RAM than typical servers. ESXi 8.0 requires a minimum of 8 GB of physical RAM. Provide at least 12 GB of RAM to take full advantage of ESXi features and run virtual machines in typical production environments. An ESXi host must have sufficient RAM to run concurrent virtual machines. The following examples are provided to help you calculate the RAM required by the virtual machines running on the ESXi host. Operating four virtual machines with Red Hat Enterprise Linux or Windows XP requires at least 3 GB of RAM for baseline performance. This figure includes 1024 MB for the virtual machines, 256 MB minimum for each operating system as recommended by vendors. Running these four virtual machines with 512 MB RAM requires that the ESXi host have 4 GB RAM, which includes 2048 MB for the virtual machines. These calculations do not include possible memory savings from using variable overhead memory for each virtual machine. See <i>vSphere Resource Management</i>.
Dedicated Fast Ethernet adapters for virtual machines	Place the management network and virtual machine networks on different physical network cards. Dedicated Gigabit Ethernet cards for virtual machines, such as Intel PRO 1000 adapters, improve throughput to virtual machines with high network traffic.
Disk location	Place all data that your virtual machines use on physical disks allocated specifically to virtual machines. Performance is better when you do not place your virtual machines on the disk containing the ESXi boot image. Use physical disks that are large enough to hold disk images that all the virtual machines use.

System Element	Recommendation
VMFS6 partitioning	The ESXi installer creates the initial VMFS volumes on the first blank local disk found. To add disks or modify the original configuration, use the vSphere Client. This practice ensures that the starting sectors of partitions are 64K-aligned, which improves storage performance. Note: For SAS-only environments, the installer might not format the disks. For some SAS disks, it is not possible to identify whether the disks are local or remote. After the installation, you can use the vSphere Client to set up VMFS.
Processors	Faster processors improve ESXi performance. For certain workloads, larger caches improve ESXi performance.
Hardware compatibility	Use devices in your server that are supported by ESXi drivers. See the <i>Hardware Compatibility Guide</i> at http://www.vmware.com/resources/compatibility.

Incoming and Outgoing Firewall Ports for ESXi Hosts

Open and close firewall ports for each service by using either the vSphere Client or the VMware Host Client.

includes a firewall that is enabled by default. At installation time, the firewall is configured to block incoming and outgoing traffic, except traffic for services that are enabled in the host's security profile. For the list of supported ports and protocols in the firewall, see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com/>.

The VMware Ports and Protocols Tool lists port information for services that are installed by default. If you install other VIBs on your host, additional services and firewall ports might become available. The information is primarily for services that are visible in the but the VMware Ports and Protocols Tool includes some other ports as well.

Required Free Space for System Logging

See the recommended minimum size and rotation configuration for hostd, vpxa, and fdm logs.

If you used Auto Deploy to install your ESXi8.0 host, or if you set up a log directory separate from the default location in a scratch directory on the VMFS volume, you might need to change your current log size and rotation settings to ensure that enough space is available for system logging. All vSphere components use this infrastructure. The default values for log capacity in this infrastructure vary, depending on the amount of storage available and on how you have configured system logging. Hosts that are deployed with Auto Deploy store logs on a RAM disk, which means that the amount of space available for logs is small.

If your host is deployed with Auto Deploy, reconfigure your log storage in one of the following ways:

- Redirect logs over the network to a remote collector.
- Redirect logs to a NAS or NFS store.

If you redirect logs to non-default storage, such as a NAS or NFS store, you might also want to reconfigure log sizing and rotations for hosts that are installed to disk.

You do not need to reconfigure log storage for ESXi hosts that use the default configuration, which stores logs in a scratch directory on the VMFS volume. For these hosts, ESXi8.0 configures logs to best suit your installation, and provides enough space to accommodate log messages.

Table 32: Recommended Minimum Size and Rotation Configuration for hostd, vpxa, and fdm Logs

Log	Maximum Log File Size	Number of Log Files to Preserve	Minimum Disk Space Required
Management Agent (hostd)	10 MB	10	100 MB
VirtualCenter Agent (vpxa)	5 MB	10	50 MB
vSphere HA agent (Fault Domain Manager, fdm)	5 MB	10	50 MB

You can optionally install VMware vCenter Log Insight, which provides log aggregation and analytics.

VMware Host Client System Requirements

Make sure that your browser supports the VMware Host Client.

The following guest operating systems and Web browser versions are supported for the VMware Host Client.

Supported Browsers	Mac OS	Windows 32-bit and 64-bit	Linux
Google Chrome	89+	89+	75+
Mozilla Firefox	80+	80+	60+
Microsoft Edge	90+	90+	N/A
Safari	9.0+	N/A	N/A

ESXi Passwords and Account Lockout

For ESXi hosts, you must use a password with predefined requirements. You can change the required length and the character class requirement or allow pass phrases using the `Security.PasswordQualityControl` advanced system setting. You can also set the number of passwords to remember for each user using the `Security.PasswordHistory` advanced system setting.

Note: The default requirements for ESXi passwords can change from one release to the next. You can check and change the default password restrictions by using the `Security.PasswordQualityControl` advanced system setting.

ESXi Passwords

ESXi enforces password requirements for access from the Direct Console User Interface, the ESXi Shell, SSH, or the VMware Host Client.

- By default, you must include a mix of at least three from the following four character classes: lowercase letters, uppercase letters, numbers, and special characters such as underscore or dash when you create a password.
- By default, password length is at least 7 characters and less than 40.
- Passwords must not contain a dictionary word or part of a dictionary word.
- Passwords must not contain the user name or parts of the user name.

Note: An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used. A dictionary word used inside a password reduces the overall password strength.

Example ESXi Passwords

The following password candidates illustrate potential passwords if the option is set as follows.

```
retry=3 min=disabled,disabled,disabled,7,7
```


With this setting, a user is prompted up to three times (retry=3) for a new password that is not sufficiently strong or if the password was not entered correctly twice. Passwords with one or two character classes and pass phrases are not allowed, because the first three items are deactivated. Passwords from three- and four-character classes require seven characters. See the `pam_passwdqc` man page for details on other options, such as max, passphrase, and so on.

With these settings, the following passwords are allowed.

- xQaTEhb!: Contains eight characters from three character classes.
- xQaT3#A: Contains seven characters from four character classes.

The following password candidates do not meet requirements.

- Xqat3hi: Begins with an uppercase character, reducing the effective number of character classes to two. The minimum number of required character classes is three.
- xQaTEh2: Ends with a number, reducing the effective number of character classes to two. The minimum number of required character classes is three.

ESXi Pass Phrase

Instead of a password, you can also use a pass phrase. However, pass phrases are deactivated by default. You can change the default setting and other settings by using the `Security.PasswordQualityControl` advanced system setting from the vSphere Client.

For example, you can change the option to the following.

```
retry=3 min=disabled,disabled,16,7,7
```

This example allows pass phrases of at least 16 characters and at least three words.

Changing Default Password Restrictions

You can change the default restriction on passwords or pass phrases by using the `Security.PasswordQualityControl` advanced system setting for your ESXi host. See the *vCenter Server and Host Management* documentation for information on changing ESXi advanced system settings.

You can change the default, for example, to require a minimum of 15 characters and a minimum number of four words (passphrase=4), as follows:

```
retry=3 min=disabled,disabled,15,7,7 passphrase=4
```

See the man page for `pam_passwdqc` for details.

Note: Not all possible combinations of password options have been tested. Perform testing after you change the default password settings.

This example sets the password complexity requirement to require eight characters from four character classes that enforce a significant password difference, a remembered history of five passwords, and a 90 day rotation policy:

```
min=disabled,disabled,disabled,disabled,8 similar=deny
```

ESXi Account Lockout Behavior

Account locking is supported for access through SSH and through the vSphere Web Services SDK. The Direct Console Interface (DCUI) and the ESXi Shell do not support account lockout. By default, a maximum of five failed attempts is allowed before the account is locked. The account is unlocked after 15 minutes by default.

Configuring Login Behavior

You can configure the login behavior for your ESXi host with the following advanced system settings:

- `Security.AccountLockFailures`. Maximum number of failed login attempts before a user's account is locked. Zero deactivates account locking.
- `Security.AccountUnlockTime`. Number of seconds that a user is locked out.
- `Security.PasswordHistory`. Number of passwords to remember for each user. Starting in vSphere 8.0 Update 1, the default is five. Zero deactivates password history.

See the *vCenter Server and Host Management* documentation for information on setting ESXi advanced options.

Before Upgrading ESXi Hosts

For a successful upgrade of your ESXi hosts, understand and prepare for the changes that are involved.

For a successful upgrade, follow these best practices:

1. Make sure that you understand the upgrade process, the effect of that process on your existing deployment, and the preparation required for the upgrade.
 - If your vSphere system includes VMware solutions or plug-ins, make sure they are compatible with the vCenter Server version that you are upgrading to. See the VMware Product Interoperability Matrix at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php.
 - Read [Overview of the ESXi Host Upgrade Process](#) to understand the upgrade scenarios that are supported, and the options and tools that are available to perform the upgrade.
 - Read the VMware vSphere Release Notes for known installation issues.
2. Prepare the system for the upgrade.
 - Make sure that the current version is supported for the upgrade. See [Overview of the ESXi Host Upgrade Process](#).
 - Make sure that the system hardware complies with requirements. See [ESXi Requirements](#) and VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php>. Check for system compatibility, I/O compatibility with network and host bus adapter (HBA) cards, storage compatibility, and backup software compatibility.
 - Make sure that sufficient disk space is available on the host for the upgrade.
 - If a SAN is connected to the host, detach the Fibre Channel system before continuing with the upgrade. Do not deactivate HBA cards in the BIOS.
3. Back up the host before performing an upgrade. If the upgrade fails, you can restore the host.
4. If you are using Auto Deploy to provision hosts, the user who is running the process must have local administrator privileges on the ESXi host that is being provisioned. By default the installation process has these privileges and certificate provisioning happens as expected. However, if you are using another method than the installer, you must run it as a user who has the local administrator privileges.
5. Depending on the upgrade option you choose, you might need to migrate or power off all virtual machines on the host. See the instructions for your upgrade method.
 - For an interactive upgrade from CD, DVD, or USB drive: see [Upgrade Hosts Interactively](#).
 - For a scripted upgrade: see [Installing or Upgrading Hosts by Using a Script](#).
 - For vSphere Auto Deploy: see [Using vSphere Auto Deploy to Reprovision Hosts](#). If the ESXi6.7x or 7.0.x host was deployed by using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with a 8.0 image.
 - For the `esxcli` command method: see [Upgrading Hosts by Using ESXCLI Commands](#).
6. Plan for the tasks that must be performed after the ESXi host upgrade:
 - Test the system to ensure that the upgrade completed successfully.
 - Apply a host's licenses. See [Licensing ESXi Hosts After Upgrade](#).
 - Consider setting up a syslog server for remote logging, to ensure sufficient disk storage for log files. Setting up logging on a remote host is especially important for hosts with limited local storage. You can optionally install VMware vCenter Log Insight, which provides log aggregation and analytics. See [Required Free Space for System Logging](#).

7. If the upgrade was unsuccessful and you backed up the host, you can restore the host.

Upgrading Hosts That Have Third-Party Custom VIBs

A host can have custom vSphere Installation Bundles (VIBs) installed, for example, third-party drivers or management agents.

When you upgrade an ESXi host to 8.0, all supported custom VIBs are migrated, regardless of whether the VIBs are included in the installer ISO. If the host or the installer ISO image contains a VIB that creates a conflict and prevents the upgrade, an error message identifies the VIB that created the conflict. To upgrade the host, take one of the following actions:

- Remove the VIB that created the conflict from the ESXi host and retry the upgrade. You can remove a VIB from the host by using `esxcli` commands. For more information, see [Remove VIBs from a Host](#).
- Use the vSphere ESXi Image Builder CLI to create a custom installer ISO image that resolves the conflict. For more information about vSphere ESXi Image Builder CLI, see the *vCenter Server Installation and Setup* documentation.

Upgrading ESXi Hosts in an Environment With VMware NSX

In a VMware NSX setup, before an upgrade to ESXi 7.0 and later, make sure that the NSX kernel module is part of the image or upgrade baseline.

If your vSphere system includes VMware NSX, before you start an upgrade of your ESXi hosts, you must ensure that the NSX kernel module is part of the desired software specification or baseline that you use for the upgrade. When you upgrade an ESXi host to 7.0 or later, all supported custom VIBs are migrated, regardless of whether the VIBs are included in the installer ISO. However, the NSX kernel module is not automatically migrated to the installer ISO image. Before you proceed to the upgrade operation, see the knowledge base articles [325109](#), [78679](#), and [312013](#).

Media Options for Booting the ESXi Installer

The ESXi installer must be accessible to the system on which you are installing ESXi.

The following boot media are supported for the ESXi installer:

- Boot from a CD/DVD. See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#).
- Boot from a USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).
- Boot from a network. See [Network Booting the ESXi Installer](#).
- Boot from a remote location using a remote management application. See [Using Remote Management Applications](#).

Download and Burn the ESXi Installer ISO Image to a CD or DVD

If you do not have an ESXi installation CD/DVD, you can create one.

You can also create an installer ISO image that includes a custom installation script. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).

1. Follow the procedure [Download the ESXi Installer](#).
2. Burn the ISO image to a CD or DVD.

Format a USB Flash Drive to Boot the ESXi Installation or Upgrade

You can format a USB flash drive to boot the ESXi installation or upgrade.

- Linux machine with superuser access to it

- USB flash drive that can be detected by the Linux machine
- The ESXi ISO image, `VMware-VMvisor-Installer-version_number-build_number.x86_64.iso`, which includes the `isolinux.cfg` file

The instructions in this procedure assume that the USB flash drive is detected as `/dev/sdb`.

Note: The `ks.cfg` file that contains the installation script cannot be located on the same USB flash drive that you are using to boot the installation or upgrade. The kickstart file does not have any dependency on BIOS or UEFI boot.

1. Boot Linux, log in, and enter superuser mode by using a `su` or `sudo root` command.
2. If your USB flash drive is not detected as `/dev/sdb`, or you are not sure how your USB flash drive is detected, determine how it is detected.
 - a) Plug in your USB flash drive.
 - b) At the command line, run the command for displaying the current log messages.

```
tail -f /var/log/messages
```

You see several messages that identify the USB flash drive in a format similar to the following message.

```
Oct 25 13:25:23 ubuntu kernel: [ 712.447080] sd 3:0:0:0: [sdb] Attached SCSI
removable disk
```

In this example, `sdb` identifies the USB device. If your device is identified differently, use that identification in place of `sdb`.

3. Overwrite the entire USB drive with the ISO image. This overwrites the partition table and any previous content on the USB drive.

```
dd bs=10M if=VMware-VMvisor-Installer-version_number-build_number.x86_64.iso of=/dev/
sdb
```

4. Eject the USB drive.

```
eject /dev/sdb
```

You can use the USB flash drive to boot the ESXi installer.

Create a USB Flash Drive to Store the ESXi Installation Script or Upgrade Script

You can use a USB flash drive to store the ESXi installation script or upgrade script that is used during scripted installation or upgrade of ESXi.

- Linux machine
- ESXi installation or upgrade script, the `ks.cfg` kickstart file
- USB flash drive

When multiple USB flash drives are present on the installation machine, the installation software searches for the installation or upgrade script on all attached USB flash drives.

The instructions in this procedure assume that the USB flash drive is detected as `/dev/sdb`.

Note: Do not store the `ks` file containing the installation or upgrade script on the same USB flash drive that you are using to boot the installation or upgrade.

1. Attach the USB flash drive to a Linux machine that has access to the installation or upgrade script.
2. Create a partition table.

```
/sbin/fdisk /dev/sdb
```

- a) Type `d` to delete partitions until they are all deleted.
- b) Type `n` to create primary partition 1 that extends over the entire disk.
- c) Type `t` to set the type to an appropriate setting for the FAT32 file system, such as `c`.
- d) Type `p` to print the partition table.

The result should be similar to the following text:

```
Disk /dev/sdb: 2004 MB, 2004877312 bytes
255 heads, 63 sectors/track, 243 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

   Device Boot      Start         End      Blocks   Id  System
/dev/sdb1            1          243       1951866    c   W95 FAT32 (LBA)
```

- e) Type `w` to write the partition table and quit.

3. Format the USB flash drive with the FAT32 file system.

```
/sbin/mkfs.vfat -F 32 -n USB /dev/sdb1
```

4. Create a destination directory and mount the USB flash drive to it.

```
mkdir -p /usbdisk
mount /dev/sdb1 /usbdisk
```

5. Copy the ESXi installation script to the USB flash drive.

```
cp ks.cfg /usbdisk
```

6. Unmount the USB flash drive.

```
umount /usbdisk
```

The USB flash drive contains the installation or upgrade script for ESXi.

When you boot the ESXi installer, point to the location of the USB flash drive for the installation or upgrade script. See [Enter Boot Options to Run an Installation or Upgrade Script](#) and [PXELINUX Configuration Files](#).

Create an Installer ISO Image with a Custom Installation or Upgrade Script

You can customize the standard ESXi installer ISO image with your own installation or upgrade script. This customization enables you to perform a scripted, unattended installation or upgrade when you boot the resulting installer ISO image.

- Linux machine
- The ESXi ISO image `VMware-VMvisor-Installer-x.x.x-XXXXXX.x86_64.iso`, where `x.x.x` is the version of ESXi you are installing, and `XXXXXX` is the build number of the installer ISO image
- Your custom installation or upgrade script, the `KS_CUST.CFG` kickstart file

See also [Installing ESXi by Using a Script](#) and [About the boot.cfg File](#).

1. Download the ESXi ISO image from the Broadcom Support Portal.

2. Mount the ISO image in a folder:

```
mount -o loop VMware-VMvisor-Installer-x.x.x-XXXXXX.x86_64.iso /esxi_cdrom_mount
```

`XXXXXX` is the ESXi build number for the version that you are installing or upgrading to.

3. Copy the contents of `esxi_cdrom` to another folder:

```
cp -r /esxi_cdrom_mount/* /esxi_cdrom
```

4. Copy the kickstart file to `/esxi_cdrom`.

```
cp KS_CUST.CFG /esxi_cdrom
```

5. Modify the `boot.cfg` file in both `/esxi_cdrom/efi/boot/boot.cfg` (for UEFI boot) and `/esxi_cdrom/boot.cfg` (for legacy BIOS boot) to specify the location of the installation or upgrade script by using the `kernelopt` option.

You must use uppercase characters to provide the path of the script, for example,

```
kernelopt=runweasel ks=cdrom:/KS_CUST.CFG
```

The installation or upgrade becomes completely automatic, without the need to specify the kickstart file during the installation or upgrade.

6. Recreate the ISO image using the `mkisofs` or the `genisoimage` command.

Command	Syntax
mkisofs	<code>mkisofs -relaxed-filenames -J -R -o custom_esxi.iso -b ISOLINUX.BIN -c BOOT.CAT -no-emul-boot -boot-load-size 4 -boot-info-table -eltorito-alt-boot -eltorito-platform efi -b EFIBOOT.IMG -no-emul-boot /esxi_cdrom</code>
genisoimage	<code>genisoimage -relaxed-filenames -J -R -o custom_esxi.iso -b ISOLINUX.BIN -c BOOT.CAT -no-emul-boot -boot-load-size 4 -boot-info-table -eltorito-alt-boot -e EFIBOOT.IMG -no-emul-boot /esxi_cdrom</code>

You can use this ISO installer image for regular boot or UEFI secure boot. However, the vSphere Lifecycle Manager cannot verify the checksum of such an ISO image and you cannot use it for upgrades by using vSphere Lifecycle Manager workflows.

The ISO image includes your custom installation or upgrade script.

Install ESXi from the ISO image.

Download the ESXi Installer

You can obtain the ESXi installer software either from an OEM or from the Broadcom Support Portal.

Register on the Broadcom Support Portal. For more information, see [Register for an account on the Broadcom Support Portal and Communities](#).

For product download instructions, see [Download Broadcom products and software](#).

For download of offline bundle ZIP files for ESXi patches and updates, see [Downloading Broadcom PTF files and solutions](#).

For more information, see [VMware to Broadcom Support Frequently Asked Questions](#).

ESXi Storage Device Names and Identifiers

In the ESXi environment, each storage device is identified by several names.

Device Identifiers

Depending on the type of storage, the ESXi host uses different algorithms and conventions to generate an identifier for each storage device.

Storage-provided identifiers

The ESXi host queries a target storage device for the device name. From the returned metadata, the host extracts or generates a unique identifier for the device. The identifier is based on specific storage standards, is unique and persistent across all hosts, and has one of the following formats:

- `naa.xxx`
- `eui.xxx`
- `t10.xxx`

Path-based identifier

When the device does not provide an identifier, the host generates an `mpx.path` name, where *path* represents the first path to the device, for example, `mpx.vmhba1:C0:T1:L3`. This identifier can be used in the same way as the storage-provided identifiers.

The `mpx.path` identifier is created for local devices on the assumption that their path names are unique. However, this identifier is not unique or persistent, and can change after every system restart.

Typically, the path to the device has the following format:

`vmhbaAdapter:CChannel:TTarget:LLUN`

- `vmhbaAdapter` is the name of the storage adapter. The name refers to the physical adapter on the host, not to the SCSI controller used by the virtual machines.
- `CChannel` is the storage channel number. Software iSCSI adapters and dependent hardware adapters use the channel number to show multiple paths to the same target.
- `TTarget` is the target number. Target numbering is determined by the host and might change when the mappings of targets visible to the host change. Targets that are shared by different hosts might not have the same target number.
- `LLUN` is the LUN number that shows the position of the LUN within the target. The LUN number is provided by the storage system. If a target has only one LUN, the LUN number is always zero (0).

For example, `vmhba1:C0:T3:L1` represents LUN1 on target 3 accessed through the storage adapter `vmhba1` and channel 0.

Legacy identifier

In addition to the device-provided identifiers or `mpx.path` identifiers, ESXi generates an alternative legacy name for each device. The identifier has the following format:

`vmi.number`

The legacy identifier includes a series of digits that are unique to the device. The identifier can be derived in part from the metadata obtained through the SCSI INQUIRY command. For nonlocal devices that do not provide SCSI INQUIRY identifiers, the `vmi.number` identifier is used as the only available unique identifier.

Displaying Device Names in the vSphere CLI

You can use the `esxcli storage core device list` command to display all device names in the vSphere CLI. The output is similar to the following example:

```
# esxcli storage core device list
naa.XXX
    Display Name: DGC Fibre Channel Disk(naa.XXX)
    ...
    Other UUIDs: vml.000XXX
mpx.vmhbal:C0:T0:L0
    Display Name: Local VMware Disk (mpx.vmhbal:C0:T0:L0)
    ...
    Other UUIDs: vml.00000000000XYZ
```

Upgrade Hosts Interactively

To upgrade ESXi6.7 hosts or ESXi7.0 hosts to ESXi8.0, you can boot the ESXi installer from a CD, DVD, or USB flash drive.

- Verify that the ESXi installer ISO is in one of the following locations.
 - On CD or DVD. If you do not have the installation CD or DVD, you can create one. See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#)
 - On a USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#)
- **Note:** You can also use PXE to boot the ESXi installer to run an interactive installation or a scripted installation. See [Overview of the Network Boot Installation Process](#).
- Verify that the server hardware clock is set to UTC. This setting is in the system BIOS or UEFI.
- ESXi Embedded must not be on the host. ESXi Installable and ESXi Embedded cannot exist on the same host.
- If you are upgrading an ESXi host, supported custom VIBs that are not included in the ESXi installer ISO are migrated. See [Upgrading Hosts That Have Third-Party Custom VIBs](#)
- See your hardware vendor documentation for information about changing the boot order.

Before upgrading, consider disconnecting the network storage. This action decreases the time it takes the installer to search for available disk drives. When you disconnect network storage, any files on the disconnected disks are unavailable at installation. Do not disconnect a LUN that contains an existing ESXi installation.

Note: Interactive upgrade is not supported on ESXi hosts with a data processing unit (DPU).

1. Insert the ESXi installer CD or DVD in the CD-ROM or DVD-ROM drive, or attach the Installer USB flash drive and restart the machine.
2. Set the BIOS or UEFI to boot from the CD-ROM device or the USB flash drive.
3. In the Select a Disk panel, select the drive on which to install or upgrade ESXi and press Enter. Press F1 for information about the selected disk.

Note: Do not rely on the disk order in the list to select a disk. The disk order is determined by the BIOS or UEFI. On systems where drives are continuously being added and removed, they might be out of order.

4. Upgrade or install ESXi if the installer finds an existing ESXi installation and VMFS datastore.

If an existing VMFS datastore cannot be preserved, you can choose only to install ESXi and overwrite the existing VMFS datastore, or to cancel the installation. If you choose to overwrite the existing VMFS datastore, back up the datastore first.

5. Press F11 to confirm and start the upgrade.
6. Remove the installation CD or DVD or USB flash drive when the upgrade is complete.
7. Press Enter to reboot the host.
8. Set the first boot device to be the drive which you selected previously when you upgraded ESXi.

Installing or Upgrading Hosts by Using a Script

You can quickly deploy hosts by using scripted, unattended installations or upgrades.

Scripted installations or upgrades provide an efficient way to deploy multiple hosts. The installation or upgrade script contains the installation settings for . You can apply the script to all hosts that you want to have a similar configuration.

For a scripted installation or upgrade, you must use the supported commands to create a script. You can edit the script to change settings that are unique for each host.

The installation or upgrade script can reside in one of the following locations:

- FTP server
- HTTP/HTTPS server
- NFS server
- USB flash drive
- CD-ROM drive

Enter Boot Options to Run an Installation or Upgrade Script

You can start an installation or upgrade script by typing boot options at the ESXi installer boot command line.

At boot time you might need to specify options to access the kickstart file. You can enter boot options by pressing Shift+O in the boot loader. For a PXE boot installation, you can pass options through the `kernelopts` line of the `boot.cfg` file. See [About the boot.cfg File](#) and [Network Booting the ESXi Installer](#).

To specify the location of the installation script, set the `ks=filepath` option, where `filepath` indicates the location of your kickstart file. Otherwise, a scripted installation or upgrade cannot start. If `ks=filepath` is omitted, the text installer is run.

Supported boot options are listed in [Boot Options](#).

1. Start the host.
2. When the ESXi installer window appears, press Shift+O to edit boot options.



3. At the `runweasel` command prompt, type `ks=location of installation script plus boot command-line options`.

Boot Option

You type the following boot options:

```
ks=http://00.00.00.00/kickstart/ks-osdc-pdp101.cfg nameserver=00.00.0.0 ip=00.00.00.000 net-
mask=255.255.255.0 gateway=00.00.00.000
```

Boot Options

When you perform a scripted installation, you might need to specify options at boot time to access the kickstart file.

Supported Boot Options

Table 33: Boot Options for ESXi Installation

Boot Option	Description
<code>BOOTIF=hwtype-MAC address</code>	Similar to the <code>netdevice</code> option, except in the PXELINUX format as described in the IPAPPEND option under SYSLINUX at the syslinux.org site.
<code>gateway=ip address</code>	Sets this network gateway as the default gateway to be used for downloading the installation script and installation media.
<code>ip=ip address</code>	Sets up a static IP address to be used for downloading the installation script and the installation media. Note: the PXELINUX format for this option is also supported. See the IPAPPEND option under SYSLINUX at the syslinux.org site.
<code>ks=cdrom:/path</code>	Performs a scripted installation with the script at <code>path</code> , which resides on the CD in the CD-ROM drive. Each CDROM is mounted and checked until the file that matches the path is found. Important: If you have created an installer ISO image with a custom installation or upgrade script, you must use uppercase characters to provide the path of the script, for example, <code>ks=cdrom:/KS_CUST.CFG</code> .
<code>ks=file://path</code>	Performs a scripted installation with the script at <code>path</code> .
<code>ks=protocol://serverpath</code>	Performs a scripted installation with a script located on the network at the given URL. <code>protocol</code> can be <code>http</code> , <code>https</code> , <code>ftp</code> , or <code>nfs</code> . An example using NFS protocol is <code>ks=nfs://host/porturl-path</code> . The format of an NFS URL is specified in RFC 2224.
<code>ks=usb</code>	Performs a scripted installation, accessing the script from an attached USB drive. Searches for a file named <code>ks.cfg</code> . The file must be located in the root directory of the drive. If multiple USB flash drives are attached, they are searched until the <code>ks.cfg</code> file is found. Only FAT16 and FAT32 file systems are supported.
<code>ks=usb:/path</code>	Performs a scripted installation with the script file at the specified path, which resides on USB.

Boot Option	Description
<code>ksdevice=device</code>	Tries to use a network adapter <i>device</i> when looking for an installation script and installation media. Specify as a MAC address, for example, 00:50:56:C0:00:01. This location can also be a vmnicNN name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.
<code>nameserver=ip address</code>	Specifies a domain name server to be used for downloading the installation script and installation media.
<code>netdevice=device</code>	Tries to use a network adapter <i>device</i> when looking for an installation script and installation media. Specify as a MAC address, for example, 00:50:56:C0:00:01. This location can also be a vmnicNN name. If not specified and files need to be retrieved over the network, the installer defaults to the first discovered network adapter that is plugged in.
<code>netmask=subnet mask</code>	Specifies subnet mask for the network interface that downloads the installation script and the installation media.
<code>vlanid=vlanid</code>	Configure the network card to be on the specified VLAN.
<code>systemMediaSize=small</code>	Limits the size of system storage partitions on the boot media. The selected value must fit the purpose of your system. You can select from the following values: <i>min</i> (32 GB, for single disk or embedded servers) <i>small</i> (64 GB, for servers with at least 512 GB RAM) <i>default</i> (128 GB) <i>max</i> (consume all available space, for multi-terabyte servers)

For more information on ESXi booting options post installation, see VMware knowledge base article [77009](#).

Installing ESXi by Using a Script

The installation/upgrade script is a text file, for example `ks.cfg`, that contains supported commands.

The command section of the script contains the ESXi installation options. This section is required and must appear first in the script.

About the boot.cfg File

The boot loader configuration file `boot.cfg` specifies the kernel, the kernel options, and the boot modules that the `mboot.c32` or `mboot.efi` boot loader uses in an ESXi installation.

The `boot.cfg` file is provided in the ESXi installer. You can modify the `kernelopt` line of the `boot.cfg` file to specify the location of an installation script or to pass other boot options.

The `boot.cfg` file has the following syntax:

```
# boot.cfg -- mboot configuration file
#
# Any line preceded with '#' is a comment.
```

```
title=STRING
prefix=DIRPATH
```

```
kernel=FILEPATH
kernelopt=STRING
modules=FILEPATH1---FILEPATH2... ---FILEPATHn

# Any other line must remain unchanged.
```

The commands in `boot.cfg` configure the boot loader.

Table 34: Commands in boot.cfg.

Command	Description
<code>title=STRING</code>	Sets the boot loader title to <i>STRING</i> .
<code>prefix=STRING</code>	(Optional) Adds <i>DIRPATH/</i> in front of every <i>FILEPATH</i> in the <code>kernel=</code> and <code>modules=</code> commands that do not already start with <code>/</code> or with <code>http://</code> .
<code>kernel=FILEPATH</code>	Sets the kernel path to <i>FILEPATH</i> .
<code>kernelopt=STRING</code>	Appends <i>STRING</i> to the kernel boot options.
<code>modules=FILEPATH1 --- FILEPATH2... --- FILEPATHn</code>	Lists the modules to be loaded, separated by three hyphens (<code>---</code>).

See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#) and [Network Booting the ESXi Installer](#).

Locations Supported for Installation or Upgrade Scripts

In scripted installations and upgrades, the ESXi installer can access the installation or upgrade script, also called the kickstart file, from several locations.

The following locations are supported for the installation or upgrade script:

- CD/DVD. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).
- USB Flash drive. See [Create a USB Flash Drive to Store the ESXi Installation Script or Upgrade Script](#).
- A network location accessible through the following protocols: NFS, HTTP, HTTPS, FTP

Path to the Installation or Upgrade Script

You can specify the path to an installation or upgrade script.

`ks=http://XXX.XXX.XXX.XXX/kickstart/KS.CFG` is the path to the ESXi installation script, where `XXX.XXX.XXX.XXX` is the IP address of the machine where the script resides. See [Installing ESXi by Using a Script](#).

To start an installation script from an interactive installation, you enter the `ks=` option manually. See [Enter Boot Options to Run an Installation or Upgrade Script](#).

Installation and Upgrade Script Commands

To modify the default installation or upgrade script or to create your own script, use supported commands. Use supported commands in the installation script, which you specify with a boot command when you boot the installer.

To determine which disk to install or upgrade ESXi on, the installation script requires one of the following commands: `install`, `upgrade`, or `installorupgrade`. The `install` command creates the default partitions, including a VMFS datastore that occupies all available space after the other partitions are created.

With vSphere 8.0, if your system has supported data processing units (DPU), always consider the installation, re-installation or upgrade of ESXi on the DPUs along with ESXi on hosts. ESXi update and upgrade on DPUs is not supported by the interactive or scripted method, you can only use vSphere Lifecycle Manager.

Note: The use of SD and USB devices for storing ESX-OSData partitions is being deprecated. You can use SD and USB devices only to create boot bank partitions, `boot-bank 0` and `boot-bank 1`. Additionally, you can provide a persistent disk of minimum 32 GB on which to install the ESX-OSData partition. You define such disks by using the parameter `systemDisk` in the `install` command.

accepteula or vmaccepteula (Required)

Accepts the ESXi license agreement.

clearpart (Optional)

Clears any existing partitions on the disk. Requires the `install` command to be specified. Carefully edit the `clearpart` command in your existing scripts.

--drives=

Remove partitions on the specified drives.

--alldrives

Ignores the `--drives=` requirement and allows clearing of partitions on every drive.

--ignoredrives=

Removes partitions on all drives except those specified. Required unless the `--drives=` or `--alldrives` flag is specified.

--overwritevmfs

Allows overwriting of VMFS partitions on the specified drives. By default, overwriting VMFS partitions is not allowed.

--firstdisk=

disk-type1

[*disk-type2*, ...]

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot: : `install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>`. For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (`local`)
2. Network storage (`remote`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESXi installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name `ST3120814A` and any disk that uses the `mptsas` driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `localesx` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

dryrun (Optional)

Parses and checks the installation script. Does not perform the installation.

install

Specifies that this is a fresh installation. Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

--disk= or --drive=

Specifies the disk to partition. In the command `--disk=diskname`, the `diskname` can be a disk name or a full disk filesystem path in ESXi, for example:

- Disk name: `--disk=naa.6d09466044143600247aee55ca2a6405` or
- Device path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`

For accepted disk name formats, see [Disk Device Names](#).

--firstdisk=

`disk-type1,`

`[disk-type2,...]`

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot: `install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>`. For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (`local`)
2. Network storage (`remote`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including esx for the first disk with ESX installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the mptsas driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `localesx` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

--ignoressd

Excludes solid-state disks from eligibility for partitioning. This option can be used with the `install` command and the `--firstdisk` option. This option takes precedence over the `--firstdisk` option. This option is invalid with the `--drive` or `--disk` options and with the `upgrade` and `installorupgrade` commands. See the *vSphere Storage* documentation for more information about preventing SSD formatting during auto-partitioning.

--overwritevsan

You must use the `--overwritevsan` option when you install ESXi on a disk, either SSD or HDD (magnetic), that is in a vSAN disk group. If you use this option and no vSAN partition is on the selected disk, the installation fails. When you install ESXi on a disk that is in vSAN disk group, the result depends on the disk that you select:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group is wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD is wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD is wiped.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

--overwritevmfs

Required to overwrite an existing VMFS datastore on the disk before installation.

--preservevmfs

Preserves an existing VMFS datastore on the disk during installation.

--novmfsondisk

Prevents a VMFS partition from being created on this disk. Must be used with `--overwritevmfs` if a VMFS partition exists on the disk.

--systemdisk

If you use an USB or SD device, `systemDisk` specifies local persistent disk on which to install the ESX-OSData partition. For example, `install --firstdisk = usb --systemDisk=<diskID>`. As a result, boot bank partitions are placed on the USB device, while the OSData partition is on the disk specified in the `systemDisk` parameter.

--repartitionssystemdisk

If you use an USB or SD device and the local disk that you specify with the `systemDisk` parameter is not empty or contains a datastore, you can use `repartitionSystemDisk` to make sure that the persistent disk is repartitioned before use.

Note: If a local persistent disk is not available or the disk size is less than 32GB, you see warning messages, but installation continues.

--forceunsupportedinstall

Blocks the installation of deprecated CPUs.

installorupgrade

Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

--disk= or --drive=

Specifies the disk to partition. In the command `--disk=diskname`, the `diskname` can be a disk name or a full disk filesystem path in ESXi, for example:

- Disk name: `--disk=naa.6d09466044143600247aee55ca2a6405` or
- Device path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`

For accepted disk name formats, see [Disk Device Names](#).

--firstdisk=

disk-type1,

[*disk-type2*,...]

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot: `install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>`. For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (`local`)
2. Network storage (`remote`)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including esx for the first disk with ESX installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the mptsas driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use `localesx` for local storage that contains ESXi image or `remoteesx` for remote storage that contains ESXi image.

--overwritevsan

You must use the `--overwritevsan` option when you install ESXi on a disk, either SSD or HDD (magnetic), that is in a vSAN disk group. If you use this option and no vSAN partition is on the selected disk, the installation fails. When you install ESXi on a disk that is in a vSAN disk group, the result depends on the disk that you select:

- If you select an SSD, the SSD and all underlying HDDs in the same disk group is wiped.
- If you select an HDD, and the disk group size is greater than two, only the selected HDD is wiped.
- If you select an HDD disk, and the disk group size is two or less, the SSD and the selected HDD is wiped.

For more information about managing vSAN disk groups, see the *vSphere Storage* documentation.

--overwritevmfs

Install ESXi if a VMFS partition exists on the disk, but no ESX or ESXi installation exists. Unless this option is present, the installer fails if a VMFS partition exists on the disk, but an ESX or ESXi installation is missing.

keyboard (Optional)

Sets the keyboard type for the system.

keyboardType

Specifies the keyboard map for the selected keyboard type. *keyboardType* must be one of the following types.

- Belgian
- Brazilian
- Croatian
- Czechoslovakian
- Danish
- Estonian
- Finnish
- French
- German
- Greek
- Icelandic
- Italian
- Japanese
- Latin American
- Norwegian
- Polish
- Portuguese
- Russian
- Slovenian
- Spanish
- Swedish
- Swiss French
- Swiss German
- Turkish
- Ukrainian
- United Kingdom
- US Default
- US Dvorak

serialnum or vmserialnum (Optional)

The command is supported in ESXi version 5.1 and later. Configures licensing. If not included, ESXi installs in evaluation mode.

--esx=<license-key>

Specifies the vSphere license key to use. The format is 5 five-character groups (XXXXX-XXXXX-XXXXX-XXXXX-XXXXX).

network (Optional)

Specifies a network address for the system.

--bootproto=[dhcp|static]

Specifies whether to obtain the network settings from DHCP or set them manually.

--device=

Specifies either the MAC address of the network card or the device name, in the form `vmnicNN`, as in `vmnic0`. This option refers to the uplink device for the virtual switch.

--ip=

Sets an IP address for the machine to be installed, in the form `xxx.xxx.xxx.xxx`. Required with the `--bootproto=static` option and ignored otherwise.

--gateway=

Designates the default gateway as an IP address, in the form `xxx.xxx.xxx.xxx`. Used with the `--bootproto=static` option.

--nameserver=

Designates the primary name server as an IP address. Used with the `--bootproto=static` option. Omit this option if you do not intend to use DNS.

The `--nameserver` option can accept two IP addresses. For example: `--nameserver="10.126.87.104[,10.126.87.120]"`

--netmask=

Specifies the subnet mask for the installed system, in the form `255.xxx.xxx.xxx`. Used with the `--bootproto=static` option.

--hostname=

Specifies the host name for the installed system.

--vlanid= *vlanid*

Specifies which VLAN the system is on. Used with either the `--bootproto=dhcp` or `--bootproto=static` option. Set to an integer from 1 to 4096.

--addvmportgroup=(0|1)

Specifies whether to add the VM Network port group, which is used by virtual machines. The default value is 1.

paranoid (Optional)

Causes warning messages to interrupt the installation. If you omit this command, warning messages are logged.

part or partition (Optional)

Creates an extra VMFS datastore on the system. Only one datastore per disk can be created. Cannot be used on the same disk as the `install` command. Only one partition can be specified per disk and it can only be a VMFS partition.

datastore name

Specifies where the partition is to be mounted.

--ondisk= or --ondrive=

Specifies the disk or drive where the partition is created.

--onfirstdisk=

disk-type1,
[*disk-type2*,...]

Note: If your vSphere system is of version earlier than 8.0 Update 3 and has DPUs, you also specify a PCI slot : `install --firstdisk --overwritevmfs --dpupcislots=<PCIeSlotID>` . For systems of version 8.0 Update 3 and later, the `dpupcislots` parameter is deprecated.

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (*local*)
2. Network storage (*remote*)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including `esx` for the first disk with ESX installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the `mptsas` driver rather than a normal local disk, the argument is `--onfirstdisk=ST3120814A,mptsas,local`. You can use `localesx` for local storage that contains ESXi image or `remotesx` for remote storage that contains ESXi image.

reboot (Optional)

Reboots the machine after the scripted installation is complete.

<--noeject>

The CD is not ejected after the installation.

rootpw (Required)

Sets the root password for the system.

--iscrypted

Specifies that the password is encrypted.

password

Specifies the password value.

upgrade

Either the `install`, `upgrade`, or `installorupgrade` command is required to determine which disk to install or upgrade ESXi on.

--disk= or --drive=

Specifies the disk to partition. In the command `--disk=diskname` , the *diskname* can be a disk name or a full disk filesystem path in ESXi, for example:

- Disk name: `--disk=naa.6d09466044143600247aee55ca2a6405` or
- Device path: `--disk=/vmfs/devices/disks/mpx.vmhba1:C0:T0:L0`

For accepted disk name formats, see [Disk Device Names](#).

--firstdisk=
disk-type1,
[*disk-type2*,...]

Partitions the first eligible disk found. By default, the eligible disks are set to the following order:

1. Locally attached storage (*local*)
2. Network storage (*remote*)

You can change the order of the disks by using a comma-separated list appended to the argument. If you provide a filter list, the default settings are overridden. You can combine filters to specify a particular disk, including *esx* for the first disk with ESX installed on it, model and vendor information, or the name of the VMkernel device driver. For example, to prefer a disk with the model name ST3120814A and any disk that uses the *mptsas* driver rather than a normal local disk, the argument is `--firstdisk=ST3120814A,mptsas,local`. You can use *localesx* for local storage that contains ESXi image or *remoteesx* for remote storage that contains ESXi image.

%include or include (Optional)

Specifies another installation script to parse. This command is treated similarly to a multiline command, but takes only one argument.

filename

For example: `%include part.cfg`

%pre (Optional)

Specifies a script to run before the kickstart configuration is evaluated. For example, you can use it to generate files for the kickstart file to include.

--interpreter
=[python|busybox]

Specifies an interpreter to use. The default is *busybox*.

%post (Optional)

Runs the specified script after package installation is complete. If you specify multiple `%post` sections, they run in the order that they appear in the installation script.

--interpreter
=[python|busybox]

Specifies an interpreter to use. The default is *busybox*.

--timeout=secs

Specifies a timeout for running the script. If the script is not finished when the timeout expires, the script is forcefully stopped.

--ignorefailure
=[true|false]

If true, the installation is considered a success even if the `%post` script stops with an error.

%firstboot

Creates an *init* script that runs only during the first boot. The script has no effect on subsequent boots. If multiple `%firstboot` sections are specified, they run in the order that they appear in the kickstart file.

Note: You cannot check the semantics of `%firstboot` scripts until the system is booting for the first time. A `%firstboot` script might contain potentially catastrophic errors that are not exposed until after the installation is complete.

Important: The `%firstboot` script does not run, if secure boot is enabled on the ESXi host.

--interpreter
=[python|busybox]

Specifies an interpreter to use. The default is busybox.

Note: You cannot check the semantics of the `%firstboot` script until the system boots for the first time. If the script contains errors, they are not exposed until after the installation is complete.

Install or Upgrade ESXi from a CD or DVD by Using a Script


You can install or upgrade ESXi from a CD-ROM or DVD-ROM drive by using a script that specifies the installation or upgrade options.

Before you run the scripted installation or upgrade, verify that the following prerequisites are met:

- The system on which you are installing or upgrading meets the hardware requirements. See [ESXi Hardware Requirements](#).
- You have the ESXi installer ISO on an installation CD or DVD . See [Download and Burn the ESXi Installer ISO Image to a CD or DVD](#).
- The default installation or upgrade script (`ks.cfg`) or a custom installation or upgrade script is accessible to the system. See [Installing ESXi by Using a Script](#).
- You have selected a boot command to run the scripted installation or upgrade. See [Enter Boot Options to Run an Installation or Upgrade Script](#). For a complete list of boot commands, see [Boot Options](#).

You can start the installation or upgrade script by entering a boot option when you start the host. You can also create an installer ISO image that includes the installation script. With an installer ISO image, you can perform a scripted, unattended installation when you boot the resulting installer ISO image. See [Create an Installer ISO Image with a Custom Installation or Upgrade Script](#).

1. Boot the ESXi installer from the local CD-ROM or DVD-ROM drive.
2. When the ESXi installer window appears, press Shift+O to edit boot options.



3. Type a boot option that calls the default installation or upgrade script or an installation or upgrade script file that you created.

The boot option has the form `ks=` .

4. Press Enter.

The installation, upgrade, or migration runs, using the options that you specified.

Install or Upgrade ESXi from a USB Stick by Using a Script


You can install or upgrade ESXi from a USB flash drive by using a script that specifies the installation or upgrade options.

Before running the scripted installation or upgrade, verify that the following prerequisites are met:

- The system that you are installing or upgrading to ESXi meets the hardware requirements for the installation or upgrade. See [ESXi Hardware Requirements](#).
- You have the ESXi installer ISO on a bootable USB flash drive. See [Format a USB Flash Drive to Boot the ESXi Installation or Upgrade](#).
- The default installation or upgrade script (`ks.cfg`) or a custom installation or upgrade script is accessible to the system. See [Installing ESXi by Using a Script](#).
- You have selected a boot option to run the scripted installation, upgrade, or migration. See [Enter Boot Options to Run an Installation or Upgrade Script](#).

Supported boot options are listed in [Boot Options](#).

1. Boot the ESXi installer from the USB flash drive.
2. When the ESXi installer window appears, press Shift+O to edit boot options.



3. Type a boot option that calls the default installation or upgrade script or an installation or upgrade script file that you created.

The boot option has the form `ks=`.

4. Press Enter.

The installation, upgrade, or migration runs, using the options that you specified.

Network Boot the ESXi Installer for a Scripted Installation or Upgrade

ESXi8.0 provides many options for booting the installer over a network and using an installation or upgrade script.

- For information about setting up a network infrastructure, see [Network Booting the ESXi Installer](#).
- For information about creating and locating an installation script, see [Installing ESXi by Using a Script](#).
- For specific procedures to network boot the ESXi installer and use an installation script, see one of the following topics:
 - [Boot the ESXi Installer by Using Native UEFI HTTP](#)
 - [Boot the ESXi Installer by Using iPXE and HTTP](#)
 - [Boot the ESXi Installer by Using PXE and TFTP](#)

Disk Device Names

The `install`, `upgrade`, and `installorupgrade` installation script commands require the use of disk device names.

Table 35: Disk Device Names

Format	Example	Description
NAA	naa.6d09466044143600247aee55ca2a6405	SCSI INQUIRY identifier
EUI	eui.3966623838646463	SCSI INQUIRY identifier
T10	t10.SanDisk00Cruzer_Blade000000004C530001171118101244	SCSI INQUIRY identifier
VML	vml.00025261	Legacy VMkernel identifier
MPX	mpx.vmhba0:C0:T0:L0	Path-based identifier

For more information on storage device names, see *Storage Device Names and Identifiers* in the *vSphere Storage* documentation.

How to Boot an ESXi Host from a Network Device

Network Booting the ESXi Installer

You can use preboot execution environment (PXE) to boot an ESXi host from a network device, if your host uses legacy BIOS or UEFI.

Alternatively, if your ESXi host supports native UEFI HTTP, you can use hypertext transfer protocol (HTTP) to boot the host from a network device. ESXi is distributed in an ISO format that is used to install to flash memory or to a local hard drive. You can extract the files and boot them over a network interface.

PXE uses Dynamic Host Configuration Protocol (DHCP) and Trivial File Transfer Protocol (TFTP) to boot an operating system over a network.

PXE booting requires some network infrastructure and a machine with a PXE-capable network adapter. Most machines that can run ESXi have network adapters that can PXE boot.

Native UEFI HTTP uses DHCP and HTTP to boot over a network. UEFI HTTP boot requires a network infrastructure, UEFI firmware version on the ESXi host that includes HTTP boot feature, and a network adapter that supports UEFI networking.


Booting by using HTTP is faster and more reliable than using TFTP. This is due to the capabilities of the TCP protocol that underlies the HTTP, such as built-in streaming and lost packet recovery. If your ESXi hosts do not support native UEFI HTTP, you can use iPXE HTTP for the boot process.

Note: Network booting with legacy BIOS firmware is possible only over IPv4. Network booting with UEFI BIOS firmware is possible over IPv4 or IPv6.

Overview of the Network Boot Installation Process

You can boot an ESXi host from a network interface. The network boot process varies depending on whether the target host is using legacy BIOS or UEFI firmware, and whether the boot process uses PXE TFTP, iPXE HTTP, or UEFI HTTP.

When you boot a target host, it interacts with the different servers in the environment to get a network adapter, boot loader, kernel, IP address for the kernel, and finally an installation script. When all components are in place, installation starts, as shown in the following illustration.

Figure 26: Overview of PXE Boot Installation Process

The interaction between the ESXi host and other servers proceeds as follows:

1. The user boots the target ESXi host.
2. The target ESXi host makes a DHCP request.
3. The DHCP server responds with the IP information, the location of the TFTP or HTTP server, and the filename or URL of the initial network boot loader.
4. The ESXi host contacts the TFTP or HTTP server and requests the filename or URL that the DHCP server specified.
5. The TFTP or HTTP server sends the network boot loader, and the ESXi host runs it. The initial boot loader might load additional boot loader components from the server.
6. The boot loader searches for a configuration file on the TFTP or HTTP server, downloads the kernel and other ESXi components as specified in the configuration file, and boots the kernel on the ESXi host.
7. The installer runs interactively or using a kickstart script, as specified in the configuration file.

Network Boot Background Information

Understanding the network boot process can help you during troubleshooting.

TFTP Server

Trivial File Transfer Protocol (TFTP) is similar to the FTP service, and is typically used only for network booting systems or loading firmware on network devices such as routers. TFTP is available on Linux and Windows.

- Most Linux distributions include a copy of the `tftp-hpa` server. If you require a supported solution, purchase a supported TFTP server from your vendor of choice. You can also acquire a TFTP server from one of the packaged appliances on the VMware Marketplace.
- If your TFTP server runs on a Microsoft Windows host, use `tfptd32` version 2.11 or later. See <http://tftpd32.jounin.net/>.

SYSLINUX and PXELINUX

If you are using PXE in a legacy BIOS environment, you must understand the different boot environments.

- SYSLINUX is an open-source boot environment for machines that run legacy BIOS firmware. The ESXi boot loader for BIOS systems, `mboot.c32`, runs as a SYSLINUX plugin. You can configure SYSLINUX to boot from several types of media, including disk, ISO image, and network. You can find the SYSLINUX package at <http://www.kernel.org/pub/linux/utils/boot/syslinux/>.
- PXELINUX is a SYSLINUX configuration for booting from a TFTP server according to the PXE standard. If you use PXELINUX to boot the ESXi installer, the `pxelinux.0` binary file, `mboot.c32`, the configuration file, the kernel, and other files are transferred by TFTP.

Note: VMware builds the `mboot.c32` plugin to work with SYSLINUX version 3.86 and tests PXE booting only with that version. Other versions might be incompatible. *The Open Source Disclosure Package for VMware vSphere Hypervisor* includes bug fixes for SYSLINUX version 3.86.

iPXE

iPXE is open-source software that provides an implementation of HTTP. You can use the software to perform an initial boot. For more information, see <https://ipxe.org/>.

VMware includes a build of iPXE as part of Auto Deploy. The source tree for this build is available in *The Open Source Disclosure Package for VMware vCenter Server*.

UEFI PXE and UEFI HTTP

Most UEFI firmware natively includes PXE support that allows booting from a TFTP server. The firmware can directly load the ESXi boot loader for UEFI systems, `mboot.efi`. Additional software such as PXELINUX is not required.

Some UEFI firmware support native UEFI HTTP boot. The feature is introduced in version 2.5 of the UEFI specification. The firmware can load the ESXi boot loader from an HTTP server, without additional software, such as iPXE.

Note: Apple Macintosh products do not include PXE boot support. They include support for network booting through an Apple-specific protocol instead.

Alternative Approaches to Network Booting

Alternative approaches to network booting different software on different hosts are also possible, for example:

- Configuring the DHCP server to provide different initial boot loader filenames to different hosts depending on MAC address or other criteria. See your DHCP server's documentation.
- Approaches using iPXE as the initial bootloader with an iPXE configuration file that selects the next bootloader based on the MAC address or other criteria.

PXELINUX Configuration Files

You need a PXELINUX configuration file to boot the ESXi installer on a legacy BIOS system. The configuration file defines the menu displayed to the target ESXi host as it starts.

This section gives general information about PXELINUX configuration files.

For syntax details, see the SYSLINUX website at <http://www.syslinux.org/>.

Required Files

In the PXE configuration file, you must include paths to the following files:

- `mboot.c32` is the boot loader.
- `boot.cfg` is the boot loader configuration file.

See [About the boot.cfg File](#)

Filename for the PXE Configuration File

For the filename of the PXE configuration file, select one of the following options:

- `01-mac_address_of_target_ESXi_host`. For example, `01-23-45-67-89-0a-bc`
- The target ESXi host IP address in a hexadecimal notation.
- `default`

The initial boot file `pxelinux.0` tries to load a PXE configuration file in the following order:

1. It tries with the MAC address of the target ESXi host, prefixed with its ARP type code, which is 01 for Ethernet.
2. If that attempt fails, it tries with the hexadecimal notation of target ESXi system IP address.
3. Ultimately, it tries to load a file named `default`.

File Location for the PXE Configuration File

Save the file in `/tftpboot/pxelinux.cfg/` on the TFTP server.

For example, you might save the file on the TFTP server at `/tftpboot/pxelinux.cfg/01-00-21-5a-ce-40-f6`. The MAC address of the network adapter on the target ESXi host is `00-21-5a-ce-40-f6`.

Boot the ESXi Installer by Using PXE and TFTP

You can use a TFTP server to PXE boot the ESXi installer. The process differs slightly depending on whether you use UEFI or boot from a legacy BIOS.

Because most environments include ESXi hosts that support UEFI boot and hosts that support only legacy BIOS, this topic discusses prerequisites and steps for both types of hosts.

Verify that your environment meets the following prerequisites.

- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with PXE support on the target ESXi host.
- DHCP server that you can configure for PXE booting. See [Sample DHCP Configurations](#).
- TFTP server.
- Network security policies to allow TFTP traffic (UDP port 69).
- For legacy BIOS, you can use only IPv4 networking. For UEFI PXE boot, you can use IPv4 or IPv6 networking.

- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

For legacy BIOS systems, obtain version 3.86 of the SYSLINUX package. For more information, see [Network Boot Background Information](#).

- For legacy BIOS machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `pxelinux.0` initial boot loader for all target machines, but potentially different PXELINUX configuration files depending on the target machine's MAC address.
- For UEFI machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

1. If your ESXi host runs legacy BIOS firmware only, obtain and configure PXELINUX.
 - a) Obtain SYSLINUX version 3.86, unpack it, and copy the `pxelinux.0` file to the top-level `/tftpboot` directory on your TFTP server.
 - b) Create a PXELINUX configuration file using the following code model.

ESXi-8.x.x-XXXXXX is the name of the TFTP subdirectory that contains the ESXi installer files.

```
DEFAULT install
NOHALT 1
LABEL install
  KERNEL ESXi-8.x.x-XXXXXX/mboot.c32
  APPEND -c ESXi-8.x.x-XXXXXX/boot.cfg
  IPAPPEND 2
```

- c) Save the PXELINUX file in the `/tftpboot/pxelinux.cfg` directory on your TFTP server with a filename that will determine whether all hosts boot this installer by default:

Option	Description
Same installer	Name the file <code>default</code> if you want all host to boot this ESXi installer by default.
Different installers	Name the file with the MAC address of the target host machine (01-mac_address_of_target_ESXi_host) if you want only a specific host to boot with this file, for example, 01-23-45-67-89-0a-bc.

2. If your ESXi host runs UEFI firmware, copy the `efi/boot/bootx64.efi` and `efi/boot/crypto64.efi` files from the ESXi installer ISO image to the `/tftpboot` folder on your TFTP server.
3. Rename the `efi/boot/bootx64.efi` file to `mboot.efi`.

Note: Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.
4. Configure the DHCP server.
5. Create a subdirectory of your TFTP server's top-level `/tftpboot` directory and name it after the version of ESXi it will hold, for example, `/tftpboot/ESXi-8.x.x-xxxxx`.
6. Copy the contents of the ESXi installer image to the newly created directory.
7. Modify the `boot.cfg` file
 - a) Add the following line:

```
prefix=ESXi-7.x.x-xxxxxx
```

Here, `ESXi-7.x.x-xxxxxx` is the pathname of the installer files relative to the TFTP server's root directory.

- b) If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (/) character, delete that character.
- c) If the `kernelopt=` line contains the string `cdromBoot`, remove the string only.

8. Optional: For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.

Use the following code as a model, where `XXX.XXX.XXX.XXX` is the IP address of the server where the installation script resides, and `esxi_ksFiles` is the directory that contains the `ks.cfg` file.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```

9. If your ESXi host runs UEFI firmware, specify whether you want all UEFI hosts to boot the same installer.

Option	Description
Same installer	Copy or link the <code>boot.cfg</code> file to <code>/tftpboot/boot.cfg</code>
Different installers	1. Create a subdirectory of <code>/tftpboot</code> named after the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), for example, <code>01-23-45-67-89-0a-bc</code>. 2. Place a copy of (or a link to) the host's <code>boot.cfg</code> file in that directory, for example, <code>/tftpboot/01-23-45-67-89-0a-bc/boot.cfg</code>.

Boot the ESXi Installer by Using iPXE and HTTP

You can use iPXE to boot the ESXi installer from an HTTP server.

The prerequisites and steps depend on the support of UEFI boot or legacy BIOS only. You can use iPXE to boot the ESXi installer from an HTTP server. The following topic discusses prerequisites and steps for ESXi hosts that support UEFI boot and hosts that support legacy BIOS only.

Verify that your environment has the following components:

- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with PXE support on the target ESXi host.
- DHCP server that you can configure for PXE booting. See [Sample DHCP Configurations](#).
- TFTP server.
- Network security policies to allow TFTP traffic (UDP port 69).
- For legacy BIOS, you can use only IPv4 networking. For UEFI PXE boot, you can use IPv4 or IPv6 networking.
- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

Verify that your environment also meets the following prerequisites required for PXE boot using an HTTP Server:

- Verify that the HTTP server is accessible by your target ESXi hosts.
- If your ESXi host runs legacy BIOS firmware only, obtain version 3.86 of the SYSLINUX package. For more information, see [Network Boot Background Information](#).
- For legacy BIOS machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `pxelinux.0` initial boot loader for all target machines, but potentially different PXELINUX configuration files depending on the target machine's MAC address.

- For UEFI machines, the procedure supports booting multiple different versions of the ESXi installer by using the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

1. Obtain and configure iPXE.

- Obtain the iPXE source code.
- On the iPXE download page, follow the build instructions, but run one of the following commands.
 - For ESXi hosts that run legacy BIOS firmware only, run `make bin/undionly.kpxe`.
 - For ESXi hosts that run UEFI firmware, run `make bin-x86_64-efi/snponly.efi`.
- Copy the `undionly.kpxe` or `snponly.efi` file to the `/tftpboot` directory on your TFTP server.

2. If your ESXi host runs legacy BIOS firmware only, obtain and configure PXELINUX.

- Obtain SYSLINUX version 3.86, unpack it, and copy the `pxelinux.0` file to the `/tftpboot` directory on your TFTP server.
- Create a PXELINUX configuration file using the following code model.

`ESXi-8.x.x-XXXXXX` is the name of the TFTP subdirectory that contains the ESXi installer files.

```
DEFAULT install
NOHALT 1
LABEL install
    KERNEL ESXi-8.x.x-XXXXXX/mboot.c32
    APPEND -c ESXi-8.x.x-XXXXXX/boot.cfg
    IPAPPEND 2
```

- Save the PXELINUX file in the `/tftpboot/pxelinux.cfg` directory on your TFTP server.

The filename determines whether all hosts boot this installer by default.

Option	Description
Same installer	Name the file <code>default</code> if you want all host to boot this ESXi installer by default.
Different installers	Name the file with the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), if only a specific host must boot this file. For example, <code>01-23-45-67-89-0a-bc</code> .

- If your ESXi host runs UEFI firmware, copy the `efi/boot/bootx64.efi` file from the ESXi installer ISO image to the `/tftpboot` folder on your TFTP server, and rename the file to `mboot.efi`.

Note: Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.

4. Configure the DHCP server.

- Create a directory on your HTTP server with the same name as the version of ESXi it will hold. For example, `/var/www/html/ESXi-8.x.x-XXXXXX`.

- Copy the contents of the ESXi installer image to the newly created directory.

7. Modify the `boot.cfg` file

- Add the following line:

```
prefix=http://XXX.XXX.XXX.XXX/ESXi-8.x.x-XXXXXX
```


where `http://XXX.XXX.XXX.XXX/ESXi-8.x.x-XXXXXX` is the location of the installer files on the HTTP server.

- b) If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (/) character, delete that character.
- c) If the `kernelopt=` line contains the string `cdromBoot`, remove the string only.

8. Optional: For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.

Use the following code as a model, where `XXX.XXX.XXX.XXX` is the IP address of the server where the installation script resides, and `esxi_ksFiles` is the directory that contains the `ks.cfg` file.

```
kernelopt=ks=http://XXX.XXX.XXX.XXX/esxi_ksFiles/ks.cfg
```

9. If your ESXi host runs UEFI firmware, specify whether you want all UEFI hosts to boot the same installer.

Option	Description
Same installer	Copy or link the <code>boot.cfg</code> file to <code>/tftpbboot/boot.cfg</code>
Different installers	1. Create a subdirectory of <code>/tftpbboot</code> named after the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), for example, <code>01-23-45-67-89-0a-bc</code>. 2. Place a copy of (or a link to) the host's <code>boot.cfg</code> file in that directory, for example, <code>/tftpbboot/01-23-45-67-89-0a-bc/boot.cfg</code>.

Boot the ESXi Installer by Using Native UEFI HTTP

You can boot the ESXi installer directly from an HTTP server, without additional software to support the process.

Verify that your environment has the following components:

- ESXi host with UEFI firmware that supports the HTTP boot feature.
- ESXi installer ISO image, downloaded from the VMware Web site.
- Target host with a hardware configuration that is supported for your version of ESXi. See the *VMware Compatibility Guide*.
- Network adapter with UEFI networking support on the target ESXi host.
- DHCP server that you can configure for UEFI HTTP booting. See [Sample DHCP Configurations](#)
- (Optional) Installation script (kickstart file).
- Use a native VLAN in most cases. If you want to specify the VLAN ID to be used with PXE booting, check that your NIC supports VLAN ID specification.

UEFI HTTP supports booting multiple versions of the ESXi installer. You use the same `mboot.efi` initial boot loader for all target machines, but potentially different `boot.cfg` files depending on the target machine's MAC address.

Note: Do not mix IPv4 or IPv6 networking during the boot process. Use either IPv4 or IPv6 networking.

1. Copy the `efi/boot/bootx64.efi` file from the ESXi installer ISO image to a directory on your HTTP server and rename the file to `mboot.efi`. For example, `http://www.example.com/esxi/mboot.efi`.

Note: Newer versions of `mboot.efi` can generally boot older versions of ESXi, but older versions of `mboot.efi` might be unable to boot newer versions of ESXi. If you plan to configure different hosts to boot different versions of the ESXi installer, use the `mboot.efi` from the newest version.

2. Configure the DHCP server.
3. Create a directory on your HTTP server with the same name as the version of ESXi it will hold. For example, `http://www.example.com/esxi/ESXi-8.x.x-XXXXXX`.
4. Copy the contents of the ESXi installer image to the newly created directory.
5. Modify the `boot.cfg` file.
 - a) Add the following line with the URL of the newly created directory.

```
prefix=http://www.example.com/esxi/ESXi-8.x.x-XXXXXX
```
 - b) If the filenames in the `kernel=` and `modules=` lines begin with a forward slash (/) character, delete that character.
 - c) If the `kernelopt=` line contains the string `cdromBoot`, remove the string only.
6. Optional: For a scripted installation, in the `boot.cfg` file, add the `kernelopt` option to the line after the kernel command, to specify the location of the installation script.
 For example, `kernelopt=ks=http://www.example.com/esxi_ksFiles/ks.cfg`
7. Optional: You can use the virtual machine configuration parameters `networkBootProtocol` and `networkBootUri` to specify from where a virtual machines can boot. The setting `networkBootProtocol` specifies the boot protocol, IPv4 or IPv6. For example, `networkBootProtocol = httpv4`. The setting `networkBootUri` specifies the HTTP URL to the ESXi bootloader (`bootx64.efi`). For example, `networkBootUri = http://xxx.xxx.xx.x/esxi80uc1/efi/boot/bootx64.efi`.
8. Specify whether you want all UEFI hosts to boot the same installer.

Option	Description
Same installer	Add the <code>boot.cfg</code> file to the same directory as <code>mboot.efi</code> . For example, <code>http://www.example.com/esxi/boot.cfg</code>
Different installers	1. Create a subdirectory of the directory that contains the <code>mboot.efi</code> file. Name the directory as the MAC address of the target host machine (<code>01-mac_address_of_target_ESXi_host</code>), for example, <code>01-23-45-67-89-0a-bc</code>. 2. Add the custom <code>boot.cfg</code> file in the directory. For example, <code>http://www.example.com/esxi/01-23-45-67-89-0a-bc/boot.cfg</code>.

You can use both installer types. ESXi hosts without custom `boot.cfg` file on your HTTP server, boot from the default `boot.cfg` file.

Sample DHCP Configurations

The DHCP server must send the address of the TFTP or HTTP server and the filename of the initial boot loader to the ESXi host.

When the target machine first boots, it broadcasts a packet across the network requesting information to boot itself. The DHCP server responds. The DHCP server must be able to determine whether the target machine is allowed to boot and the location of the initial boot loader binary. For PXE boot, the location is a file on a TFTP server. For UEFI HTTP boot, the location is a URL.

CAUTION: Do not set up a second DHCP server if your network already has one. If multiple DHCP servers respond to DHCP requests, machines can obtain incorrect or conflicting IP addresses, or can fail to receive the proper boot information. Talk to a network administrator before setting up a DHCP server. For support on configuring DHCP, contact your DHCP server vendor.

There are many DHCP servers that you can use. The following examples are for an ISC DHCP server. If you are using a version of DHCP for Microsoft Windows, see the DHCP server documentation to determine how to pass the next-server and filename arguments to the target machine.

Example of Booting Using PXE and TFTP with IPv4

This example shows how to configure an ISC DHCP server to PXE boot ESXi using a TFTP server at IPv4 address xxx.xxx.xxx.xxx.

```
#
# ISC DHCP server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server xxx.xxx.xxx.xxx;
    if option client-system-arch = 00:07 or option client-system-arch = 00:09 {
        filename = "mboot.efi";
    } else {
        filename = "pxelinux.0";
    }
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `pxelinux.0` or `mboot.efi` binary file on the TFTP server.

Example of Booting Using PXE and TFTP with IPv6

This example shows how to configure an ISC DHCPv6 server to PXE boot ESXi using a TFTP server at IPv6 address xxxx:xxxx:xxxx:xxxx::xxxx.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option dhcp6.bootfile-url code 59 = string;
option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx::xxxx]/mboot.efi";
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `mboot.efi` binary file on the TFTP server.

Example of Booting Using iPXE and HTTP with IPv4

This example shows how to configure an ISC DHCP server to boot ESXi by loading iPXE from a TFTP server at IPv4 address xxx.xxx.xxx.xxx.

```
#
# ISC DHCP server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
```

```
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "pxeclients" {
    match if substring(option vendor-class-identifier, 0, 9) = "PXEClient";
    next-server xxx.xxx.xxx.xxx;
    if option client-system-arch = 00:07 or option client-system-arch = 00:09 {
        if exists user-class and option user-class = "iPXE" {
            # Instruct iPXE to load mboot.efi as secondary bootloader
            filename = "mboot.efi";
        } else {
            # Load the snponly.efi configuration of iPXE as initial bootloader
            filename = "snponly.efi";
        }
    } else {
        if exists user-class and option user-class = "iPXE" {
            # Instruct iPXE to load pxelinux.0 as secondary bootloader
            filename = "pxelinux.0";
        } else {
            # Load the undionly configuration of iPXE as initial bootloader
            filename = "undionly.kpxe";
        }
    }
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `undionly.kpxe` or `snponly.efi` binary file on the TFTP server. In the legacy BIOS case, iPXE then asks the DHCP server for the next file to load, and the server returns `pxelinux.0` as the filename. In the UEFI case, iPXE then asks the DHCP server for the next file to load, and this time the server returns `mboot.efi` as the filename. In both cases, iPXE is resident and the system has HTTP capability. As a result, the system can load additional files from an HTTP server.

Example of Booting Using iPXE and HTTP with IPv6

This example shows how to configure an ISC DHCPv6 server to boot ESXi by loading iPXE from a TFTP server at IPv6 address `xxxx:xxxx:xxxx:xxxx:xxxx`.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;

option dhcp6.bootfile-url code 59 = string;
if exists user-class and option user-class = "iPXE" {
    # Instruct iPXE to load mboot.efi as secondary bootloader
    option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx:xxxx]/mboot.efi";
} else {
    # Load the snponly.efi configuration of iPXE as initial bootloader
    option dhcp6.bootfile-url "tftp://[xxxx:xxxx:xxxx:xxxx:xxxx]/snponly.efi";
}
```

When a machine attempts to PXE boot, the DHCP server provides an IP address and the location of the `snponly.efi` (iPXE) binary file on the TFTP server. iPXE then asks the DHCP server for the next file to load, and this time the server returns `mboot.efi` as the filename. iPXE is resident and the system has HTTP capability. As a result, the system can load additional files from an HTTP server.

Example of Booting Using UEFI HTTP with IPv4

This example shows how to configure an ISC DHCP server to boot ESXi by using native UEFI HTTP over IPv4 from Web server `www.example.com`.

```
#
# ISC DHCP server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;
option client-system-arch code 93 = unsigned integer 16;
class "httpclients" {
    match if substring(option vendor-class-identifier, 0, 10) = "HTTPClient";
    option vendor-class-identifier "HTTPClient";

    if option client-system-arch = 00:10 {
        # x86_64 UEFI HTTP client
        filename = http://www.example.com/esxi/mboot.efi;
    }
}
```

Example of Booting Using UEFI HTTP with IPv6

This example shows how to configure an ISC DHCPv6 server to boot ESXi by using native UEFI HTTP over IPv6 from Web server `www.example.com`.

```
#
# ISC DHCPv6 server configuration file snippet. This is not a complete
# configuration file; see the ISC server documentation for details on
# how to configure the DHCP server.
#
allow booting;
allow bootp;

option dhcp6.bootfile-url code 59 = string;
option dhcp6.user-class code 15 = { integer 16, string };
option dhcp6.vendor-class code 16 = { integer 32, integer 16, string };

if option dhcp6.client-arch-type = 00:10 {
    # x86_64 HTTP clients
    option dhcp6.vendor-class 0 10 "HTTPClient";
    option dhcp6.bootfile-url "http://www.example.com/esxi/mboot.efi";
}
```

How to Upgrade Hosts by Using ESXCLI Commands

Upgrading Hosts by Using ESXCLI Commands

By using ESXCLI, you can upgrade an ESXi6.7 host or ESXi7.0 host to version 8.0 and update or patch ESXi6.7, ESXi7.0, and ESXi8.0 hosts.

vSphere 8.0 introduces configuration files, components, base images, and add-ons as new software deliverables that you can use to update or patch ESXi8.0 hosts. For information about managing components, base images, and add-ons on ESXi, see [ESXCLI Concepts and Examples](#).

To use ESXCLI commands, you must install Standalone ESXCLI. For more information about installing and using ESXCLI, see the following documents.

- [Using ESXCLI](#)
- [Getting Started with ESXCLI](#)
- [ESXCLI Reference](#)

Note:

If you press Ctrl+C while an `esxcli` command is running, the command-line interface exits to a new prompt without displaying a message. However, the command continues to run to completion.

For ESXi hosts deployed with vSphere Auto Deploy, the tools VIB must be part of the base booting image used for the initial Auto Deploy installation. The tools VIB cannot be added separately later.

VIBs, Image Profiles, and Software Depots

Upgrading ESXi with `esxcli` commands requires an understanding of VIBs, image profiles, and software depots.

The following technical terms are used throughout the vSphere documentation set in discussions of installation and upgrade tasks.

VIB	A VIB is an ESXi software package. VMware and its partners package solutions, drivers, CIM providers, and applications that extend the ESXi platform as VIBs. VIBs are available in software depots. You can use VIBs to create and customize ISO images or to upgrade ESXi hosts by installing VIBs asynchronously onto the hosts.
Image Profile	An image profile defines an ESXi image and consists of VIBs. An image profile always includes a base VIB, and might include more VIBs. You examine and define an image profile by using vSphere ESXi Image Builder.
Software Depot	A software depot is a collection of VIBs and image profiles. The software depot is a hierarchy of files and folders and can be available through an HTTP URL (online depot) or a ZIP file (offline depot). VMware and VMware partners make depots available. Companies with large VMware installations might create internal depots to provision ESXi hosts with vSphere Auto Deploy, or to export an ISO for ESXi installation.

Understanding Acceptance Levels for VIBs and Hosts

Each VIB is released with an acceptance level that cannot be changed. The host acceptance level determines which VIBs can be installed to a host.

The acceptance level applies to individual VIBs installed by using the `esxcli software vib install` and `esxcli software vib update` commands, to VIBs installed using vSphere Lifecycle Manager, and to VIBs in image profiles.

The acceptance level of all VIBs on a host must be at least as high as the host acceptance level. For example, if the host acceptance level is VMwareAccepted, you can install VIBs with acceptance levels of VMwareCertified and VMwareAccepted, but you cannot install VIBs with acceptance levels of PartnerSupported or CommunitySupported. To install a VIB with a restrictive acceptance level that is less than the acceptance level of the host, you can change the setting of the host by using the vSphere Client or by running `esxcli software acceptance` commands.

Setting host acceptance levels is a best practice that allows you to specify which VIBs can be installed on a host and used with an image profile, and the level of support you can expect for a VIB. For example, you might set a more restrictive acceptance level for hosts in a production environment than for hosts in a testing environment.

VMware supports the following acceptance levels.

VMwareCertified

The VMwareCertified acceptance level has the most stringent requirements. VIBs with this level go through thorough testing fully equivalent to VMware in-house Quality Assurance testing for the same technology. Today, only I/O Vendor Program (IOVP) program drivers are published at this level. VMware takes support calls for VIBs with this acceptance level.

VMwareAccepted

VIBs with this acceptance level go through verification testing, but the tests do not fully test every function of the software. The partner runs the tests and VMware verifies the result. Today, CIM providers and PSA plug-ins are among the VIBs published at this level. VMware directs customers with support calls for VIBs with this acceptance level to contact the partner's support organization.

PartnerSupported

VIBs with the PartnerSupported acceptance level are published by a partner that VMware trusts. The partner performs all testing. VMware does not verify the results. This level is used for a new or nonmainstream technology that partners want to enable for VMware systems. Today, driver VIB technologies such as Infiniband, ATAoE, and SSD are at this level with nonstandard hardware drivers. VMware directs customers with support calls for VIBs with this acceptance level to contact the partner's support organization.

CommunitySupported

The CommunitySupported acceptance level is for VIBs created by individuals or companies outside of VMware partner programs. VIBs at this level have not gone through any VMware-approved testing program and are not supported by VMware Technical Support or by a VMware partner.

Table 36: VIB Acceptance Levels Required to Install on Hosts

Host Acceptance Level	VMwareCertified VIB	VMwareAccepted VIB	PartnerSupported VIB	CommunitySupported VIB
VMwareCertified	x			
VMwareAccepted	x	x		
PartnerSupported	x	x	x	
CommunitySupported	x	x	x	x

Match a Host Acceptance Level with an Update Acceptance Level

You can change the host acceptance level to match the acceptance level for a VIB or image profile that you want to install. The acceptance level of all VIBs on a host must be at least as high as the host acceptance level.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Use this procedure to determine the acceptance levels of the host and the VIB or image profile to install, and to change the acceptance level of the host, if necessary for the update.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

1. Retrieve the acceptance level for the VIB or image profile.

Option	Description
List information for all VIBs	<code>esxcli --server=<server_name> software sources vib list --depot=<depot_URL></code>
List information for a specified VIB	<code>esxcli --server=<server_name> software sources vib list --viburl=<vib_URL></code>
List information for all image profiles	<code>esxcli --server=<server_name> software sources profile list --depot=<depot_URL></code>
List information for a specified image profile	<code>esxcli --server=<server_name> software sources profile get --depot=<depot_URL> --profile=<profile_name></code>

2. Retrieve the host acceptance level.

```
esxcli --server=<server_name> software acceptance get
```

3. Optional: If the acceptance level of the VIB is more restrictive than the acceptance level of the host, change the acceptance level of the host.

```
esxcli --server=<server_name> software acceptance set --level=<acceptance_level>
```

The `acceptance_level` can be `VMwareCertified`, `VMwareAccepted`, `PartnerSupported`, or `CommunitySupported`. The values for `acceptance_level` are case-sensitive.

Note: You can use the `--force` option for the `esxcli software vib` or `esxcli software profile` command to add a VIB or image profile with a lower acceptance level than the host. A warning will appear. Because your setup is no longer consistent, the warning is repeated when you install VIBs, remove VIBs, and perform certain other operations on the host.

See if ESXi Host Update Requires Maintenance Mode or Reboot

VIBs that you can install with a live install do not require the host to be rebooted, but might require the host to be placed in maintenance mode.

Other VIBs and profiles might require the host to be rebooted after the installation or update.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

1. Check whether the VIB or image profile that you want to install requires the host to be placed in maintenance mode or to be rebooted after the installation or update.

Run one of the following commands.

Option	Description
Check the VIB	<code>esxcli --server=<server_name> software sources vib get -v <absolute_path_to_vib></code>

Option	Description
Check the VIBs in a depot	<code>esxcli --server=<server_name> software sources vib get --depot=<depot_name></code>
Check the image profile in a depot	<code>esxcli --server=<server_name> software sources profile get --depot=<depot_name></code>

2. Review the return values.

The return values, which are read from the VIB metadata, indicate whether the host must be in maintenance mode before installing the VIB or image profile, and whether installing the VIB or profile requires the host to be rebooted.

Note: vSphere Lifecycle Manager relies on an internal ESXi software scan API to determine whether maintenance mode is required or not. When you install a VIB on a live system, if the value for `Live-Install-Allowed` is set to false, the installation result instructs vSphere Lifecycle Manager to reboot the host. When you remove a VIB from a live system, if the value for `Live-Remove-Allowed` is set to false, the removal result instructs vSphere Lifecycle Manager to reboot the host. In either case, when the remediation starts, vSphere Lifecycle Manager automatically puts the host into maintenance mode.

If necessary, place the host in maintenance mode. See [Place a Host in Maintenance Mode](#). If a reboot is required, and if the host belongs to a VMware HA cluster, remove the host from the cluster or disable HA on the cluster before the installation or update. Also, place the host in maintenance mode to minimize boot disk activity during the upgrade.

Place a Host in Maintenance Mode

Some installation and update operations that use live install require the host to be in maintenance mode.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Maintenance mode is required when an update operation requires a reboot. However, you only put the host in maintenance mode manually when you use `esxcli` commands for update and upgrade operations.

To determine whether an upgrade operation requires the host to be in maintenance mode, see [See if ESXi Host Update Requires Maintenance Mode or Reboot](#)

Note: If the host is a member of a vSAN cluster, and any virtual machine object on the host uses the "Number of failures to tolerate=0" setting in its storage policy, the host might experience unusual delays when entering maintenance mode. The delay occurs because vSAN has to evacuate this object from the host for the maintenance operation to complete successfully.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

1. Check whether the host is in maintenance mode.

```
esxcli --server=<server_name> system maintenanceMode get
```

2. Power off each virtual machine running on the ESXi host.

Note: You can list all running virtual machines and retrieve the World ID of each one by running the following command.

```
esxcli --server=<server_name> vm process list
```

Option	Command
To shut down the guest operating system and then power off the virtual machine	<code>esxcli --server=<server_name> vm process kill --type soft --world-id <vm_ID></code>
To power off the virtual machine immediately	<code>esxcli --server=<server_name> vm process kill --type hard --world-id <vm_ID></code>

Option	Command
To force the power off operation	<code>esxcli --server=<server_name> vm process kill --type force --world-id <vm_ID></code>

Alternatively, to avoid powering off virtual machines, you can migrate them to another host. See the topic [Migrating Virtual Machines](#) in the *vCenter Server and Host Management* documentation.

- Place the host in maintenance mode.

```
esxcli --server=<server_name> system maintenanceMode set --enable true
```

- Verify that the host is in maintenance mode.

```
esxcli --server=<server_name> system maintenanceMode get
```

Update a Host with Individual VIBs

You can update a host with VIBs stored in a software depot that is accessible through a URL or in an offline ZIP depot.

- Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.
- Determine whether the update requires the host to be in maintenance mode or to be rebooted. If necessary, place the host in maintenance mode.
See [See if ESXi Host Update Requires Maintenance Mode or Reboot](#). See [Place a Host in Maintenance Mode](#).
- If the update requires a reboot, and if the host belongs to a vSphere HA cluster, remove the host from the cluster or deactivate HA on the cluster.

Important: If you are updating ESXi from a zip bundle in a VMware-supplied depot, either online from the VMware Web site or downloaded locally, VMware supports only the update method specified for VMware-supplied depots in the topic [Upgrade or Update a Host with Image Profiles](#).

Note: The `esxcli software vib update` and `esxcli software vib install` commands are not supported for upgrade operations. See [Upgrade or Update a Host with Image Profiles](#).

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

- Determine which VIBs are installed on the host.

```
esxcli --server=<server_name> software vib list
```

- Find out which VIBs are available in the depot.

Option	Description
From a depot accessible by URL	<code>esxcli --server=<server_name> software sources vib list --depot=http://<web_server>/<depot_name></code>
From a local depot ZIP file	<code>esxcli --server=<server_name> software sources vib list --depot=<absolute_path_to_depot_zip_file></code>

You can specify a proxy server by using the `--proxy` option.

- Update the existing VIBs to include the VIBs in the depot or install new VIBs.

Option	Description
Update VIBs from a depot accessible by URL	<code>esxcli --server=<server_name> software vib update --depot=http://<web_server>/<depot_name></code>
Update VIBs from a local depot ZIP file	<code>esxcli --server=<server_name> software vib update --depot=<absolute_path_to_depot_ZIP_file></code>

Option	Description
Install all VIBs from a ZIP file on a specified offline depot (includes both VMware VIBs and partner-supplied VIBs)	<code>esxcli --server=<server_name> software vib install --depot <path_to_VMware_vib_ZIP_file> \<VMware_vib_ZIP_file> --depot <path_to_partner_vib_ZIP_file>\<partner_vib_ZIP_file></code>

Options for the `update` and `install` commands allow you to perform a dry run, to specify a specific VIB, to bypass the acceptance level verification, and so on. Do not bypass verification on production systems. See the *ESXCLI Reference*.

- Verify that the VIBs are installed on your ESXi host.

```
esxcli --server=<server_name> software vib list
```

Upgrade or Update a Host with Image Profiles

You can upgrade or update a host with image profiles stored in a software depot that is accessible through a URL or in an offline ZIP depot.

- Install Standalone ESXCLI. See *Getting Started with ESXCLI*. For troubleshooting, run `esxcli` commands in the ESXi Shell.
- Determine whether the update requires the host to be in maintenance mode or to be rebooted. If necessary, place the host in maintenance mode.

See [See if ESXi Host Update Requires Maintenance Mode or Reboot](#). See [Place a Host in Maintenance Mode](#).

Important: When you use ESXCLI to update or upgrade the host, put the host manually in maintenance mode to ensure that the boot disk is not actively in use before the upgrade begins.

- If the update requires a reboot, and if the host belongs to a vSphere HA cluster, remove the host from the cluster or deactivate HA on the cluster.

You can use the `esxcli software profile update` or `esxcli software profile install` command to upgrade or update an ESXi host.

When you upgrade or update a host, the `esxcli software profile update` or `esxcli software profile install` command applies a later version (major or minor) of a full image profile onto the host. After this operation and a reboot, the host can join a vCenter Server environment of the same or later version.

The `esxcli software profile update` command brings the entire content of the ESXi host image to the same level as the corresponding upgrade method using an ISO installer. However, the ISO installer performs a pre-upgrade check for potential problems, such as insufficient memory or unsupported devices. The `esxcli` upgrade method only performs such checks when upgrading from ESXi 6.7 Update 1 or later to a newer version.

Note: Do not use the `--dry-run` option for upgrades from ESXi 6.7.x and ESXi 7.0.x versions earlier than 7.0 Update 3i, to ESXi 8.0 and later. When the `--dry-run` option is removed, you can still use the `esxcli` upgrade method to upgrade from ESXi 6.7 Update 1 or later to ESXi 8.0 or later. For ESXi versions earlier than 6.7 Update 1, you must first upgrade to 6.7 Update 1 or later before you upgrade to ESXi 8.0 or later.

For more about the ESXi upgrade process and methods, see [Overview of the ESXi Host Upgrade Process](#).

Important:

If you are upgrading or updating ESXi from a VMware online depot or an offline .zip bundle downloaded locally from the Broadcom Support Portal, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. ESXCLI commands support 2 types of depot location: either a full path to a locally downloaded offline .zip bundle, such as `/vmfs/volumes/<datastore_name>/<depot_name>.zip`, or an `http/https` link to the `index.xml` in an online depot, such as `https://<web_server>/<depot_name>/index.xml`. You can extract an offline .zip bundle and host it as an online depot on a private `http/https` server accessible to an ESXi host.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

Note:

Options to the `update` and `install` commands allow you to perform a dry run, to bypass acceptance level verification, to ignore hardware compatibility check warnings, and so on. The option to bypass hardware compatibility check warnings is only available for ESXi 6.7 Update 1 or later. Do not bypass verification on production systems.

For options help, type `esxcli software profile install --help` or `esxcli software profile update --help`. For the complete listing of available command-line options, see the [ESXCLI Reference](#).

1. Determine which VIBs are installed on the host.

```
esxcli --server=<server_name> software vib list
```

2. Determine which image profiles are available in the depot.

```
esxcli --server=<server_name> software sources profile list --depot=http://<depot_location>/<depot_name>
```

You can specify a proxy server by using the `--proxy` option.

3. Update the existing image profile to include the VIBs or install new VIBs.

Important: The `software profile update` command updates existing VIBs with the corresponding VIBs from the specified profile, but does not affect other VIBs installed on the target server. The `software profile install` command installs the VIBs present in the depot image profile, and removes any other VIBs installed on the target server.

Option	Description
Update the image profile from a VMware-supplied ZIP bundle, in a depot, accessible online from the Broadcom Support Portal or downloaded to a local depot	<pre>esxcli software profile update --depot=<depot_location> --profile=<profile_name></pre> Important: This is the only update method that VMware supports for zip bundles supplied by VMware. VMware-supplied ZIP bundle names take the form: <code>VMware-ESXi-<version_number>-<build_number>-depot.zip</code>. The profile name for VMware-supplied zip bundles takes one of the following forms. • <code>ESXi-<version_number>-<build_number>-standard</code> • <code>ESXi-<version_number>-<build_number>-notools</code> (does not include VMware Tools)
Update the image profile from a depot accessible by URL	<pre>esxcli --server=<server_name> software profile update --depot=http://<depot_location>/<depot_name> --profile=<profile_name></pre>
Update the image profile from ZIP file stored locally on the target server	<pre>esxcli --server=<server_name> software profile update --depot=file:///<path_to_profile_ZIP_file>/<profile_ZIP_file> --profile=<profile_name></pre>
Update the image profile from a ZIP file on the target server, copied into a datastore	<pre>esxcli --server=<server_name> software profile update --depot=<datastore_name>/<profile_ZIP_file> --profile=<profile_name></pre>

Option	Description
Update the image profile from a ZIP file copied locally and applied on the target server	<code>esxcli --server=<server_name> software profile update --depot=/<>root_dir>/<path_to_profile_ZIP_file>/<profile_ZIP_file> --profile=<profile_name></code>
Install all new VIBs in a specified profile accessible by URL	<code>esxcli --server=<server_name> software profile install --depot=http://<depot_location>/<depot_name> --profile=<profile_name></code>
Install all new VIBs in a specified profile from a ZIP file stored locally on the target.	<code>esxcli --server=<server_name> software profile install --depot=file:///<path_to_profile_ZIP_file>/<profile_ZIP_file> --profile=<profile_name></code>
Install all new VIBs from a ZIP file on the target server, copied into a datastore	<code>esxcli --server=<server_name> software profile install --depot=<datastore_name>/<profile_ZIP_file> --profile=<profile_name></code>
Install all new VIBs from a ZIP file copied locally and applied on the target server	<code>esxcli --server=<server_name> software profile install --depot=/<>root_dir>/<path_to_profile_ZIP_file>/<profile_ZIP_file> --profile=<profile_name></code>

4. Reboot the ESXi host immediately if in the ESXCLI output you see `Reboot Required: true`.

Note: If you do not reboot the ESXi host and continue, further modifications to the host, such as installation of a VIB, might fail or not take effect due to the pending reboot.

5. Verify that the VIBs are installed on your ESXi host.

```
esxcli --server=<server_name> software vib list
```

Update ESXi Hosts by Using Zip Files

You can update hosts with VIBs or image profiles by downloading a ZIP file of a depot.

- Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.
- Download the ZIP file of a depot bundle from a third-party VMware partner.
- Determine whether the update requires the host to be in maintenance mode or to be rebooted. If necessary, place the host in maintenance mode.
See [See if ESXi Host Update Requires Maintenance Mode or Reboot](#). See [Place a Host in Maintenance Mode](#).
- If the update requires a reboot, and if the host belongs to a vSphere HA cluster, remove the host from the cluster or deactivate HA on the cluster.

VMware partners prepare third-party VIBs to provide management agents or asynchronously released drivers.

Important: If you are upgrading or updating ESXi from a zip bundle in a VMware-supplied depot, either online from the VMware Web site or downloaded locally, the `esxcli software vib update` and `esxcli software vib install` commands are not supported. VMware supports only the update and upgrade methods specified for VMware-supplied depots in the topic [Upgrade or Update a Host with Image Profiles](#). If the depot is not supplied by VMware and does not have an image profile, you can use the `esxcli software vib update` and `esxcli software vib install` commands for individual or multiple VIB updates as described in [Update a Host with Individual VIBs](#) and the procedure in this topic.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

Install the ZIP file.

```
esxcli --server=<server_name> software vib update --depot=/<>path_to_vib_ZIP>/<ZIP_file_name>.zip
```

Remove VIBs from a Host

You can uninstall third-party VIBs or VMware VIBs from your ESXi host, unless the VIB is part of a Component that is required for the ESXi lifecycle or a driver.

- If the removal requires a reboot, and if the host belongs to a vSphere HA cluster, deactivate HA for the host.
- Determine whether the update requires the host to be in maintenance mode or to be rebooted. If necessary, place the host in maintenance mode.

See [See if ESXi Host Update Requires Maintenance Mode or Reboot](#). See [Place a Host in Maintenance Mode](#).

Note: To ensure that the boot disk is not actively in use when you use ESXCLI to update or upgrade the host, put the host manually in maintenance mode.

- Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

VMware partners prepare third-party VIBs to provide management agents or asynchronously released drivers.

Important: Starting with ESXi 8.0 Update 2, removing some VIBs from their Components can cause runtime issues with ESXi.

Table 37: VIBs required for ESXi image completeness.

Component	VIBs
ESXi	bmcal
	bmcal-esxio
	clusterstore
	cpu-microcode
	crx
	drivervm-gpu
	esx-base
	esx-dvfilter-generic-fastpath
	esx-ui
	esx-xserver
	esxio
	esxio-base
	esxio-combiner
	esxio-combiner-esxio
	esxio-dvfilter-generic-fastpath
	gc
	gc-esxio
	native-misc-drivers
	infravisor
	native-misc-drivers-esxio
	pensandoatlas
	trx
	vdfs

Component	VIBs
	vsan
	vsanhealth
esx-update	esx-update
	loadesx
esxio-update	esxio-update
	loadesxio
Intel-ne1000	ne1000
Mellanox-nmlx5	nmlx5-core
	nmlx5-rdma
	nmlx5-rdma-esxio
	nmlx5-core-esxio
	nmlx5-cc
	nmlx5-cc-esxio
Mellanox-nmlx5bf-gige	nmlx5bf-gige-esxio
nipmi	nipmi
Pensando-ionic-en	ionic-en-esxio
	ionic-en
VMware-ahci	vmw-ahci
VMware-bfedac	bfedac-esxio
VMware-dwi2c	dwi2c
	dwi2c-esxio
VMware-iser	iser
VMware-mlnx-bfbootctl	mlnx-bfbootctl-esxio
VMware-mnet	mnet-esxio
VMware-nmlx5bf-pmc	nmlx5bf-pmc-esxio
VMware-NVMe-PCIe	nvme-pcie
	nvme-pcie-esxio
VMware-nvme-pcie-plugin	lsuv2-nvme-pcie-plugin
VMware-nvme-plugin	vmware-esx-esxcli-nvme-plugin-esxio
	vmware-esx-esxcli-nvme-plugin
VMware-NVMeoF-RDMA	nvmerdma
VMware-NVMeoF-TCP	nvmetcp
	nvmetcp-esxio
VMware-nvmxnet3	nvmxnet3
	nvmxnet3-esxio
VMware-nvmxnet3-ens	nvmxnet3-ens-esxio
	nvmxnet3-ens

Component	VIBs
VMware-penedac	penedac-esxio
VMware-pengpio	pengpio-esxio
VMware-penspi	penspi-esxio
VMware-pvscsi	pvscsi
	pvscsi-esxio
VMware-rd1173	rd1173-esxio
VMware-SDHCI-Driver	vmksdhci
	vmksdhci-esxio
VMware-spidev	spidev-esxio
VMware-vmkata	vmkata
VMware-vmkusb	vmkusb
	vmkusb-esxio

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

1. Power off each virtual machine running on the ESXi host.

Note: You can list all running virtual machines and retrieve the World ID of each one by running the following command.

```
esxcli --server=<server_name> vm process list
```

Option	Command
To shut down the guest operating system and then power off the virtual machine	<code>esxcli --server=<server_name> vm process kill --type soft --world-id <vm_ID></code>
To power off the virtual machine immediately	<code>esxcli --server=<server_name> vm process kill --type hard --world-id <vm_ID></code>
To force the power off operation	<code>esxcli --server=<server_name> vm process kill --type force --world-id <vm_ID></code>

Alternatively, to avoid powering off virtual machines, you can migrate them to another host. See the topic [Migrating Virtual Machines](#) in the *vCenter Server and Host Management* documentation.

2. Place the host in maintenance mode.

```
esxcli --server=<server_name> system maintenanceMode set --enable true
```

3. If necessary, shut down or migrate virtual machines.

4. Determine which VIBs are installed on the host.

```
esxcli --server=<server_name> software vib list
```

5. Remove the VIB.

```
esxcli --server=<server_name> software vib remove --vibname=<name>
```

Specify one or more VIBs to remove in one of the following forms.

- `<name>`
- `<name>:<version>`
- `<vendor>:<name>`
- `<vendor>:<name>:<version>`

For example, the command to remove a VIB specified by vendor, name and version can take the following form.

```
esxcli --server myEsxiHost software vib remove --vibname=PatchVendor:patch42:version3
```

Note: The `remove` command supports several more options. See the *ESXCLI Reference*.

Add Third-Party Extensions to Hosts with an ESXCLI Command

You can use the `esxcli software vib` command to add a third-party extension released as a VIB package to the system.

When you use this command, the VIB system updates the firewall rule set and refreshes the host daemon after you reboot the system.

Otherwise, you can use a firewall configuration file to specify port rules for host services to enable for the extension. The *vSphere Security* documentation discusses how to add, apply, and refresh a firewall rule set and lists the `esxcli network firewall` commands.

Perform a Dry Run of an ESXCLI Installation or Upgrade

You can use the `--dry-run` option to preview the results of an installation or upgrade operation.

A dry run of the installation or update procedure does not make any changes, but reports the VIB-level operations that will be performed if you run the command without the `--dry-run` option.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

1. Enter the installation or upgrade command, adding the `--dry-run` option.

- `esxcli --server=<server_name> software vib install --dry-run`
- `esxcli --server=<server_name> software vib update --dry-run`
- `esxcli --server=<server_name> software profile install --dry-run`
- `esxcli --server=<server_name> software profile update --dry-run`

2. Review the output that is returned.

The output shows which VIBs will be installed or removed and whether the installation or update requires a reboot.

List the Available VIBs and Profiles After Host Reboot

You can use the `--rebooting-image` option to list the VIBs and profiles that are installed on the host and will be active after the next host reboot.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

1. Enter one of the following commands.

Option	Description
For VIBs	<code>esxcli --server=<server_name> software vib list --rebooting-image</code>
For Profiles	<code>esxcli --server=<server_name> software profile get --rebooting-image</code>

2. Review the output that is returned.

The output displays information for the ESXi image that will become active after the next reboot. If the pending-reboot image has not been created, the output returns nothing.

Display the Image Profile and Acceptance Level of the Host

You can use the `software profile get` command to display the currently installed image profile and acceptance level for the specified host.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

This command also shows details of the installed image profile history, including profile modifications.

When you specify a target server by using `--server=<server_name>`, the server prompts you for a user name and password. Other connection options, such as a configuration file or session file, are supported. For a list of connection options, see [Getting Started with ESXCLI](#), or run `esxcli --help` at the ESXCLI command prompt.

1. Enter the following command.

```
esxcli --server=<server_name> software profile get
```

2. Review the output.

After You Upgrade ESXi Hosts

To complete a host upgrade, you ensure that the host is reconnected to its managing vCenter Server system and reconfigured if necessary. You also check that the host is licensed correctly.

After you upgrade an ESXi host, take the following actions:

- View the upgrade logs. You can use the vSphere Client to export the log files.
- If a vCenter Server system manages the host, you must reconnect the host to vCenter Server by right-clicking the host in the vCenter Server inventory and selecting **Connect**.
- When the upgrade is complete, the ESXi host is in evaluation mode. The evaluation period is 60 days. You must assign a vSphere 8.0 license before the evaluation period expires. You can upgrade existing licenses or acquire new ones from My VMware. Use the vSphere Client to configure the licensing for the hosts in your environment. See the *vCenter Server and Host Management* documentation for details about managing licenses in vSphere.
- The host sdX devices might be renumbered after the upgrade. If necessary, update any scripts that reference sdX devices.
- Upgrade virtual machines on the host. See [Upgrading Virtual Machines and VMware Tools](#).
- Set up the vSphere Authentication Proxy service. Earlier versions of the vSphere Authentication Proxy are not compatible with vSphere 8.0. See the *vSphere Security* documentation for details about configuring the vSphere Authentication Proxy service.

About ESXi Evaluation and Licensed Modes

You can use evaluation mode to explore a set of features equal to the vSphere Enterprise Plus license.

You can use evaluation mode to explore the entire set of features for ESXi hosts. The evaluation mode provides the set of features equal to a vSphere Enterprise Plus license. Before the evaluation mode expires, you must assign to your hosts a license that supports all the features in use. For example, in evaluation mode, you can use vSphere vMotion technology, the vSphere HA feature, the vSphere DRS feature, and other features. If you want to continue using these features, you must assign a license that supports them.

The installable version of ESXi hosts is always installed in evaluation mode. ESXi Embedded is preinstalled on an internal storage device by your hardware vendor. It might be in evaluation mode or prelicensed.

The evaluation period is 60 days and begins when you turn on the ESXi host. At any time during the 60-day evaluation period, you can convert from licensed mode to evaluation mode. The time available in the evaluation period is decreased by the time already used.

For example, suppose that you use an ESXi host in evaluation mode for 20 days and then assign a vSphere Standard Edition license key to the host. If you set the host back in evaluation mode, you can explore the entire set of features for the host for the remaining evaluation period of 40 days.

For ESXi hosts, license or evaluation period expiry leads to disconnection from vCenter Server. All powered on virtual machines continue to work, but you cannot power on virtual machines after they are powered off. You cannot change the current configuration of the features that are in use. You cannot use the features that remained unused before the license expiration.

For information about managing licensing for ESXi hosts, see the *vCenter Server and Host Management* documentation.

Licensing ESXi Hosts After Upgrade

After you upgrade to ESXi8.0, you must apply a vSphere 8 license.

If you upgrade an ESXi host to a version that starts with the same number, you do not need to replace the existing license with a new one. For example, if you upgrade a host from ESXi 6.5 to 6.7, you can use the same license for the host.

If you upgrade an ESXi host to a version that starts with a different number, you must apply a new license. For example, if you upgrade an ESXi host from 7.x to 8.0, you need to license the host with a vSphere 8 license.

When you upgrade ESXi6.7 or ESXi7.0 hosts to ESXi8.0 hosts, the hosts are in a 60-day evaluation mode period until you apply the correct vSphere 8 licenses. See [About ESXi Evaluation and Licensed Modes](#).

You can acquire vSphere 8 licenses from the Broadcom Support Portal. After you have vSphere 8 licenses, you must assign them to all upgraded ESXi8.0 hosts by using the license management functionality in the vSphere Client. See the *vCenter Server and Host Management* documentation for details. If you use the scripted method to upgrade to ESXi8.0, you can provide the license key in the kickstart (ks) file.

For more information, see [Licensing and Subscription in vSphere](#).

Run the Secure Boot Validation Script After ESXi Upgrade

After you upgrade an ESXi host from a version that does not support UEFI secure boot, you must check if you can activate secure boot.

After you upgrade an ESXi host from an older version of ESXi that did not support UEFI secure boot, you might be able to activate secure boot. Whether you can activate secure boot depends on how you performed the upgrade and whether the upgrade replaced all the existing VIBs or left some VIBs unchanged. You can run a validation script after you perform the upgrade to determine whether the upgraded installation supports secure boot.

- Verify that the hardware supports UEFI secure boot.

- Verify that all VIBs are signed with an acceptance level of at least `PartnerSupported`. If you include VIBs at the `CommunitySupported` level, you cannot use secure boot.

For secure boot to succeed, the signature of every installed VIB must be available on the system. Older versions of ESXi do not save the signatures when installing VIBs.

- If you upgrade using ESXCLI commands, the old version of ESXi performs the installation of the new VIBs, so their signatures are not saved and secure boot is not possible.
- If you upgrade using the ISO, new VIBs do have their signatures saved. This is true also for vSphere Lifecycle Manager upgrades that use the ISO.
- If old VIBs remain on the system, the signatures of those VIBs are not available and secure boot is not possible.
 - If the system uses a third-party driver, and the VMware upgrade does not include a new version of the driver VIB, then the old VIB remains on the system after upgrade.
 - In rare cases, VMware might drop ongoing development of a specific VIB without providing a new VIB that replaces or obsoletes it, so the old VIB remains on the system after upgrade.

Note: UEFI secure boot also requires an up-to-date bootloader. This script does not check for an up-to-date bootloader.

1. Upgrade the ESXi and run the following command.

```
/usr/lib/vmware/secureboot/bin/secureBoot.py -c
```

2. Check the output.

The output either includes `Secure boot can be enabled` or `Secure boot CANNOT be enabled`.

Configure Syslog on ESXi Hosts

You can use the vSphere Client, the VMware Host Client, or the `esxcli system syslog` command to configure the syslog service.

The syslog service receives, categorizes, and stores log messages for analyses that help you take preventive action in your environment.

Set ESXi Syslog by using the vSphere Client

You can use the vSphere Client command to configure the syslog service globally and edit various advanced settings.

1. Browse to the ESXi host in the vSphere Client inventory.
2. Click **Configure**.
3. Under **System**, click **Advanced System Settings**.
4. Click **Edit**.
5. Filter for `syslog`.
6. To set up logging globally and configure various advanced settings, see [ESXi Syslog Options](#).
7. Optional: To overwrite the default log size and log rotation for any of the logs:
 - a) Click the name of the log that you want to customize.
 - b) Enter the number of rotations and the log size you want.
8. Click **OK**.

Changes to the syslog options take effect.

Note: Syslog parameter settings that you define by using the vSphere Client or VMware Host Client are effective immediately. However, most settings you define by using ESXCLI require an additional command to take effect. For more details, see [ESXi Syslog Options](#).

Set ESXi Syslog by using the vSphere Client

You can use the vSphere Client to configure and edit syslog service parameters on ESXi hosts.

1. In the vSphere Client, under **Host**, click **Manage > System > Advanced settings**.
2. In the **Search** panel, type a syslog setting you want to define. See [ESXi Syslog Options](#).
3. Select the setting and click **Edit option**.
4. Set the value as described in the table of parameters in [ESXi Syslog Options](#).
5. Click **Save**.

Set ESXi Syslog by using ESXCLI

You can configure the syslog service on ESXi hosts by using the ESXCLI command: `esxcli system syslog config set <syslog option>`.

For information about using the `esxcli system syslog` command and other ESXCLI commands, see [Getting Started with ESXCLI](#). For details how to open the ESXi firewall for the port specified in each remote host specification, see [Configuring the ESXi Firewall](#).

Note:

Using ESXCLI requires ESXi to open SSH logins, which is a security risk, and is not recommended. If you chose to use ESXCLI, make sure you use the `esxcli system syslog reload` command after setting each parameter to make sure that it takes effect.

Use the ESXCLI command `esxcli system syslog config set <syslog option>` to set a syslog option that you decide to enable. For example, to set the `Syslog.global.logHost` option, use the command `esxcli system syslog config set --loghost=<str>`

After setting `Syslog.global.logHost`, ESXi hosts open and maintain connections to the syslog collectors, and the transmission of messages begins immediately. When ESXi generates a syslog message, it writes it to the appropriate log file on the ESXi host and also forwards it to all configured syslog collectors.

ESXi Syslog Options

You can define the behavior of ESXi syslog files and transmissions by using a set of syslog options.

Apart from the base settings, such as `Syslog.global.logHost`, starting from ESXi 7.0 Update 1, a list of advanced options is available for customizations, and NIAP compliance.

Note: Always configure persistent storage before you set any of the audit record parameters or the `Syslog.global.logDir` parameter.

Note: All audit record settings, beginning with `Syslog.global.auditRecord`, take effect immediately. However, for other settings that you define by using ESXCLI, make sure to run the `esxcli system syslog reload` command to enable the changes.

Table 38: Legacy Syslog Options

Option	ESXCLI command	Description
Syslog.global.logHost	esxcli system syslog config set --loghost=<str>	Defines a comma-delimited list of remote hosts and specifications for message transmissions. If the loghost=<str> field is blank, no logs are forwarded. While no hard limit to the number of remote hosts to receive syslog messages exists, good practice is to keep the number of remote hosts to five or less. The format of a remote host specification is: protocol://hostname ipv4 '['ipv6']':port[?formatter=value[&framing=value]]. The protocol must be one of tcp,udp, or ssl. The value of a port can be any decimal number between 1 and 65535. If a port is not provided, UDP and TCP use 514; SSL uses 1514. For example: ssl://hostName1:1514. The formatter must be RFC_3164 or RFC_5424; RFC_3164 is the default. The framing must be non_transparent or octet_counting; the default is non_transparent. For more information, see Fine-tune Syslog on ESXi Hosts .
Syslog.global.defaultRotate	esxcli system syslog config set --default-rotate=<long>	Maximum number of old log files to keep. You can set this number globally and for individual subloggers (see Syslog.global.defaultSize).
Syslog.global.defaultSize	esxcli system syslog config set --default-size=<long>	Default size of log files, in KiB. After a file reaches the default size, the syslog service creates a new file. You can set this number globally and for individual subloggers.
Syslog.global.logDir	esxcli system syslog config set --logdir=<str>	Directory where logs reside. The directory can be on mounted NFS or VMFS volumes. Only the /scratch directory on the local file system is persistent across reboots. Specify the directory as [datastorename] path_to_file, where the path is relative to the root of the volume backing the datastore. For example, the path [storage1] /systemlogs maps to the path /vmfs/volumes/storage1/systemlogs.
Syslog.global.logDirUnique	esxcli system syslog config set --logdir-unique=<bool>	Specifies the ESXi host name to be concatenated to the value of Syslog.global.logDir. It is critical that you enable this setting when multiple ESXi hosts log to a shared file system. Selecting this option

Option	ESXCLI command	Description
		creates a subdirectory with the name of the ESXi host under the directory specified by Syslog.global.LogDir . A unique directory is useful if the same NFS directory is used by multiple ESXi hosts.
Syslog.global.certificate.checkSSL	esxcli system syslog config set --check-ssl-certs=<bool>	Enforces checking of SSL certificates when transmitting messages to remote hosts.

Table 39: Syslog Options Available Starting from ESXi 7.0 Update 1

Option	ESXCLI command	Description
Syslog.global.auditRecord.storageCapacity	esxcli system auditrecords local set --size=<long>	Specifies the capacity of the audit record storage directory located on the ESXi host, in MiB. You cannot decrease the capacity of the audit record storage. You can increase the capacity before or after the audit record storage is enabled (see Syslog.global.auditRecord.storageEnable).
Syslog.global.auditRecord.remoteEnable	esxcli system auditrecords remote enable	Enables sending audit records to remote hosts. Remote hosts are specified by using the Syslog.global.logHost parameter.
Syslog.global.auditRecord.storageDirectory	esxcli system auditrecords local set --directory=<dir>	Creates an audit record storage directory and unless specified, sets /scratch/auditLog as the default location. You must not manually create an audit record storage directory and you cannot change the audit record storage directory while audit record storage is enabled (see Syslog.global.auditRecord.storageEnable).
Syslog.global.auditRecord.storageEnable	esxcli system auditrecords local enable	Enables the storage of audit records on an ESXi host. If the audit record storage directory does not exist, it is created with the capacity specified by Syslog.global.auditRecord.storageCapacity.
Syslog.global.certificate.checkCRL	esxcli system syslog config set --crl-check=<bool>	Enables checking the revocation status of all the certificates in an SSL certificate chain. Enables verification of X.509 CRLs, which are not checked by default in compliance with industry conventions. A NIAP-validated configuration requires CRL checks. Due to implementation limitations, if CRL checks are enabled, then all certificates in a certificate chain must provide a CRL link. Do not enable the <code>crl-check</code> option for installations not related to certification, because of the difficulty in properly configuring an environment that uses CRL checks.

Option	ESXCLI command	Description
<code>Syslog.global.certificate.strictX509Compliance</code>	<code>esxcli system syslog config set --x509-strict=<bool></code>	Enables strict compliance with X.509. Performs additional validity checks on CA root certificates during verification. These checks are generally not performed, as CA roots are inherently trusted, and might cause incompatibilities with existing, misconfigured CA roots. A NIAP-validated configuration requires even CA roots to pass validations. Do not enable the <code>x509-strict</code> option for installations not related to certification, because of the difficulty in properly configuring an environment that uses CRL checks.
<code>Syslog.global.droppedMsgs.fileRotate</code>	<code>esxcli system syslog config set --drop-log-rotate=<long></code>	Specifies the number of old dropped message log files to keep.
<code>Syslog.global.droppedMsgs.fileSize</code>	<code>esxcli system syslog config set --drop-log-size=<long></code>	Specifies the size of each dropped message log file before switching to a new one, in KiB.
<code>Syslog.global.logCheckSSLCerts</code>	<code>esxcli system syslog config set --check-ssl-certs=<bool></code>	Enforces checking of SSL certificates when transmitting messages to remote hosts. Note: Deprecated. Use <code>Syslog.global.certificate.checkSSLCerts</code> in ESXi 7.0 Update 1 and later.
<code>Syslog.global.logFilters</code>	<code>esxcli system syslog config logfilter [add remove set] ...</code>	Specifies one or more log filtering specifications. Each log filter must be separated by a double vertical bar " ". The format of a log filter is: <code>numLogs ident logRegexp.numLogs</code> sets the maximum number of log entries for the specified log messages. After reaching this number, the specified log messages are filtered and ignored. <code>ident</code> specifies one or more system components to apply the filter to the log messages that these components generate. <code>logRegexp</code> specifies a case-sensitive phrase with Python regular expression syntax to filter the log messages by their content.
<code>Syslog.global.logFiltersEnable</code>		Enables the use of log filters.
<code>Syslog.global.logLevel</code>	<code>esxcli system syslog config set --log-level=<str></code>	Specifies the log filtering level. You must change this parameter only when troubleshooting an issue with the syslog daemon. You can use the values <code>debug</code> for the most detailed level, <code>info</code> for the default detail level, <code>warning</code> for only warnings or errors, or <code>error</code> , only for errors.

Option	ESXCLI command	Description
<code>Syslog.global.msgQueueDropMark</code>	<code>esxcli system syslog config --queue-drop-mark=<long>)</code>	Specifies the percent of the message queue capacity at which messages are dropped.
<code>Syslog.global.remoteHost.connectRetryDelay</code>	<code>esxcli system syslog config set --default-timeout=<long></code>	Specifies the delay before retrying to connect to a remote host after a connection attempt fails, in seconds.
<code>Syslog.global.remoteHost.maxMsgLen</code>	<code>esxcli system syslog config set --remote-host-max-msg-len=<long></code>	For the TCP and SSL protocols, this parameter specifies the maximum length of a syslog transmission before truncation occurs, in bytes. The default maximum length for remote host messages is 1 KiB. You can increase the maximum message length to up to 16 KiB. However, raising this value above 1 KiB does not ensure that long transmissions arrive untruncated to a syslog collector. For example, when the syslog infrastructure that issues a message is external to ESXi. This setting does not affect the UDP protocol. RFC 5426 specifies the UDP message lengths that can be safely accepted at 480 bytes for IPV4 and 1180 bytes for IPV6. Because of these restrictions, and because the networking infrastructure can arbitrary drop UDP packets, the use of UDP for transmitting critical syslog messages is not recommended.
<code>Syslog.global.vsanBacking</code>	<code>esxcli system syslog config set --vsan-backing=<bool></code>	Allows log files and the audit record storage directory to be placed on a vSAN cluster. However, enabling this parameter might cause the ESXi host to become unresponsive.

Fine-tune Syslog on ESXi Hosts

By using the right syslog settings, you can achieve proactive monitoring of your environment, reduce downtime and take preventive action on servers.

While setting up syslog, you need to consider several parameters that affect log file retention, syslog transmission, transmission length, error handling, and the set up of SSL certificates for secure syslog message transmission. What follows are recommendations for fine-tuning your syslog parameters. You can see a description of all available parameters at [ESXi Syslog Options](#).

How to Specify Log File Retention

By default, log files cannot expand past a configured size. Once a log file reaches the configured size, logging is routed to a new log file and the oldest log file is deleted.

Note: Best practice is to balance the rotate and size settings. Increasing the rotate setting ensures that syslog files are generated often enough to prevent any potential corruption or destruction from the other log files. Increasing the size setting reduces the time for switching to another log file. Optimal size settings are a multiple of 1024 KiB.

Use the `Syslog.global.defaultSize` setting to specify the log file maximum size in KiB, and `Syslog.global.defaultRotate` to set the maximum number of old log files to keep before rotating to a new log file. To change the log file retention parameters associated with a specific program, use the `Syslog.loggers.<progName>.rotate` and `Syslog.loggers.<progName>.size` settings, where `<progName>` is the name of the program whose parameters you want to adjust.

Manage Settings that Affect the Virtual Machine Log File

You can configure some settings that affect the virtual machine log file, `vmware.log`, either in the `vmx` file or in the `/etc/vmware/config` file. You must power off a virtual machine to edit the `vmx` file and edits take effect only on that virtual machine. If you use the `/etc/vmware/config` file, you must add the prefix `"vmx"` to the setting, for example `vmx.log.keepOld = "20"`, and edits affect all virtual machines on the ESXi host.

Table 40: Configurable settings for the `vmware.log` file

Parameter	Description	Example	Notes
<code>logging</code>	Disables all virtual machine logging.	The default value is <code>logging = "TRUE"</code> To disable virtual machine logging: <code>logging = "FALSE"</code>	Do not use this setting, because disabling virtual machine logging makes it difficult or impossible to get support for virtual machine problems. If you need to use this setting for some reason, you can only place it in the <code>vmx</code> file of a virtual machine.
<code>log.throttleBytesPerSec</code>	Controls when a log file throttles. Log file throttling occurs when writes to the <code>vmware.log</code> exceed the specified rate for a significant amount of time. This occurs when code within the VMX process, which controls a virtual machine, creates excessive log messages. The default value for this setting is 1 KB/sec. In case of log throttling, you see <code><<< Log Throttled >>></code> in the <code>vmware.log</code> file.	<code>log.throttleBytesPerSec = "1500"</code> To disable log throttling, use <code>log.throttleBytesPerSec = "0xFFFFFFFF"</code>	Log file throttling might obscure information necessary to diagnose problems with the affected virtual machine. If you need to disable log throttling, place the line in the example in the <code>vmx</code> file of the affected virtual machine. Remove the line after the debugging session ends.
<code>log.keepOld</code>	Controls the number of older <code>vmware.log</code> file to retain.	<code>log.keepOld = "20"</code>	Do not put the value of this setting below the default value (10). If virtual machines are frequently modified or moved, consider raising this setting to 20 or more.
<code>log.rotateSize</code>	Controls the maximum size of a <code>vmware.log</code> file in bytes.	<code>log.rotateSize = "2500000"</code> To disable limiting the maximum size of a <code>vmware.log</code> file, use <code>log.rotateSize = "0"</code>	A value of this setting below 100,000 can cause a loss of critical log messages and affect virtual machine performance. In ESXi 7.x and earlier, the default

Parameter	Description	Example	Notes
			value of this setting places no limit on the size of a <code>vmware.log</code> file. In ESXi 8.x and later, the default value of this setting is 2,048,000.
<code>log.fileName</code>	Controls the name and location of virtual machine log files.	<pre>log.fileName = "myVMLog"</pre> This setting changes the name of the virtual machine log files from <code>vmware.log</code> to <code>myVMlog</code>. <pre>log.fileName = "/vmfs/volumes/vol1/myVM/myVM.log"</pre> This setting directs virtual machine log files to a directory on a different VMFS volume (vol1) by using <code>myVM</code> for a file name.	Do not place a log file outside the virtual machine directory to make sure that the collection of host support bundles picks up the log file, which can be critical to debug virtual machine problems.
<code>log.fileLevel</code>	Controls the minimum level at which messages are written to <code>vmware.log</code>. Every log message has a level associated with it. Levels below the specified setting are not added to a virtual machine log file. The virtual machine message log levels (from most to least restricted) are: • error • warning • notice • info (default) • trivia • debug • debug1 • debug2 • debug3 • debug4 • debug5 • debug6 • debug7 • debug8 • debug9 • debug10	<pre>log.fileLevel = "debug1"</pre>	Do not set a more restrictive level than "info" to avoid filtering out messages that are necessary for debugging virtual machine problems. Lower the level below "info" only upon request by licensed support. Restore the setting to "info" after debugging ends.
<code>log.filter.minLogLevel.<groupName></code>	Controls the output of specialized debugging messages.	<pre>log.filter.minLogLevel = "debug5"</pre>	Use this setting only upon request by licensed support, who should provide one or more <code><groupName></code> parameters. Remove the setting after debugging ends.

Parameter	Description	Example	Notes
<code>log.syslogID</code>	Activates the sending of virtual machine log messages to the system logger of an ESXi host, such as the syslog.	<code>log.syslogID = "vmx"</code>	Use "vmx" as value for this setting to allow the ESXi syslog daemon, <code>vmssyslogd</code> , to send these messages to a separate log file.
<code>log.syslogLevel</code>	Controls the minimum level at which messages are output to the system logger of an ESXi host, such as the syslog.	<code>log.syslogLevel = "debug"</code>	The levels and functioning of this setting are identical to those for the <code>log.fileLevel</code> setting.

How to Specify Message Transmission to Remote Hosts

Optionally, you can configure ESXi to send syslog messages to one or more remote hosts, called syslog collectors, such as VMware Aria Operations for Logs (formerly VMware vRealize Log Insight and vCenter Log Insight), to collect syslog messages.

Note: Best practice is that you configure each ESXi host to send syslog messages to at least one syslog collector. This helps ensure that the messages are preserved in case of a catastrophic system event and that you can process syslog messages in various ways, such as real-time categorization and analysis (for example, by type, time span, or machine), or archive messages.

Use the `Syslog.global.logHost` setting to define remote host specifications. Separate multiple remote host specifications with a comma (.). After setting `Syslog.global.logHost`, ESXi hosts open and maintain connections to the syslog collectors, and the transmission of messages begins immediately. When ESXi generates a syslog message, it writes it to the appropriate log file on the ESXi host and also forwards it to all configured syslog collectors.

In addition to syslog messages, audit messages can also be transmitted to syslog collectors for security purposes. Audit records track security-related activity on the ESXi host. For more information about audit records, see [Audit Records](#).

Note: Consult with your company security response team if and how to set audit records. Certified configurations usually require audit records to be activated.

What follows is the syntax for `Syslog.global.logHost` remote host specifications:

```
protocol://target[:port][?formatter=value[&framing=value]]
```

Parameter	Description	Notes
<code>protocol</code>	Specifies the networking protocol. Valid values are <code>udp</code> , <code>tcp</code> , and <code>ssl</code> .	The <code>ssl</code> protocol specifies that transmission of syslog messages is encrypted. The <code>tcp</code> and <code>udp</code> protocols do not encrypt the transmission. Note: If capturing syslog messages or audit messages is critical to your system, avoid using the <code>udp</code> protocol because the networking infrastructure external to ESXi might drop UDP messages.
<code>target</code>	Specifies the remote host. You can use either an IPV4 or IPV6 address, or a host name.	When you use an IPV6 address, you must embed it in square brackets <code>[xxx]</code> , where <code>xxx</code> is the IPV6 address.

Parameter	Description	Notes
<code>port</code>	(Optional) Specifies the remote host port to use. If you use UDP or TCP, the default port is 514. If you use SSL, the default port is 1514. If you choose to use different ports from 514 or 1514, you must adjust the ESXi firewall to open the port.	For details how to open the ESXi firewall for the port specified in each remote host specification, see Configuring the ESXi Firewall .
<code>formatter</code>	Specifies how transmissions are formatted. The formatter must be RFC 3164 or RFC 5424.	RFC 3164 is the default.
<code>framing</code>	Specifies if transmissions are framed. Framing must be <code>non_transparent</code> or <code>octet_counting</code> .	The default is <code>non_transparent</code> . Transmissions in RFC 5424 format must specify <code>octet_counting</code> framing. For more information, see Protocols, Formats and Framing of ESXi Syslog Messages .

Examples of remote machine specifications:

Syslog.global.logHost string example	Notes
<code>tcp://10.176.130.7:12345?formatter=RFC_3164</code>	Transmits syslog messages to 10.176.130.7 using TCP/IP and port 12345. Transmission format is RFC 3164 with no framing.
<code>tcp://10.176.130.7:12345?formatter=RFC_3164&framing=octet_counting</code>	Transmits syslog messages to 10.176.130.7 using TCP/IP and port 12345. Transmission format is RFC 3164 and framing is <code>octet_counting</code> .
<code>tcp://10.176.130.7:12345?formatter=RFC_5424&framing=octet_counting</code>	Transmits syslog messages to 10.176.130.7 using TCP/IP and port 12345. Transmission format is RFC 5424 and framing is <code>octet_counting</code> .
<code>tcp://[2001:db8:85a3:8d3:1319:8a2e:370:7348]</code>	Transmits syslog messages to an IPV6 address using port 1514.
<code>tcp://[2001:db8:85a3:8d3:1319:8a2e:370:7348]:5432</code>	Transmit syslog messages to an IPV6 address using port 54321.
<code>udp://company.com</code>	Transmits syslog messages to <code>company.com</code> using UDP and port 514.
<code>udp://company.com,tcp://10.20.30.40:1050</code>	Transmits syslog messages to two remote hosts. The first remote host uses UDP to communicate with <code>company.com</code> using port 514. The second remote host uses TCP to communicate with the IPV4 address <code>10.20.30.40</code> using port 1050.
<code>ssl://company.com</code>	Transmits syslog messages to <code>company.com</code> using SSL (TLS) and port 514.

Maximum Message Transmission Length

If you use UDP, the maximum syslog message transmission length is 480 bytes for IPV4 and 1180 bytes for IPV6.

For TCP or SSL, the default maximum syslog message transmission length is 1 kibibyte (KiB). You can increase this length by using the `Syslog.global.remoteHost.maxMsgLen` parameter. The maximum value is 16 KiB. Messages longer than 16 KiB are truncated.

Note: If increasing the maximum transmission length is necessary, best practice is to increase the length only as much as specifically necessary.

Increasing the maximum syslog message length can cause problems if the networking and syslog infrastructure external to ESXi is unable to handle messages longer than 1 KiB.

Note: Best practice is that you do not use UDP to transmit syslog messages due to the packet length constraints and the possibility that the external networking infrastructure might drop the messages.

Considerations for Certificates When Configuring SSL Transmissions to Remote Hosts

When configuring ESXi to transmit syslog messages to remote hosts using SSL, you must add an SSL certificate for each remote host to the ESXi host CA store. For more information, see [Managing Certificates for ESXi Hosts](#).

Note: Consult with your syslog collector documentation on how to configure the collector for secure receipt of syslog messages using SSL and a private key.

Additional SSL Transmission Parameters

An ESXi system complying with security certification requirements might require activating X509 CRL checks. You turn on the advanced settings `Syslog.global.certificate.strictX509Compliance` and `Syslog.global.certificate.checkCRL` by changing the default value of `false` to `true`. Due to implementation limitations, if you activate CRL checks by using the setting `Syslog.global.certificate.checkSSLCerts`, then all certificates in a certificate chain must provide a CRL link. By default, the setting is active. You can deactivate SSL certificate checks by changing the setting to `false`, but this is not a best practice. You might need to turn off SSL certificate checks when troubleshooting communications with a remote host, but do this only for a limited time.

Where to Find Syslog Daemon Error and Status Information

The ESXi syslog daemon uses the log file `/var/run/log/vmsyslogd.log` to store status and error information, including dropped messages. If audit record transmission is active, the syslog daemon also emits audit records related to its operation, such as daemon start, stop, and error conditions, which allows you to verify that the syslog daemon runs properly.

How to Change the Default Syslog Log File Storage Area

The default syslog log file storage area is `/var/run/log`, local to each ESXi host. Use the `Syslog.global.logDir` syslog configuration variable to change the default syslog log file storage area, as long as the location resides on persistent storage. If `Syslog.global.logDir` is configured to a persistent store shared by multiple ESXi hosts to store their syslog log files, change the `Syslog.global.logDirUnique` setting to `true` to prevent mixing logs. The `Syslog.global.logDirUnique` setting makes sure that each ESXi machine gets a unique name added to the `Syslog.global.logDir` path, separating the log files from other hosts.

Syslog Message Queueing for Remote Hosts and Message Drops

Once syslog emissions start, they never stop except for ESXi reboots and failures, or a syslog reconfiguration to stop.

To avoid dropping messages, ESXi uses an in-memory queue that allows the `vmsyslogd` service to handle the following conditions for a short time:

- ESXi generates log messages at a faster rate than the `vmsyslogd` service can process and transmit
- Network connectivity between ESXi and the remote host fails

If either of these conditions continues for a long period, the capacity of the in-memory queue might not be sufficient and `vmsyslog` might stop transmitting messages and audit logs to remote hosts. No data is lost, because dropped messages are written to the `/var/run/log/vmsyslogd-dropped.log` file.

To minimize the chance that message dropping occurs, place the `vmsyslogd` log files on the fastest available storage and configure `vmsyslogd` and its syslog collectors on a network with end-to-end bandwidth greater than 1 GigE.

Note: Consider an end-to-end bandwidth of 2.5 GigE or more, preferably 10 GigE, to optimize performance and prevent message drops.

If you see excessive logs from an application, file a support service request for analysis and correction.

You can see statistics about message drops in the syslog daemon log file.

You can see dropped messages at `/var/run/log/vmsyslogd-dropped.log`. This log file has retention settings specific to it, similar to those for the program-specific retention parameters. The dropped messages log file retention parameters are: `Syslog.global.droppedMsgs.fileRotate` and `Syslog.global.droppedMsgs.fileSize`.

Configure Log Filtering on ESXi Hosts

The log filtering capability lets you modify the logging policy of the syslog service that is running on an ESXi host.

You can create log filters to reduce the number of repetitive entries in the ESXi logs and to denylist specific log events entirely.

Install ESXCLI. See *Getting Started with ESXCLI*. For troubleshooting, run `esxcli` commands in the ESXi Shell.

Starting with vSphere 7.0 Update 2, you can add log filters and enable log filtering by using ESXCLI. A log filter, once established, remains in place until it is removed, even across ESXi reboots.

Log filters affect all log events that are processed by the ESXi host `vmsyslogd` service, whether they are recorded to a log directory or to a remote syslog server.

You must enable the log filtering capability and reload the syslog daemon to activate the log filters on the ESXi host.

ESXCLI commands to configure log filters follow this pattern: `esxcli system syslog config logfilter {cmd} [cmd options]`.

For example, to get the list of available log filters, run the following command: `[root@xxx-xx-dhcp-xx-xx:~] esxcli system syslog config logfilter list`.

Use the `set` command to activate or deactivate log filtering: `[root@xxx-xx-dhcp-xx-xx:~] esxcli system syslog config logfilter set`.

Use the `add` command to add a log filter and the `remove` command to remove a log filter.

Use the `get` command to determine if log filtering is enabled.

A log filter is specified by three components and uses the following syntax: `numLogs | ident | logRegexp`.

Parameter	Description
<code>numLogs</code>	Specifies the number of matches of the <code>logRegexp</code> Python regular expression that will be allowed before filtering begins.
<code>ident</code>	The <code>ident</code> string is how an application identifies itself to the syslog facility. The <code>logRegexp</code> filter must be associated with the same application. You can find the <code>ident</code> string of an application by inspecting the log files in <code>/var/run/log</code> . The third field of each log file begins with the <code>ident</code> string and ends with <code>[</code> .
<code>logRegexp</code>	Python regular expression that identifies the messages which you want to filter out.

For example, to filter out all messages from the `hostd` daemon that contain the word "mark" after the tenth occurrence, use the following command: `esxcli system syslog config logfilter add --filter="10|Hostd|mark"`.

To remove the log filter, use the command `esxcli system syslog config logfilter remove --filter="10|Hostd|mark"`.

For more information, see [ESXi Syslog Options](#).

Using vSphere Auto Deploy to Reprovision Hosts

You can reprovision ESXi hosts deployed by using vSphere Auto Deploy with a new image profile for a different version of ESXi.

Note: If you upgrade the host to use an ESXi 6.0 or later image, the vSphere Auto Deploy server provisions the ESXi host with certificates that are signed by VMCA. If you are currently using custom certificates, you can set up the host to use the custom certificates after the upgrade. See *vSphere Security*.

If a host was deployed using vSphere Auto Deploy, you can use vSphere Auto Deploy to reprovision the host with a new image profile that contains a different version of ESXi. You can use vSphere ESXi Image Builder to create and manage image profiles. The vSphere Auto Deploy server is automatically upgraded if you upgrade the corresponding vCenter Server system. Starting with version 6.0, the vSphere Auto Deploy server is always on the same management node as the vCenter Server system.

Introduction to vSphere Auto Deploy

vSphere Auto Deploy uses PXE boot infrastructure with host profiles, a desired image, or configuration on a cluster level to provision ESXi hosts.

State Information for ESXi Hosts

Note: You cannot use Auto Deploy on ESXi hosts configured with DPUs as part of the vSphere Distributed Services Engine feature.

When you start a physical host that is set up for vSphere Auto Deploy, vSphere Auto Deploy uses PXE boot infrastructure in conjunction with vSphere host profiles, a desired image, or configuration on a cluster level to provision and customize that host. No state is stored on the host itself. Instead, the vSphere Auto Deploy server manages state information for each host. vSphere Auto Deploy stores the information for the ESXi hosts to be provisioned in different locations. Information about the location of image profiles, host profiles, or clusters that you manage by either a single image or by a configuration on a cluster level is initially specified in the rules that map machines to image profiles and host profiles.

Table 41: vSphere Auto Deploy Stores Information for Deployment


Information Type	Description	Source of Information
Image state	The executable software to run on an ESXi host.	Image profile, created with vSphere ESXi Image Builder or a vSphere Lifecycle Manager image.
Configuration state	The configurable settings that determine how the host is configured, for example, virtual switches and their settings, driver settings, boot parameters, and so on.	Host profile, created by using the host profile UI, or a configuration that you create when setting up a cluster that manages all ESXi host settings at a cluster level in the Inventory UI.
Dynamic state	The runtime state that is generated by the running software, for example, generated private keys or runtime databases.	Host memory, lost during reboot.
Virtual machine state	The virtual machines stored on a host and virtual machine autostart information (subsequent boots only).	Virtual machine information sent by vCenter Server to vSphere Auto Deploy must be available to supply virtual machine information to vSphere Auto Deploy.
User input	State that is based on user input, for example, an IP address that the user provides when the system starts up, cannot automatically be included in the host profile.	Host customization information, stored by vCenter Server during first boot. You can create a host profile that requires user input for certain values.

Information Type	Description	Source of Information
		When vSphere Auto Deploy applies a host profile that requires user provided information, the host is placed in maintenance mode. Use the host profile UI to check the host profile compliance, and respond to the prompt to customize the host.

vSphere Auto Deploy Architecture

The vSphere Auto Deploy infrastructure consists of several components.

Figure 27: vSphere Auto Deploy Architecture



vSphere Auto Deploy server

vSphere Auto Deploy rules engine

Image profiles

Serves images and host profiles to ESXi hosts.

Apart from legacy image profiles that you create by using the VMware Image Builder and host profiles, you can also create vSphere Auto Deploy rules to deploy ESXi by using a single vSphere Lifecycle Manager image or a configuration on a cluster level.

Define the set of VIBs to boot ESXi hosts with.

- VMware and VMware partners make image profiles and VIBs available in public depots. Use vSphere ESXi Image Builder to examine the depot and use the vSphere Auto Deploy rules engine to specify which image profile to assign to which host.
- You use vSphere Lifecycle Manager images to apply software and firmware updates to the ESXi hosts in a cluster. Using a

single image to manage all hosts in a cluster ensures cluster-wide host image homogeneity.

- With ESXi 8.0, you can set up a cluster that manages all ESXi host settings at a cluster level.

Host profiles

Host customization

Stores information that the user provides when host profiles are applied to the host. Host customization might contain an IP address or other information that the user supplied for that host. For more information about host customizations, see the *vSphere Host Profiles* documentation.

Host customization was called answer file in earlier releases of vSphere Auto Deploy.

Auto Deploy Certificates

By default, the Auto Deploy server provisions each host with certificates that are signed by the VMware Certificate Authority (VMware CA). For more information, see [Managing Certificates for ESXi Hosts](#).

Alternatively, if your corporate policy requires that you use custom certificates, you can set up the Auto Deploy server to provision all hosts with custom certificates that are not signed by VMware CA. The Auto Deploy server becomes a subordinate certificate authority of your third-party CA. In the Custom Certificate Authority mode, you are responsible for managing the certificates. You cannot refresh and renew certificates from the vSphere Client. In this mode, you also cannot select only a set of hosts to provision with custom certificates, and you can manually sign custom certificates only for stateful hosts. For more information, see [Use Custom Certificates with Auto Deploy](#).

With ESXi8.0, Auto Deploy provides a third option that allows you to generate a certificate outside vSphere and become independent of the certificate management in vCenter Server. For example, you can generate a custom certificate by using a custom script or by using a provider of domain name registry services such as Verisign. You can use custom certificates for only a set of ESXi hosts. You can provide custom certificates for stateless hosts as well. ESXi hosts are identified by the MAC address of the NIC used for network booting, or the BIOS UUID of the ESXi host. You update the VMware Endpoint Certificate Store (VECS) with the custom certificate by using PowerCLI. For more information on the new PowerCLI cmdlets, see [vSphere Auto Deploy PowerCLI Cmdlet Overview](#). The VMware CA must trust the custom ESXi certificates so you must add the CA public certificate for the custom certificates to the TRUSTED_ROOTS store in VECS. Auto Deploy also stores the custom certificates and when it recognizes a booting host with the respective MAC address of the NIC used for network booting, or the BIOS UUID of the ESXi host, it automatically provides the custom certificate. You do not need to stop or restart Auto Deploy or vCenter Server when you add a custom certificate to VECS, only restart the host for which you upload a custom certificate. For more information, see [Use Custom Certificates with Auto Deploy](#).

Install and Configure vSphere Auto Deploy

Before you can start using vSphere Auto Deploy, you must prepare your environment in several steps.

You start with server setup and hardware preparation. You must configure the vSphere Auto Deploy service startup type in the vCenter Server system that you plan to use for managing the hosts you provision, and install PowerCLI.

vSphere Auto Deploy Preinstallation Checklist

Before you can start the tasks in this vSphere Auto Deploy scenario, make sure that your environment meets the hardware and software requirements, and that you have the necessary permissions for the components included in the setup.

Table 42: Preinstallation Checklist

Required Software and Hardware	Details
vCenter Server	The vSphere Auto Deploy server is part of vCenter Server. You must enable and start the vSphere Auto Deploy service on the vCenter Server system. You can perform many of the setup tasks by logging in to vCenter Server. See Prepare Your System for vSphere Auto Deploy .
Storage	Storage for ESXi datastores NFS, iSCSI, or Fibre Channel, with servers and storage arrays that are configured so the servers can detect the LUNs. - A list of target IP addresses for NFS or iSCSI. - A list of target volume information for NFS or iSCSI.
Host information (for four ESXi hosts)	A list of target IP addresses for NFS or iSCSI. A list of target volume information for NFS or iSCSI. - Default route, net mask, and primary and secondary DNS server IP addresses. - IP address and net mask for the VMkernel primary management network. - IP address and net mask for other VMkernel networks such as storage, vSphere FT, or VMware vMotion. vSphere Auto Deploy does not overwrite existing partitions by default.
PowerCLI	See Install PowerCLI .
ESXi software depot	The location of the ESXi software depot on the Downloads page of the VMware website. You use a URL to point to the image profile stored at that location, or you download a ZIP file to work with a local depot. Do not download the ESXi image.
TFTP server	TFTP installer software such as WinAgents TFTP server.
DHCP server	The DHCP server is included in the vSphere supported Windows Server versions.
DNS server	A working DNS server. You must add entries in both Forward (A Record) and Reverse (PTR Record) Zones for each target host.

You also need information about and administrator privileges to the core servers of the environment, including the ActiveDirectory server, DNS server, DHCP server, NTP server, and so on.

You must have complete control of the broadcast domain of the subnet in which you deploy the setup. Ensure that no other DHCP, DNS, or TFTP server are on this subnet.

Prepare Your System for vSphere Auto Deploy

Before you can PXE boot an ESXi host with vSphere Auto Deploy, you must install prerequisite software and set up the DHCP and TFTP servers that vSphere Auto Deploy interacts with.

- Verify that the hosts that you plan to provision with vSphere Auto Deploy meet the hardware requirements for ESXi. See [ESXi Hardware Requirements](#).

- Verify that the ESXi hosts have network connectivity to vCenter Server and that all port requirements are met. See *vCenter Server Upgrade*.
- Verify that you have a TFTP server and a DHCP server in your environment to send files and assign network addresses to the ESXi hosts that Auto Deploy provisions. See [vSphere PowerCLI Scenario for vSphere Auto Deploy](#).
- Verify that the ESXi hosts have network connectivity to DHCP, TFTP, and vSphere Auto Deploy servers.
- If you want to use VLANs in your vSphere Auto Deploy environment, you must set up the end to end networking properly. When the host is PXE booting, the firmware driver must be set up to tag the frames with proper VLAN IDs. You must do this set up manually by making the correct changes in the UEFI/BIOS interface. You must also correctly configure the ESXi port groups with the correct VLAN IDs. Ask your network administrator how VLAN IDs are used in your environment.
- Verify that you have enough storage for the vSphere Auto Deploy repository. The vSphere Auto Deploy server uses the repository to store data it needs, including the rules and rule sets you create and the VIBs and image profiles that you specify in your rules.
Best practice is to allocate 2 GB to have enough room for four image profiles and some extra space. Each image profile requires approximately 400 MB. Determine how much space to reserve for the vSphere Auto Deploy repository by considering how many image profiles you expect to use.
- Obtain administrative privileges to the DHCP server that manages the network segment you want to boot from. You can use a DHCP server already in your environment, or install a DHCP server. For your vSphere Auto Deploy setup, replace the `gpxelinux.0` filename with `snponly64.efi.vmw-hardwired` for UEFI or `undionly.kpxe.vmw-hardwired` for BIOS. For more information on DHCP configurations, see [Sample DHCP Configurations](#).
- Secure your network as for any other PXE-based deployment method. vSphere Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or the vSphere Auto Deploy server is not checked during a PXE boot.
- If you want to manage vSphere Auto Deploy with PowerCLI cmdlets, verify that Microsoft .NET Framework 4.5 or 4.5.x and Windows PowerShell 3.0 or 4.0 are installed on a Windows machine. See the *vSphere PowerCLI User's Guide*.
- Set up a remote Syslog server. See the *vCenter Server and Host Management* documentation for Syslog server configuration information. Configure the first host you boot to use the remote Syslog server and apply that host's host profile to all other target hosts. Optionally, install and use VMware vCenter Log Insight, which provides log aggregation and analytics for VMware and non-VMware products, virtual and physical, with near real-time search and analytics of log events.
- If the hosts that you plan to provision with vSphere Auto Deploy are with legacy BIOS, verify that the vSphere Auto Deploy server has an IPv4 address. PXE booting with legacy BIOS firmware is possible only over IPv4. PXE booting with UEFI firmware is possible with either IPv4 or IPv6.

If you want to manage vSphere Auto Deploy with PowerCLI cmdlets, see *Set Up vSphere Auto Deploy and Provision Hosts with vSphere PowerCLI*.

1. Navigate to **Home > Auto Deploy**.

By default, only the administrator role has privileges to use the vSphere Auto Deploy service.

2. On the **Auto Deploy** page, select your vCenter Server from the drop-down menu at the top.

3. Click **Enable Auto Deploy and Image Builder** to activate the service.

If the **Image Builder** service is already active, select the **Configure** tab and click **Enable Auto Deploy Service**.

The **Software Depot** page appears.

4. Configure the TFTP server.

a) Click the **Configure** tab.

b) Click **Download TFTP Boot Zip** to download the TFTP configuration file and unzip the file to the directory in which your TFTP server stores files.

c) Optional: To use a proxy server, click **Add** on the *Auto Deploy Runtime Summary* pane and enter a proxy server URL in the text box.

Using reverse proxy servers can offload the requests made to the vSphere Auto Deploy server.

5. Set up your DHCP server to point to the TFTP server on which the TFTP ZIP file is located.
 - a) Specify the TFTP Server's IP address in DHCP option 66, frequently called `next-server`.
 - b) Specify the boot filename, which is `snponly64.efi.vmw-hardwired` for UEFI or `undionly.kpxe.vmw-hardwired` for BIOS in the DHCP option 67, frequently called `boot-filename`.
6. Set each host you want to provision with vSphere Auto Deploy to network boot or PXE boot, following the manufacturer's instructions.
7. Optional: If you set up your environment to use Thumbprint mode, you can use your own Certificate Authority (CA) by replacing the OpenSSL certificate `rbd-ca.crt` and the OpenSSL private key `rbd-ca.key` with your own certificate and key file.
 The files are in `/etc/vmware-rbd/ssl/`.
 By default, vCenter Server uses VMware Certificate Authority (VMCA).

When you start an ESXi host that is set up for vSphere Auto Deploy, the host contacts the DHCP server and is directed to the vSphere Auto Deploy server, which provisions the host with the image profile specified in the active rule set.

- You can change the default configuration properties of the **Auto Deploy Service**. For more information, see "Configuring vCenter Server" in the *vCenter Server and Host Management* documentation.
- You can change the default configuration properties of the **Image Builder Service**. For more information, see "Configuring vCenter Server" in the *vCenter Server and Host Management* documentation.
- Configure the first host that you provision as a reference host. Use the storage, networking, and other settings you want for your target hosts to share. Create a host profile for the reference host and write a rule that assigns both the already tested image profile and the host profile to target hosts.
- If you have to configure host-specific information, set up the host profile of the reference host to prompt for user input. For more information about host customizations, see the *vSphere Host Profiles* documentation.

Using vSphere Auto Deploy Cmdlets

vSphere Auto Deploy cmdlets are implemented as Microsoft PowerShell cmdlets and included in PowerCLI. Users of vSphere Auto Deploy cmdlets can take advantage of all PowerCLI features.

Experienced PowerShell users can use vSphere Auto Deploy cmdlets just like other PowerShell cmdlets. If you are new to PowerShell and PowerCLI, the following tips might be helpful.

You can type cmdlets, parameters, and parameter values in the PowerCLI shell.

- Get help for any cmdlet by running `Get-Help cmdlet_name`.
- Remember that PowerShell is not case sensitive.
- Use tab completion for cmdlet names and parameter names.
- Format any variable and cmdlet output by using `Format-List` or `Format-Table`, or their short forms `fl` or `ft`. For more information, run the `Get-Help Format-List` cmdlet.

Passing Parameters by Name

You can pass in parameters by name in most cases and surround parameter values that contain spaces or special characters with double quotes.

```
Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile
```

Most examples in the *vCenter Server Installation and Setup* documentation pass in parameters by name.

Passing Parameters as Objects

You can pass parameters as objects if you want to perform scripting and automation. Passing in parameters as objects is useful with cmdlets that return multiple objects and with cmdlets that return a single object. Consider the following example.

1. Bind the object that encapsulates rule set compliance information for a host to a variable.

```
$tr = Test-DeployRuleSetCompliance MyEsxi42
```

2. View the `itemlist` property of the object to see the difference between what is in the rule set and what the host is currently using.

```
$tr.itemlist
```

3. Remediate the host to use the revised rule set by using the `Repair-DeployRuleSetCompliance` cmdlet with the variable.

```
Repair-DeployRuleSetCompliance $tr
```

The example remediates the host the next time you boot the host.

Set Up Bulk Licensing

You can use the vSphere Client or ESXi Shell to specify individual license keys, or you can set up bulk licensing by using PowerCLI cmdlets. Bulk licensing works for all ESXi hosts, but is especially useful for hosts provisioned with vSphere Auto Deploy.

[#unique_239](#).

Assigning license keys through the vSphere Client and assigning licensing by using PowerCLI cmdlets function differently.

Assign license keys with the vSphere Client

You can assign license keys to a host when you add the host to the vCenter Server system or when the host is managed by a vCenter Server system.

Assign license keys with LicenseDataManager PowerCLI

You can specify a set of license keys to be added to a set of hosts. The license keys are added to the vCenter Server database. Each time a host is added to the vCenter Server system or reconnects to it, the host is assigned a license key. A license key that is assigned through PowerCLI is treated as a default license key. When an unlicensed host is added or reconnected, it is assigned the default license key. If a host is already licensed, it keeps its license key.

The following example assigns licenses to all hosts in a data center. You can also associate licenses with hosts and clusters.

The following example is for advanced PowerCLI users who know how to use PowerShell variables.

1. In a PowerCLI session, connect to the vCenter Server system you want to use and bind the associated license manager to a variable.

```
Connect-VIServer -Server 192.XXX.X.XX -User username -Password password
$licenseDataManager = Get-LicenseDataManager
```

2. Run a cmdlet that retrieves the data center in which the hosts for which you want to use the bulk licensing feature are located.

```
$hostContainer = Get-Datacenter -Name Datacenter-X
```

You can also run a cmdlet that retrieves a cluster to use bulk licensing for all hosts in a cluster, or retrieves a folder to use bulk licensing for all hosts in a folder.

3. Create a `LicenseData` object and a `LicenseKeyEntry` object with associated type ID and license key.

```
$licenseData = New-Object VMware.VimAutomation.License.Types.LicenseData
$licenseKeyEntry = New-Object VMware.VimAutomation.License.Types.LicenseKeyEntry
$licenseKeyEntry.TypeId = "vmware-vmware"
$licenseKeyEntry.LicenseKey = "XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX"
```

4. Associate the `LicenseKeys` attribute of the `LicenseData` object you created in step 3 with the `LicenseKeyEntry` object.

```
$licenseData.LicenseKeys += $licenseKeyEntry
```

5. Update the license data for the data center with the `LicenseData` object and verify that the license is associated with the host container.

```
$licenseDataManager.UpdateAssociatedLicenseData($hostContainer.Uid, $licenseData)
$licenseDataManager.QueryAssociatedLicenseData($hostContainer.Uid)
```

6. Provision one or more hosts with vSphere Auto Deploy and assign them to the data center or to the cluster that you assigned the license data to.

7. You can use the vSphere Client to verify that the host is successfully assigned to the default license `XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX`.

All hosts that you assigned to the data center are now licensed automatically.

Reprovisioning Hosts

Use vSphere Auto Deploy to reprovision ESXi hosts with a different image profile or a different host profile.

vSphere Auto Deploy supports multiple reprovisioning options. You can perform a simple reboot or reprovision with a different image profile or a different host profile.

The following reprovisioning operations are available.

- Simple reboot.
- Reboot of hosts for which the user answered questions during the boot operation.
- Reprovision with a different image profile.
- Reprovision with a different host profile.

Reprovision Hosts with Simple Reboot Operations

You can reprovisions ESXi hosts with the image profile, host profile, custom script, and vCenter Server location assigned during first boot.

A simple reboot of a host that is provisioned with vSphere Auto Deploy requires only that all prerequisites are still met. The process uses the previously assigned image profile, host profile, custom script, and vCenter Server location.

- Verify that all associated items like are available. An item can be an image profile, host profile, custom script or vCenter Server inventory location.
- Verify that the host has the identifying information (asset tag, IP address) it had during previous boot operations.

1. Place the host in maintenance mode.

Host Type	Action
Host is part of a DRS cluster	VMware DRS migrates virtual machines to appropriate hosts when you place the host in maintenance mode.
Host is not part of a DRS cluster	You must migrate all virtual machines to different hosts and place each host in maintenance mode.

2. Reboot the host.

The host shuts down. When the host reboots, it uses the image profile that the vSphere Auto Deploy server provides. The vSphere Auto Deploy server also applies the host profile stored on the vCenter Server system.

Use PowerCLI To Reprovision a Host

You can use vSphere Auto Deploy to reprovision a host with a new image profile in a PowerCLI session.

- Verify that the setup you performed during the first boot operation is in place.

Several options for reprovisioning hosts exist.

- If the VIBs that you want to use support live update, you can use an `esxcli software vib update` command. In that case, you must also update the rule set to use an image profile that includes the new VIBs.
- During testing, you can apply an image profile to an individual host with the `Apply-EsxImageProfile` cmdlet and reboot the host so the change takes effect. The `Apply-EsxImageProfile` cmdlet updates the association between the host and the image profile but does not install VIBs on the host.
- In all other cases, use this procedure.

1. At the PowerShell prompt, run the `Connect-VIServerPowerCLI` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Determine the location of a public software depot that contains the image profile that you want to use, or define a custom image profile with vSphere ESXi Image Builder.
3. Run `Add-EsxSoftwareDepot` to add the software depot that contains the image profile to the PowerCLI session.

Depot Type	Cmdlet
Remote depot	Run <code>Add-EsxSoftwareDepot depot_url</code> .
ZIP file	1. Download the ZIP file to a local file path or create a mount point local to the PowerCLI machine. 2. Run <code>Add-EsxSoftwareDepot C:\file_path\my_offline_depot.zip</code>.

4. Run `Get-EsxImageProfile` to see a list of image profiles, and decide which profile you want to use.
5. Run `Copy-DeployRule` and specify the `ReplaceItem` parameter to change the rule that assigns an image profile to hosts.

The following cmdlet replaces the current image profile that the rule assigns to the host with the `my_new_imageprofile` profile. After the cmdlet completes, `myrule` assigns the new image profile to hosts. The old version of `myrule` is renamed and hidden.

```
Copy-DeployRule myrule -ReplaceItem my_new_imageprofile
```

6. Test the rule compliance for each host that you want to deploy the image to.

- a) Verify that you can access the host for which you want to test rule set compliance.

```
Get-VMHost -Name ESXi_hostname
```

- b) Run the cmdlet that tests rule set compliance for the host, and bind the return value to a variable for later use.

```
$str = Test-DeployRuleSetCompliance ESXi_hostname
```

- c) Examine the differences between the contents of the rule set and configuration of the host.

```
$str.itemlist
```

The system returns a table of current and expected items if the host for which you want to test the new rule set compliance is compliant with the active rule set.

CurrentItem	ExpectedItem
-----	-----
my_old_imageprofile	my_new_imageprofile

- d) Remediate the host to use the revised rule set the next time you boot the host.

```
Repair-DeployRuleSetCompliance $str
```

7. Reboot the host to provision it with the new image profile.

Write a Rule and Assign a Host Profile to Hosts

vSphere Auto Deploy can assign a host profile to one or more ESXi hosts.

The host profile might include information about storage configuration, network configuration, or other characteristics of the host. If you add a host to a cluster, that cluster's host profile is used.

- Install PowerCLI and all prerequisite software. For information, see [vCenter Server Installation and Setup](#).
- Export the host profile that you want to use.

In many cases, you assign a host to a cluster instead of specifying a host profile explicitly. The host uses the host profile of the cluster.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Using the vSphere Client, set up a host with the settings you want to use and create a host profile from that host.
3. Find the name of the host profile by running `Get-VMhostProfilePowerCLI` cmdlet, passing in the ESXi host from which you create a host profile.
4. At the PowerCLI prompt, define a rule in which host profiles are assigned to hosts with certain attributes, for example a range of IP addresses.

```
New-DeployRule -Name "testrule2" -Item my_host_profile -Pattern "vendor=Acme,Zven",
"ipv4=192.XXX.1.10-192.XXX.1.20"
```

The specified item is assigned to all hosts with the specified attributes. This example specifies a rule named testrule2. The rule assigns the specified host profile `my_host_profile` to all hosts with an IP address inside the specified range and with a manufacturer of Acme or Zven.

5. Add the rule to the rule set.

```
Add-DeployRule testrule2
```

By default, the working rule set becomes the active rule set, and any changes to the rule set become active when you add a rule. If you use the `NoActivate` parameter, the working rule set does not become the active rule set.

- Assign a host already provisioned with vSphere Auto Deploy to the new host profile by performing compliance test and repair operations on those hosts. For more information, see [Test and Repair Rule Compliance](#).
- Power on unprovisioned hosts to provision them with the host profile.

Test and Repair Rule Compliance

Test new or modified rules for compliance and repair accordingly, as changes in the vSphere Auto Deploy rule set are not updated automatically.

When you add a rule to the vSphere Auto Deploy rule set or modify one or more rules, hosts are not updated automatically. vSphere Auto Deploy applies the new rules only when you test their rule compliance and perform remediation.

- Prepare your system and install the Auto Deploy Server. For more information, see [Prepare Your System for vSphere Auto Deploy](#).
 - Verify that your infrastructure includes one or more ESXi hosts provisioned with vSphere Auto Deploy, and that the host on which you installed PowerCLI can access those ESXi hosts.
1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect to the vCenter Server system that vSphere Auto Deploy is registered with.

```
Connect-VIServer ipv4_or_ipv6_address
```

The cmdlet might return a server certificate warning. In a production environment, make sure no server certificate warnings result. In a development environment, you can ignore the warning.

2. Use PowerCLI to check which vSphere Auto Deploy rules are currently available.

```
Get-DeployRule
```

The system returns the rules and the associated items and patterns.

3. Modify one of the available rules.

For example, you can change the image profile and the name of the rule.

```
Copy-DeployRule -DeployRule testrule -ReplaceItem MyNewProfile
```

You cannot edit a rule already added to the active rule set. Instead, you can copy the rule and replace the item or pattern you want to change.

4. Verify that you can access the host for which you want to test rule set compliance.

```
Get-VMHost -Name MyEsxi42
```

5. Run the cmdlet that tests rule set compliance for the host, and bind the return value to a variable for later use.

```
$tr = Test-DeployRuleSetCompliance MyEsxi42
```

6. Examine the differences between the contents of the rule set and configuration of the host.

```
$tr.itemlist
```

If the host for which you want to test the new rule set compliance is compliant with the active rule set, the system returns a table of current and expected items.

CurrentItem	ExpectedItem
-----	-----
My Profile 25MyNewProfile	

7. Remediate the host to use the revised rule set the next time you boot the host.

```
Repair-DeployRuleSetCompliance $tr
```

If the rule you changed specified the inventory location, the change takes effect when you repair compliance. For all other changes, reboot your host to have vSphere Auto Deploy apply the new rule and to achieve compliance between the rule set and the host.

Collect Logs to Troubleshoot ESXi Hosts

You can collect installation or upgrade log files for ESXi that help you identify the source of the failure if an installation or upgrade fails.

1. Enter the `vm-support` command in the ESXi Shell or through SSH.
2. Navigate to the `/var/tmp/` directory.
3. Retrieve the log files from the `.tgz` file.

vCenter Server Installation and Setup

vCenter Server Installation and Setup describes how to deploy the VMware vCenter Server® appliance.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

vCenter Server Installation and Setup is for anyone who must install and configure VMware vSphere®. These topics are for experienced Microsoft Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

Introduction to vSphere Installation and Setup


vSphere 8.0 provides various options for installation and setup that define a corresponding sequence of tasks.

The two core components of vSphere are ESXi and vCenter Server. ESXi is the virtualization platform on which you can create and run virtual machines and virtual appliances. vCenter Server is a service that acts as a central administrator for ESXi hosts connected in a network. vCenter Server lets you pool and manage the resources of multiple hosts.

You deploy the vCenter Server appliance, a preconfigured virtual machine optimized for running vCenter Server and the vCenter Server components. You can deploy the vCenter Server appliance on ESXi hosts or on vCenter Server instances.

How to Install and Set Up vSphere

vSphere is a sophisticated product with multiple components to install and set up. To ensure a successful vSphere deployment, understand the sequence of tasks required.

Figure 28: vSphere Installation and Setup Workflow

Installing vSphere includes the following tasks:

1. Read the vSphere release notes.
2. Install ESXi.
3. Install vCenter Server.

vCenter Server Components and Services

vCenter Server provides a centralized platform for management, operation, resource provisioning, and performance evaluation of virtual machines and hosts.

The following components are included in the vCenter Server appliance deployments:

- The authentication services contain vCenter Single Sign-On, License service, Lookup Service, and VMware Certificate Authority.
- The vCenter Server group of services contains vCenter Server, vSphere Client, vSphere Auto Deploy, and vSphere ESXi Dump Collector. The vCenter Server appliance also contains the VMware vSphere Lifecycle Manager Extension service and the VMware vCenter Lifecycle Manager.

What Happened to the Platform Services Controller

Beginning in vSphere 7.0, deploying a new vCenter Server or upgrading to vCenter Server 7.0 requires the use of the vCenter Server appliance, a preconfigured virtual machine optimized for running vCenter Server. The new vCenter Server contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, tags, and licensing. It is no longer necessary nor possible to deploy and use an external Platform Services Controller. All Platform Services Controller services are consolidated into vCenter Server, and deployment and administration are simplified.

As these services are now part of vCenter Server, they are no longer described as a part of Platform Services Controller. In vSphere 7.0, the *vSphere Authentication* publication replaces the *Platform Services Controller Administration* publication. The new publication contains complete information about authentication and certificate management. For information about upgrading or migrating from vSphere 6.5 and 6.7 deployments using an existing external Platform Services Controller to vSphere 7.0 using vCenter Server appliance, see the *vSphere Upgrade* documentation.

Authentication Services

vCenter Single Sign-On

The vCenter Single Sign-On authentication service provides secure authentication services to the vSphere software components. By using vCenter Single Sign-On, the vSphere components communicate with each other through a secure token exchange mechanism, instead of requiring each component to authenticate a user separately with a directory service like Active Directory.

vCenter Single Sign-On can authenticate users through:

- **External identity provider federation**
You can configure vCenter Server for an external identity provider using federated authentication. In such a configuration, you replace vCenter Server as the identity provider. Currently, vSphere supports Active Directory Federation Services (AD FS) as the external identity provider. In this configuration, AD FS interacts with the identity sources on behalf of vCenter Server.
- **vCenter Server built-in identity provider**
vCenter Server includes a built-in identity provider. By default, vCenter Server uses the vsphere.local domain as the identity source (but you can change it during installation). You can configure the vCenter Server built-in identity provider to use Active Directory (AD) as its identity source using LDAP/S, OpenLDAP/S, and Integrated Windows Authentication (IWA). Such configurations allow customers to log in to vCenter Server using their AD accounts.

Authenticated users can then be assigned registered solution-based permissions or roles within a vSphere environment.

vCenter Single Sign-On is required with vCenter Server.

vSphere License Service

The vSphere License service provides common license inventory and management capabilities to all vCenter Server systems within the Single Sign-On domain.

VMware Certificate Authority

VMware Certificate Authority (VMCA) provisions each ESXi host with a signed certificate that has VMCA as the root certificate authority, by default. Provisioning occurs when the ESXi host is added to vCenter Server explicitly or as part of the ESXi host installation process. All ESXi certificates are stored locally on the host.

For information about all authentication services and capabilities, see *vSphere Authentication*.

Services Installed with vCenter Server

These additional components are installed silently when you install vCenter Server. The components cannot be installed separately as they do not have their own installers.

PostgreSQL

A bundled version of the VMware distribution of PostgreSQL database for vSphere and vCloud Hybrid Services.

vSphere Client

The HTML5-based user interface that lets you connect to vCenter Server instances by using a Web browser. This vSphere Client replaces the Flex-based vSphere Web Client from vSphere 7.0.

vSphere ESXi Dump Collector

The vCenter Server support tool. You can configure ESXi to save the VMkernel memory to a network server, rather than to a disk, when the system encounters a critical failure. The vSphere ESXi Dump Collector collects such memory dumps over the network.

vSphere Auto Deploy

The vCenter Server support tool that can provision hundreds of physical hosts with ESXi software. You can specify the image to deploy and the hosts to provision with the image. Optionally, you can specify host profiles to apply to the hosts, and a vCenter Server location (folder or cluster) for each host.

VMware vSphere Lifecycle Manager

vSphere Lifecycle Manager enables centralized, automated patch and version management for VMware vSphere and offers support for VMware ESXi hosts, virtual machines, and virtual appliances.

What is vCenter Server Appliance

The vCenter Server appliance is a preconfigured virtual machine that is optimized for running vCenter Server and the associated services.

The vCenter Server appliance package contains the following software:

- Photon OS® 3.0
- The vSphere authentication services
- PostgreSQL
- VMware vSphere Lifecycle Manager Extension
- VMware vCenter Lifecycle Manager

Version 8.0 of vCenter Server is deployed with virtual hardware version 10, which supports 64 virtual CPUs per virtual machine in ESXi.

During the deployment, you can choose the vCenter Server appliance size for your vSphere environment size and the storage size for your database requirements.

vCenter Server uses the VMware vSphere Lifecycle Manager Extension service. An external vSphere Lifecycle Manager instance on Windows is no longer required for vSphere centralized automated patch and version management. For information about the vCenter Server see [vCenter Server Components and Services](#).

vCenter Server supports high availability. For information about configuring vCenter Server in a vCenter High Availability cluster, see *vSphere Availability*.

vCenter Server supports file-based backup and restore. For information backing up and restoring, see [File-Based Backup and Restore of vCenter Server](#).

For information about the vCenter Server maximums, see [VMware Configuration Maximums](#).

Understanding vSphere Domains and Domain Names

Each vCenter Server is associated with a vCenter Single Sign-On domain. The domain name defaults to vsphere.local, but you can change it during deployment. The domain determines the local authentication space.

vCenter Single Sign-On Domain

When you deploy a vCenter Server appliance, you are prompted to create a vCenter Single Sign-On domain or join an existing domain.

The domain name is used by the VMware Directory Service (vmdir) for all Lightweight Directory Access Protocol (LDAP) internal structuring.

You can give your domain a unique name. To prevent authentication conflicts, use a name that is not used by OpenLDAP, Microsoft Active Directory, and other directory services.

After you specify the name of your domain, you can add users and groups. You can add an Active Directory or LDAP identity source and allow the users and groups in that identity source to authenticate. You can also add vCenter Server instances, or other VMware products, such as VMware Aria Operations, to the domain.

vCenter Enhanced Linked Mode

vCenter Enhanced Linked Mode allows you to log in to any single instance of vCenter Server and view and manage the inventories of all the vCenter Server systems in the group.

You can join up to 15 vCenter Server appliance deployments with vCenter Enhanced Linked Mode in a single vSphere Single Sign-On domain.

You can create a vCenter Enhanced Linked Mode group during the deployment of vCenter Server appliance.

You can also join a vCenter Enhanced Linked Mode group by moving, or repointing, a vCenter Server from one vSphere domain to another existing domain. See [Repoint vCenter Server to Another vCenter Server in a Different Domain](#) for information on repointing a vCenter Server node.

vCenter Enhanced Linked Mode for vCenter Server Appliance


When you enable Enhanced Linked Mode for vCenter Server appliance deployments, you can connect multiple vCenter Server appliance deployments together to form a domain.

Other features include:

- A simplified backup and restore process. See [File-Based Backup and Restore of vCenter Server](#) for more information.
- A simplified HA process, removing the need for load balancers.
- Up to 15 vCenter Server appliance deployments can be linked together using enhanced linked mode and displayed in a single inventory view.

- For a vCenter High Availability (vCenter HA) cluster, three nodes are considered one logical vCenter Server node. See "vCenter Architecture Overview" in *vSphere Availability* for the vCenter HA architecture overview. A single vCenter Server standard license is needed for one vCenter HA cluster.

Figure 29: Enhanced Linked Mode for vCenter Server Appliance Deployments



Enhanced Linked Mode with Read Only Replication

If a vCenter High Availability (vCenter HA) instance is connected with another vCenter Server instance with enhanced linked mode and vCenter HA failover occurs to the passive node and is unable to communicate with its replication partner on the other vCenter Server node, the replica on the vCenter HA node enters read-only mode.

Joining a vCenter Enhanced Linked Mode Domain

You can join a vCenter Server appliance to another node during deployment of the vCenter Server appliance.

Note:

You can also join a vCenter Enhanced Linked Mode group by moving, or repointing, a vCenter Server from one vSphere domain to another existing domain. See [Repoint vCenter Server to Another vCenter Server in a Different Domain](#) for information on repointing a vCenter Server node.

For example, you can deploy two vCenter Server appliance systems and join the two nodes using vCenter Enhanced Linked Mode.

Before you deploy a new vCenter Server as part of a vCenter Enhanced Linked Mode, create an image-based backup of the existing nodes in your environment. You can use the backup as a precaution in case there is a failure during the deployment process. It is recommended that you shut down all linked vCenters and take offline snapshots to ensure that their databases are in sync during the backup process.

Enhanced Linked Mode requires vCenter instances to be of the same version for database replication to work across the different linked vCenter instances. During the upgrade process, if some of the vCenter instances are of different versions, the Enhanced Link Mode replication features might not work as expected.

If you are deploying the vCenter Server appliance nodes with the UI Installer:

1. For Appliance 1, deploy the vCenter Server appliance as an instance on ESXi Host 1. Synchronize the time settings with ESXi Host 1.
2. For Appliance 2, deploy the vCenter Server appliance as an instance on ESXi Host 1 and configure the time settings so that Appliance 2 is synchronized with ESXi Host 1. In stage 2 you select to join the vCenter Single Sign-On server of the deployed appliance on Appliance 1. For specific instructions, see [Stage 2 - Set up the Newly Deployed vCenter Server Appliance](#).

If you are deploying the vCenter Server appliance nodes with the CLI:

1. Configure the JSON configuration template `embedded_vCSA_on_VC.json` (or `embedded_vCSA_on_ESXi.json`) for Appliance 1 as an instance on ESXi Host 1. See [Prepare Your JSON Configuration File for CLI Deployment](#) for specific instructions on preparing the JSON configuration file.

2. Deploy Appliance 1 by running the `vcsa-cli-installer` command. See [Deploy a vCenter Server Appliance by Using the CLI](#) for instructions.
3. Configure the JSON configuration template `embedded_vCSA_replication_on_VC.json` (or `embedded_vCSA_replication_on_ESXi.json`) for Appliance 2 as an instance on ESXi Host 1. Enter the hostname of the first embedded node in the `replication_partner_hostname` field in the `sso` section.
4. Deploy Appliance 2 by running the `vcsa-cli-installer` command using the `embedded_vCSA_replication_on_VC.json` (or `embedded_vCSA_replication_on_ESXi.json`) file.

Deploying the vCenter Server Appliance

You can deploy the vCenter Server appliance to manage your vSphere environment.

You can deploy the vCenter Server appliance on an ESXi host 6.7 or later, or on an ESXi host or DRS cluster from the inventory of a vCenter Server instance 6.7 or later.

For information about the software included in the vCenter Server appliance 8.0, see [What is vCenter Server Appliance](#).

For information about the software and hardware requirements for deploying the vCenter Server appliance, see [System Requirements for the vCenter Server Appliance](#).

The vCenter Server installer contains executable files both for GUI and CLI deployments.

- The GUI deployment is a two stage process. The first stage is a deployment wizard that deploys the OVA file of the appliance on the target ESXi host or vCenter Server instance. After the OVA deployment finishes, you are redirected to the second stage of the process that sets up and starts the services of the newly deployed appliance.
- The CLI deployment method involves running a CLI command against a JSON file that you previously prepared. The CLI installer parses the configuration parameters and their values from the JSON file and generates an OVF Tool command that automatically deploys and sets up the appliance. The CLI deployment automatically runs both stage 1 then stage 2, with no user interaction required.

The vCenter Server appliance has the following default user names:

User Name	Description
root	Use this user name to log in to the appliance operating system and the vCenter Server Management Interface. You set the password while deploying the virtual appliance.
administrator@ <i>your_domain_name</i>	Use this user name for vCenter Single Sign-On login. You set the password while creating the vCenter Single Sign-On domain. You create a vCenter Single Sign-On domain during the deployment of a vCenter Server appliance in a new vCenter Single Sign-On domain. After you create a vCenter Single Sign-On domain, only the administrator@ <i>your_domain_name</i> user has the privileges required to log in to vCenter Single Sign-On and vCenter Server. The administrator@ <i>your_domain_name</i> user can proceed as follows: • Add an identity source in which additional users and groups are defined to vCenter Single Sign-On. • Give permissions to the users and groups. For information about adding identity sources and giving permissions to the users and groups, see <i>vSphere Authentication</i> .

For information about upgrading and patching the vCenter Server appliance, see *vSphere Upgrade*.

For information about configuring vCenter Server, see *vCenter Server Configuration*.

If you want to set up vCenter Server to use an IPv6 address version, use the fully qualified domain name (FQDN) or host name of the appliance. To set up an IPv4 address, the best practice is to use the FQDN or host name of the appliance, because the IP address can change if assigned by DHCP.

System Requirements for the vCenter Server Appliance

You can deploy the vCenter Server appliance on an ESXi host 6.7 or later, or on a vCenter Server instance 6.7 or later. Your system must also meet specific software and hardware requirements.

When you use Fully Qualified Domain Names, verify that the client machine from which you are deploying the appliance and the network on which you are deploying the appliance use the same DNS server.

Before you deploy the appliance, synchronize the clocks of the target server and all vCenter Server instances on the vSphere network. Unsynchronized clocks might result in authentication problems and can cause the installation to fail or prevent the appliance services from starting. See [Synchronizing Clocks on the vSphere Network](#).

Hardware Requirements for the vCenter Server Appliance

When you deploy the vCenter Server appliance, you can select to deploy an appliance that is suitable for the size of your vSphere environment. The option that you select determines the number of CPUs and the amount of memory for the appliance.

The hardware requirements for a vCenter Server appliance depend on the size of your vSphere inventory.

	Number of vCPUs	Memory
Tiny environment (up to 10 hosts or 100 virtual machines)	2	14 GB
Small environment (up to 100 hosts or 1,000 virtual machines)	4	21 GB
Medium environment (up to 400 hosts or 4,000 virtual machine)	8	30 GB
Large environment (up to 1,000 hosts or 10,000 virtual machines)	16	39 GB
X-Large environment (up to 2,000 hosts or 35,000 virtual machines)	24	58 GB

Note: If you want to add an ESXi host with more than 512 LUNs and 2,048 paths to the vCenter Server inventory, you must deploy a vCenter Server appliance for a large or x-large environment.

Storage Requirements for the vCenter Server Appliance

When you deploy the vCenter Server appliance, the ESXi host or DRS cluster on which you deploy the appliance must meet minimum storage requirements. The required storage depends not only on the size of the vSphere environment and the storage size, but also on the disk provisioning mode.

The storage requirements are different for each vSphere environment size and depend on your database size requirements.

	Default Storage Size	Large Storage Size	X-Large Storage Size
Tiny environment (up to 10 hosts or 100 virtual machines)	579 GB	2019 GB	4279 GB

	Default Storage Size	Large Storage Size	X-Large Storage Size
Small environment (up to 100 hosts or 1,000 virtual machines)	694 GB	2044 GB	4304 GB
Medium environment (up to 400 hosts or 4,000 virtual machine)	908 GB	2208 GB	4468 GB
Large environment (up to 1,000 hosts or 10,000 virtual machines)	1358 GB	2258 GB	4518 GB
X-Large environment (up to 2,000 hosts or 35,000 virtual machines)	2283 GB	2383 GB	4643 GB

Note: The storage requirements include the requirements for the vSphere Lifecycle Manager that runs as a service in the vCenter Server appliance.

Related Links

[Determine the Oracle Database Size and the Storage Size for the New Appliance on page 1293](#)

Before upgrading a vCenter Server appliance or migrating a vCenter Server on Windows that uses an external Oracle database, you must determine the size of the existing database. Based on the size of the existing database, you can calculate the minimum storage size for the new vCenter Server appliance database using an embedded PostgreSQL database.

Software Requirements for the vCenter Server Appliance

The VMware vCenter Server appliance can be deployed on ESXi6.7 hosts or later, or on vCenter Server instances 6.7 or later.

You can deploy the vCenter Server appliance using the GUI or CLI installer. You run the installer from a network client machine that you use to connect to the target server and deploy the appliance on the server. You can connect directly to an ESXi6.7 host on which to deploy the appliance. You can also connect to a vCenter Server6.7 instance to deploy the appliance on an ESXi host or DRS cluster that resides in the vCenter Server inventory.

For information about the requirements for network client machine, see [System Requirements for the vCenter Server Installer](#).

Required Ports for vCenter Server

The vCenter Server system must be able to send data to every managed host and receive data from the vSphere Client. To enable migration and provisioning activities between managed hosts, the source and destination hosts must be able to receive data from each other through predetermined TCP and UDP ports.

vCenter Server is accessed through predetermined TCP and UDP ports. If you manage network components from outside a firewall, you might be required to reconfigure the firewall to allow access on the appropriate ports. For the list of all supported ports and protocols in vSphere, see the VMware Ports and Protocols Tool™ at <https://ports.broadcom.com>.

During installation, if a port is in use or is blocked using a denylist, the vCenter Server installer displays an error message. You must use another port number to proceed with the installation. There are internal ports that are used only for inter-process communication.

VMware uses designated ports for communication. Additionally, the managed hosts monitor designated ports for data from vCenter Server. If a built-in firewall exists between any of these elements, the installer opens the ports during the installation or upgrade process. For custom firewalls, you must manually open the required ports. If you have a firewall

between two managed hosts and you want to perform source or target activities, such as migration or cloning, you must configure a means for the managed hosts to receive data.

To configure the vCenter Server system to use a different port to receive vSphere Client data, see the *vCenter Server and Host Management* documentation.

DNS Requirements for the vCenter Server Appliance

When you deploy the vCenter Server appliance with a static IP address, you ensure that in case of system restart, the IP address of the appliance remains the same.

Before you deploy the vCenter Server appliance with a static IP address, you must verify that this IP address has a valid internal domain name system (DNS) registration.

When you deploy the vCenter Server appliance, the installation of the web server component that supports the vSphere Client fails if the installer cannot look up the fully qualified domain name (FQDN) for the appliance from its IP address. Reverse lookup is implemented using PTR records.

If you plan to use an FQDN for the appliance system name, you must verify that the FQDN is resolvable by a DNS server, by adding forward and reverse DNS A records.

You can use the `nslookup` command to verify that the DNS reverse lookup service returns an FQDN when queried with the IP address and to verify that the FQDN is resolvable.

```
nslookup -nosearch -nodefname FQDN_or_IP_address
```

If you use DHCP instead of a static IP address for the vCenter Server appliance, verify that the appliance name is updated in the domain name service (DNS). If you can ping the appliance name, the name is updated in DNS.

Ensure that the ESXi host management interface has a valid DNS resolution from the vCenter Server and all vSphere Client instances. Ensure that the vCenter Server has a valid DNS resolution from all ESXi hosts and vSphere Client.

vSphere Client Software Requirements

Use of the vSphere Client requires a supported web browser.

VMware has tested and supports the following guest operating systems and browser versions for the vSphere Client.

Supported Guest Operating Systems

- Windows 32-bit and 64-bit
- Mac OS

Supported Browser Versions

- Google Chrome 90 or later
- Mozilla Firefox 80 or later
- Microsoft Edge 89 or later

Note: Later versions of these browsers are likely to work, but have not been tested.

Preparing for Deployment of the vCenter Server Appliance

Before you deploy the vCenter Server appliance, you must download the vCenter Server installer ISO file and mount it to a network virtual machine or physical server from which you want to perform the deployment.

The machine from which you deploy the appliance must run on a Windows, Linux, or Mac operating system that meets the operating system requirements. See [System Requirements for the vCenter Server Installer](#).

System Requirements for the vCenter Server Installer

You can run the vCenter Server GUI or CLI installer from a network client machine that is running on a Windows, Linux, or Mac operating system of a supported version.

To ensure optimal performance of the GUI and CLI installers, use a client machine that meets the minimum hardware requirements.

Table 43: System Requirements for the GUI and CLI Installers

Operating System	Supported Versions	Minimum Hardware Configuration for Optimal Performance
Windows	- Windows 10, 11 - Windows 2016 x64 bit - Windows 2019 x64 bit - Windows 2022 x64 bit	4 GB RAM, 2 CPU having 4 cores with 2.3 GHz, 32 GB hard disk, 1 NIC
Linux	- SUSE 15 - Ubuntu 18.04, 20.04, 21.10	4 GB RAM, 1 CPU having 2 cores with 2.3 GHz, 16 GB hard disk, 1 NIC Note: The CLI installer requires 64-bit OS.
Mac	- macOS 10.15, 11, 12 - macOS Catalina, Big Sur, Monterey	8 GB RAM, 1 CPU having 4 cores with 2.4 GHz, 150 GB hard disk, 1 NIC

Note: For client machines that run on Mac 10.15 or later, concurrent GUI deployments of multiple appliances are unsupported. You must deploy the appliances in a sequence.

Note: Visual C++ redistributable libraries need to be installed to run the CLI installer on versions of Windows older than Windows 10. The Microsoft installers for these libraries are located in the `vcasa-cli-installer/win32/vcredist` directory.

Note: Deploying the vCenter Server appliance with the GUI requires a minimum resolution of 1024x768 to properly display. Lower resolutions can truncate the UI elements.

Download and Mount the vCenter Server Installer

The vCenter Server appliance ISO image contains GUI and CLI installers for the vCenter Server appliance.

- Create an account at <https://support.broadcom.com/web/ecx>.
- Verify that your client machine meets the system requirements for the vCenter Server installer. See [System Requirements for the vCenter Server Installer](#).

With the GUI and CLI executable files that are included in the vCenter Server installer, you can:

- Deploy the vCenter Server appliance.
- Upgrade the vCenter Server appliance.
- Converge older versions of vCenter Server with an external Platform Services Controller to the current version of vCenter Server.

- Restore a vCenter Server appliance from a file-based backup.
1. Log in to Broadcom Support Portal.
 2. Download the vCenter Server appliance ISO image. See the Knowledge Base article <https://knowledge.broadcom.com/external/article/366685/vmware-vsphere-downloads-vmware-converte.html>.
 3. Confirm that the md5sum is correct by using an MD5 checksum tool.
 4. Mount the ISO image to the client machine from which you want to deploy, upgrade, migrate, or restore the appliance.

Note: ISO mounting software that does not allow more than eight directory levels, for example, MagicISO Maker on Windows, is unsupported.

For Linux OS and Mac OS, Archive Manager is unsupported.

For Mac OS, you can use DiskImageMounter.

For Ubuntu 14.04, you can use Disk Image Mounter.

For SUSE 12 OS, you can use the terminal.

```
$ sudo mkdir mount_dir
$ sudo mount -o loop VMware-vCSA-all-version_number-build_number.iso mount_dir
```

Important: Due to a security change in MacOS Catalina, you must modify the security settings on your computer until the vCenter Server deployment completes. If you attempt to run the installer under MacOS Catalina without modifying the security settings, the vCenter Server installer reports the error: `ovftool cannot be opened because the developer cannot be verified`. For more information, see [KB 79416](#).

Open the `readme.txt` file and review the information about the other files and directories in the vCenter Server appliance ISO image.

Synchronizing Clocks on the vSphere Network

Verify that all components on the vSphere network have their clocks synchronized. If the clocks on the physical machines in your vSphere network are not synchronized, SSL certificates and SAML tokens, which are time-sensitive, might not be recognized as valid in communications between network machines.

Unsynchronized clocks can result in authentication problems, which can cause the installation to fail or prevent the vCenter Server `vmware-vpxd` service from starting.

Time inconsistencies in vSphere can cause the first boot of a component in your environment to fail at different services depending on where in the environment time is not accurate and when the time is synchronized. Problems most commonly occur when the target ESXi host for the destination vCenter Server is not synchronized with NTP or PTP. Similarly, issues can arise if the destination vCenter Server migrates to an ESXi host set to a different time due to fully automated DRS.

To avoid time synchronization issues, ensure that the following is correct before installing, migrating, or upgrading a vCenter Server instance.

- The target ESXi host where the destination vCenter Server is to be deployed is synchronized to NTP or PTP.
- The ESXi host running the source vCenter Server is synchronized to NTP or PTP.
- When upgrading or migrating from vSphere 6.7 to vSphere 8.0, if the vCenter Server appliance is connected to an external Platform Services Controller, ensure the ESXi host running the external Platform Services Controller is synchronized to NTP or PTP.
- If you are upgrading or migrating from vSphere 6.7 to vSphere 8.0, verify that the source vCenter Server or vCenter Server appliance and external Platform Services Controller have the correct time.

Verify that any Windows host machine on which vCenter Server runs is synchronized with the Network Time Server (NTP) server.

To synchronize ESXi clocks with an NTP or a PTP server, you can use the VMware Host Client. For information about editing the time configuration of an ESXi host, see topic *Edit the Time Configuration of an ESXi Host in the VMware Host Client* in the *vSphere Single Host Management - VMware Host Client* documentation.

To learn how to change time synchronization settings for vCenter Server, see topic *Configure the System Time Zone and Time Synchronization Settings* in the *vCenter Server Configuration* documentation.

To learn how to edit the time configuration for a host by using the vSphere Client, see topic *Editing the Time Configuration Settings of a Host* in the *vCenter Server and Host Management* documentation.

Synchronizing the Client and Server System Clock

To establish a secure TLS connection to a vCenter Server (the server), the system where you are running the CLI installer (the client) must not have its system clock slower or faster than the server's system clock by an acceptable limit (tolerance).

See [Table 44: Client Clock Tolerance](#) for specific values for each deployment scenario.

Note: The client clock values are applicable only for vCenter Server 6.7 and later.

Table 44: Client Clock Tolerance

Deployment Scenario	Clock Tolerance	Connection Notes
Linking one vCenter Server with another vCenter Server	When deploying the second vCenter Server, the clock tolerance for the client and the first vCenter Server must not exceed 10 minutes.	
Installing a vCenter Server appliance using a container vCenter Server with a <code>*. _on_ _vc. json</code> template.	The maximum clock tolerance between the client and the container vCenter Server is 8 hours 20 minutes.	

Prerequisites for Deploying the vCenter Server Appliance

To ensure a successful deployment of the vCenter Server appliance, you must perform some required tasks and pre-checks before running the installer.

General Prerequisites

- [Download and Mount the vCenter Server Installer](#).

Target System Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [System Requirements for the vCenter Server Appliance](#).
- If you want to deploy the appliance on an ESXi host, verify that the ESXi host is not in lockdown or maintenance mode and not part of a fully automated DRS cluster.
- If you want to deploy the appliance on a DRS cluster of the inventory of a vCenter Server instance, verify that the cluster contains at least one ESXi host that is not in lockdown or maintenance mode.
- If you plan to use NTP servers for time synchronization, verify that the NTP servers are running and that the time between the NTP servers and the target server on which you want to deploy the appliance is synchronized.
- If you want to deploy the appliance on a vSAN ESA cluster with vSAN ESA encryption, you must enable vSAN ESA encryption before installing vCenter Server. vSAN Express Storage Architecture is a next-generation architecture

designed to get the most out of high-performance storage devices, resulting in greater performance and efficiency. You can enable vSAN ESA encryption through vSAN SDK or vSAN API.

vCenter Enhanced Linked Mode Prerequisites

When deploying a new vCenter Server as part of an Enhanced Linked Mode deployment, create an image-based backup of the existing vCenter Server nodes in your environment. You can use the backup as a precaution in case there is a failure during the deployment process.

If the deployment fails, delete the newly deployed vCenter Server appliance, and restore the vCenter Server nodes from their respective image-based backups. You must restore all the nodes in the environment from their image-based backups. Failing to do so can cause the replication partners to be out of synchronization with the restored node.

- To learn more about creating vCenter Enhanced Linked Mode deployments, see [vCenter Enhanced Linked Mode](#).
- To learn about image-based backs, see [Image-Based Backup and Restore of a vCenter Server Environment](#).

Network Prerequisites

If you plan to assign a static IP address and an FQDN as a system name in the network settings of the appliance, verify that you have configured the forward and reverse DNS records for the IP address.

GUI Deployment of the vCenter Server Appliance

You can use the GUI installer to perform an interactive deployment of a vCenter Server appliance.

When you perform the GUI deployment, you download the vCenter Server installer on a network client machine, run the deployment wizard from the client machine, and provide the inputs that are required for the appliance deployment and setup.

The GUI deployment process includes a series of two stages.

Figure 30: Stage 1 - OVA Deployment



The first stage walks you through the deployment wizard to choose the deployment type and appliance settings. This stage completes the deployment of the OVA file on the target server with the deployment type and appliance settings that you provide.

As an alternative to performing the first stage of the deployment with the GUI installer, you can deploy the OVA file of the vCenter Server appliance by using the vSphere Client. After the OVA deployment, you must log in to the vCenter Server Management Interface of the newly deployed appliance to proceed with the second stage of the deployment process. See "Deploy an OVF or OVA Template" in *vSphere Virtual Machine Administration* vSphere Virtual Machine Administration for information about deploying an OVA file using the vSphere Client.

Figure 31: Stage 2 - Appliance Setup



The second stage walks you through the setup wizard to configure the appliance time synchronization and vCenter Single Sign-On. This stage completes the initial setup and starts the services of the newly deployed appliance.

As an alternative to performing the second stage of the deployment with the GUI installer, you can log in to the vCenter Server Management Interface of the newly deployed appliance, https://FQDN_or_IP_address:5480.

Required Information for Deploying a vCenter Server Appliance

When you use the GUI method to deploy a vCenter Server appliance, the wizard prompts you for deployment and setup information. It is a best practice to keep a record of the values that you enter in case you must reinstall the product.

You can use this worksheet to record the information that you need for deploying a vCenter Server appliance.

Table 45: Required Information During Stage 1 of the GUI Deployment Process

Required Information	Default	Your Entry
FQDN or IP address of the target server on which you want to deploy the appliance. The target server can be either an ESXi host or a vCenter Server instance.	-	
HTTPS port of the target server	443	
User name with administrative privileges on the target server - If your target server is an ESXi host, use root. - If your target server is a vCenter Server instance, use <code>user_name@your_domain_name</code>, for example, <code>administrator@vsphere.local</code>.	-	
Password of the user with administrative privileges on the target server	-	
Data center from the vCenter Server inventory on which you want to deploy the appliance . Target server must be a vCenter Server instance. Optionally you can provide a data center folder.	-	
ESXi host or DRS cluster from the data center inventory on which you want to deploy the appliance	-	
VM name for the appliance - Must not contain a percent sign (%), backslash (\), or forward slash (/) - Must be no more than 80 characters in length	vCenter Server	
Password for the root user of the appliance operating system - Must contain only lower ASCII characters without spaces. - Must be at least 8 characters, but no more than 20 characters in length - Must contain at least one uppercase letter - Must contain at least one lowercase letter - Must contain at least one number - Must contain at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!)	-	

Required Information	Default	Your Entry
Deployment size of the vCenter Server appliance for your vSphere environment. See Hardware Requirements for the vCenter Server Appliance for information about the deployment sizes you can select. The option that you select determines the number of CPUs and the amount of memory for the appliance.	Tiny	
Storage size of the vCenter Server appliance for your vSphere environment. The required storage depends not only on the size of the vSphere environment, but also on the disk provisioning mode. See Storage Requirements for the vCenter Server Appliance . Increase the default storage size if you want larger volume for SEAT data (stats, events, alarms, and tasks).	Default	
Name of the datastore on which you want to store the configuration files and virtual disks of the appliance Note: The installer displays a list of datastores that are accessible from your target server.	-	
Activate or deactivate Thin Disk Mode	Deactivated	
Name of the network to which to connect the appliance Note: The installer displays a drop-down menu with networks that depend on the network settings of your target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu. The network must be accessible from the client machine from which you perform the deployment.	-	
IP version for the appliance address Can be either IPv4 or IPv6.	IPv4	
IP assignment for the appliance address Can be either static or DHCP.	static	
FQDN For a static IP assignment vCenter Server uses FQDN or IP address as the system name.	-	
IP address	-	
For IPv4 networks, you can use either a subnet mask or a network prefix. Subnet mask uses a dot decimal notation (for example, 255.255.255.0). An IPv4 network prefix is an integer between 0 and 32. For IPv6 networks, you must use a network prefix. An IPv6 network prefix is an integer between 0 and 128 .	-	
Default gateway	-	
DNS servers separated by commas	-	
System name (FQDN) Only if you use DHCP with IPv4 and have a Dynamic DNS (DDNS) server available in your environment.	-	

Table 46: Required Information During Stage 2 of the GUI Deployment Process

Required Information	Default	Your Entry
Time synchronization settings You can synchronize the time of the appliance either with the time of the ESXi host or with one or more NTP servers. If you want to use more than one NTP servers, you must provide the IP addresses or FQDNs of the NTP servers as a comma-separated list.	Synchronize time with NTP servers	
Activate or deactivate SSH access Note: vCenter Server High Availability requires remote SSH access to the appliance.	Deactivate	
Name for the new vCenter Single Sign-On domain For example, vsphere.local.	-	
Password for the administrator account, administrator@ <i>your_domain_name</i> - Must be at least 8 characters, but no more than 20 characters in length - Must contain at least one uppercase letter - Must contain at least one lowercase letter - Must contain at least one number - Must contain at least one special character, such as ampersand (&), hash key (#), and percent sign (%)	-	
Password of the vCenter Single Sign On administrator user for the domain	-	
Join or do not participate in the VMware Customer Experience Improvement Program (CEIP) For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i> .	Join the CEIP	


Deploy the vCenter Server Appliance by Using the GUI

You can use the GUI installer to perform an interactive deployment of a vCenter Server appliance. You must run the GUI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the appliance.

- See [Prerequisites for Deploying the vCenter Server Appliance](#).

- See [Required Information for Deploying a vCenter Server Appliance](#).

Figure 32: Deployment Workflow of a vCenter Server Appliance



Stage 1 - Deploy the OVA File as a vCenter Server Appliance

With stage 1 of the deployment process, you deploy the OVA file, which is included in the vCenter Server installer, as a vCenter Server appliance.

1. In the vCenter Server installer, navigate to the `vcsa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
2. On the Home page, click **Install** to start the deployment wizard.
3. Review the Introduction page to understand the deployment process and click **Next**.
4. Read and accept the license agreement, and click **Next**.
5. Connect to the target server on which you want to deploy the vCenter Server appliance.

Option	Steps
You can connect to an ESXi host on which to deploy the appliance.	1. Enter the FQDN or IP address of the ESXi host. 2. Enter the HTTPS port of the ESXi host. 3. Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4. Click Next. 5. Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target ESXi host, and click Yes to accept the certificate thumbprint.

Option	Steps
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the appliance.	1. Enter the FQDN or IP address of the vCenter Server instance. 2. Enter the HTTPS port of the vCenter Server instance. 3. Enter the user name and password of user with vCenter Single Sign-On administrative privileges on the vCenter Server instance, for example, the administrator@<i>your_domain_name</i> user. 4. Click Next. 5. Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target vCenter Server instance, and click Yes to accept the certificate thumbprint. 6. Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next. Note: You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode. 7. Select the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next.

6. On the Set up appliance VM page, enter a name for the vCenter Server appliance, set the password for the root user, and click **Next**.

The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.

The password must contain only lower ASCII characters without spaces, at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

7. Select the deployment size for the vCenter Server appliance for your vSphere inventory.

See [Hardware Requirements for the vCenter Server Appliance](#) for information about the deployment sizes you can select. The option that you select determines the number of CPUs and the amount of memory for the appliance.

8. Select the storage size for the vCenter Server appliance, and click **Next**.

The required storage depends not only on the size of the vSphere environment, but also on the disk provisioning mode. See [Storage Requirements for the vCenter Server Appliance](#).

9. Select the storage location for the vCenter Server appliance where all the virtual machine configuration files and virtual disks will be stored.

Option	Action
Install on an existing datastore accessible from the target host	Select a datastore from the list of compatible datastores.
Install on a new vSAN cluster containing the target host	Specify the required details to create a new vSAN cluster or a vSAN Express Storage Architecture (vSAN ESA) cluster to store the vCenter Server appliance.
Install on an existing vSAN datastore and claim additional disks	Specify the required details to create a cluster on the vSAN datastore. This option is displayed only if your environment contains a vSAN datastore.

To enable thin provisioning, select **Enable Thin Disk Mode**. NFS datastores are thin provisioned by default.

10. Optional: If you selected vSAN as your storage location, you must claim disks for storage.

- For vSAN, claim disks separately for cache tier and capacity tier.
- For vSAN ESA, claim disks from the list of compatible disks.

11. On the Configure network settings page, set up the network settings.

The IP address or the FQDN of the appliance is used as a system name. It is recommended to use an FQDN. However, if you want to use an IP address, use static IP address allocation for the appliance, because IP addresses allocated by DHCP might change.

Option	Action
Network	Select the network to which to connect the appliance. The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu.
IP version	Select the version for the appliance IP address. You can select either IPv4 or IPv6.
IP assignment	Select how to allocate the IP address of the appliance. • static The wizard prompts you to enter the IP address and network settings. • DHCP A DHCP server is used to allocate the IP address. Select this option only if a DHCP server is available in your environment. If there is an enabled DDNS in your environment, you can enter a preferred fully qualified domain name (FQDN) for the appliance.
Common Ports	You can customize the HTTP and HTTPS ports (optional). If specifying a custom HTTP and HTTPS port number, ensure that you do not use a port number already in use by vCenter Server, or the default HTTP and HTTPS ports of 80 and 443. Note: 1. You can only customize the HTTP and HTTPS ports during installation. You cannot customize any other port. 2. You cannot change any port after installation. 3. Upgrade or migration procedures persist all port configuration. Port customizations that were previously supported in Windows OS deployments are persisted in a Windows to Linux migration.

12. On the Ready to complete stage 1 page, review the deployment settings for the vCenter Server appliance and click **Finish** to start the OVA deployment process.

13. Wait for the OVA deployment to finish, and click **Continue** to proceed with stage 2 of the deployment process to set up and start the services of the newly deployed appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface to set up and start the services.

The newly deployed vCenter Server appliance is running on the target server but the services are not started.

Stage 2 - Set up the Newly Deployed vCenter Server Appliance

When the OVA deployment finishes, you are redirected to stage 2 of the deployment process to set up and start the services of the newly deployed vCenter Server appliance.

1. Review the introduction to stage 2 of the deployment process and click **Next**.
2. Configure the time settings in the appliance, optionally enable remote SSH access to the appliance, and click **Next**.

Option	Description
Synchronize time with the ESXi host	Enables periodic time synchronization, and VMware Tools sets the time of the guest operating system to be the same as the time of the ESXi host.
Synchronize time with NTP servers	Uses a Network Time Protocol server for synchronizing the time. If you select this option, you must enter the names or IP addresses of the NTP servers separated by commas.

3. Create a new vCenter Single Sign-On domain or join an existing domain.

Option	Description
Create a new Single Sign-On domain	Creates a new vCenter Single Sign-On domain. 1. Enter the domain name, for example <code>vsphere.local</code>. Note: Ensure that the domain name does not contain any upper-case letters. 2. Set the password for the vCenter Single Sign-On administrator account. This is the password for the user <code>administrator@your_domain_name</code>. 3. Confirm the administrator password, and click Next.
Join an existing vCenter Single Sign-On domain	Joins a new vCenter Single Sign-On server to an existing vCenter Single Sign-On domain. You must provide the information about the vCenter Single Sign-On server to which you join the new vCenter Single Sign-On server. 1. Enter the fully qualified domain name (FQDN) or IP address of the vCenter Single Sign-On server to join. 2. Enter the HTTPS port to use for communication with the vCenter Single Sign-On server. 3. Enter the domain name for the vCenter Single Sign-On you are joining, for example <code>vsphere.local</code>. 4. Enter the password of the vCenter Single Sign-On administrator account. 5. Click Next.

When you select to join an existing vCenter Single Sign-On domain, you enable the Enhanced Linked Mode feature. The infrastructure data is replicated with the joined vCenter Single Sign-On server.

4. Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.

For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.

5. On the Ready to complete page, review the configuration settings for the vCenter Server appliance, click **Finish**, and click **OK** to complete stage 2 of the deployment process and set up the appliance.
6. Optional: After the initial setup finishes, enter the URL from the browser with **`https://vcenter_server_appliance_fqdn/ui`** to go to the vSphere Client and log in to the vCenter Server instance in the vCenter Server appliance, or click the **`https://vcenter_server_appliance_fqdn:443`** to go the vCenter Server appliance Getting Started page.
7. Click **Close** to exit the wizard.

You are redirected to the vCenter Server appliance Getting Started page.

You can configure high availability for the vCenter Server appliance. For information about providing vCenter Server appliance high availability, see *vSphere Availability*.

CLI Deployment of the vCenter Server Appliance

You can use the CLI installer to perform a silent deployment of a vCenter Server appliance on an ESXi host or vCenter Server instance.

The CLI deployment process includes downloading the vCenter Server installer on a network virtual machine or physical server from which you want to perform the deployment, preparing a JSON configuration file with the deployment information, and running the deployment command.

You can also watch the video about how to install the vCenter Server appliance by using the CLI installer.

Prepare Your JSON Configuration File for CLI Deployment

Before you run the CLI installer to deploy a vCenter Server appliance, you must prepare a JSON file with configuration parameters and their values for your deployment specification.

- You must be familiar with the JSON syntax.
- [Download and Mount the vCenter Server Installer](#).

The vCenter Server installer contains JSON templates for all deployment options. For information about the templates, see [JSON Templates for CLI Deployment of the vCenter Server Appliance](#).

You can deploy an appliance with minimum configurations by setting values to the configuration parameters in the JSON template for your specification. You can edit the preset values, remove configuration parameters, and add configuration parameters for custom configurations.

For a complete list of the configuration parameters and their descriptions, navigate to the installer subdirectory for your operating system and run the `vcsa-deploy install --template-help` command or see [Deployment Configuration Parameters](#).

1. In the vCenter Server installer, navigate to the `vcsa-cli-installer` directory, and open the `templates` subfolder.
2. Copy the deployment templates from the `install` subfolder to your workspace.

Important: The path to the JSON configuration files must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

3. In a text editor, open the template file for your specification.

To ensure the correct syntax of your JSON configuration file, use a JSON editor.

4. Fill in the values for the required configuration parameters and, optionally, enter additional parameters and their values.
For example, if you want to use an IPv4 DHCP assignment for the network of the appliance, in the `network` subsection of the template, change the value of the `mode` parameter to `dhcp` and remove the default configuration parameters that are for a static assignment.

```
"network": {
  "ip_family": "ipv4",
  "mode": "dhcp"
},
```

Important:

The string values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To set a value that contains a backslash (\) or quotation mark (") character, you must precede the character with the backslash (\) character. For example, `"password": "my\"password"` sets the password `my"password`, `"image": "G:\vcsa\VMware-vCenter-Server-Appliance-8.0.0.XXXX-YYYYYYY_OVF10.ova"` sets the path `G:\vcsa\VMware-vCenter-Server-Appliance-8.0.0.XXXX-YYYYYYY_OVF10.ova`.

The Boolean values must contain only lowercase characters, that is, a value can be either `true` or `false`. For example, `"ssh_enable": false`.

5. Optional: Use a JSON editor of your choice to validate the JSON file.
6. Save in UTF-8 format and close the file.

You can create and save additional templates if needed for your deployment specification.

JSON Templates for CLI Deployment of the vCenter Server Appliance

The vCenter Server installer includes JSON templates with the minimum configuration parameters for all deployment options. The templates are located in the `vcsa-cli-installer/templates/install` directory.

For each deployment option, there is one template for deploying the appliance on an ESXi host and another template for deploying the appliance on a vCenter Server instance.

Table 47: Deployment JSON Templates Included in the vCenter Server Installer

Template	Description
<code>embedded_vCSA_on_ESXi.json</code>	Contains the minimum configuration parameters required for the deployment of a vCenter Server appliance on an ESXi host.
<code>vCSA_with_cluster_on_ESXi.json</code>	Contains the minimum configuration parameters required for the deployment of a vCenter Server appliance with a single node vSAN and a vLCM managed cluster on an ESXi host.
<code>embedded_vCSA_on_VC.json</code>	Contains the minimum configuration parameters required for the deployment of a vCenter Server appliance on a vCenter Server instance.
<code>embedded_vCSA_replication_on_ESXi.json</code>	Contains the minimum configuration parameters required for the deployment of a vCenter Server appliance as a replication partner to another embedded vCenter Server on an ESXi host.

Template	Description
embedded_vCSA_replication_on_VC.json	Contains the minimum configuration parameters required for the deployment of a vCenter Server appliance replication partner to another vCenter Server appliance on a vCenter Server instance.

Deployment Configuration Parameters

When you prepare your JSON configuration files for a CLI deployment, you must set parameters and values to provide input data for the deployment of a vCenter Server appliance.

Sections and Subsections of Configuration Parameters in the JSON Deployment Files

The configuration parameters in the JSON configuration files for a CLI upgrade are organized in sections and subsections.

Table 48: Sections and Subsections of Configuration Parameters in the JSON Deployment Files

Section	Subsection	Description
new_vcsa - describes the appliance you want to deploy	esxi	Use only if you want to deploy the appliance directly on an ESXi host. Contains the configuration parameters that describe the target ESXi host. Depending on whether you want to configure a vSAN or a vLCM Managed Cluster during the deployment, see the relevant section for a detailed information about the configuration parameters. Table 49: Configuration Parameters in the new_vcsa Section, esxi Subsection for Deployments with a vSAN and a vLCM Managed Cluster #unique_531/unique_531_Connect_42_TABLE_636C07CD-1750-4887-9F0E-F3D09933B683-en Note: You must fill in either the <code>esxi</code> or the <code>vc</code> subsection.
	vc	Use only if you want to deploy the appliance on the inventory of a vCenter Server instance. Contains the configuration parameters that describe the target ESXi host or DRS cluster from the vCenter Server inventory. See Table 51: Configuration Parameters in the new_vcsa Section, vc Subsection . Note: You must fill in either the <code>vc</code> or the <code>esxi</code> subsection.
	appliance	Contains the configuration parameters that describe the appliance. See Table 52: Configuration Parameters in the new_vcsa Section, appliance Subsection .
	install_parameters	Optional subsection for the prefix base MAC address allocation options. See Configuration Parameters in the new_vcsa Section, install_parameters Subsection .
	network	Contains the configuration parameters that describe the network settings for the appliance. See Table 54: Configuration Parameters in the new_vcsa Section, network Subsection .
	os	Contains the configuration parameters that describe the operating system settings for the appliance. See Table 55: Configuration Parameters in the new_vcsa Section, os Subsection .
	sso	Contains the configuration parameters that describe the vCenter Single Sign-On settings for the appliance. See Table 56: Configuration Parameters in the new_vcsa Section, sso Subsection .

Section	Subsection	Description
	ovftool_arguments	Optional subsection for adding arbitrary arguments and their values to the OVF Tool command that the installer generates. Important: The vCenter Server installer does not validate the configuration parameters in the <code>ovftool_arguments</code> subsection. If you set arguments that the OVF Tool does not recognize, the deployment might fail.
ceip - describes joining the VMware Customer Experience Improvement Program (CEIP)	settings	Contains only the <code>ceip_enabled</code> configuration parameter to join or not to join the VMware Customer Experience Improvement Program (CEIP). See Table 57: Configuration Parameters in the ceip Section, settings Subsection . Note: If set to <code>true</code> , you must run the CLI deployment command with the <code>--acknowledge-ceip</code> argument. For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i> .

Important:

The string values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To set a value that contains a backslash (\) or a quotation mark (") character, you must precede the character with the backslash (\) character. For example, `"password": "my\"password"` sets the password my"password. Similarly, `"image": "G:\vcsa\VMware-vCenter-Server-Appliance-8.0.0.XXXX-YYYYYYY_OVF10.ova"` sets the path `G:\vcsa\VMware-vCenter-Server-Appliance-8.0.0.XXXX-YYYYYYY_OVF10.ova`.

Boolean values must contain only lowercase characters. It must be either `true` or `false`. For example, `"ssh_enable": false`.

Configuration Parameters in the new_vcsa Section

Table 49: Configuration Parameters in the new_vcsa Section, esxi Subsection for Deployments with a vSAN and a vLCM Managed Cluster

Name	Type	Description
hostname	string	The IP address or FQDN of the target ESXi host on which you want to deploy the appliance.
username	string	A user name with administrative privileges on the target ESXi host, for example, root.
password	string	The password of the user with administrative privileges on the target ESXi host.
deployment_network	string	The name of the network to which to connect the appliance. Note: The network must be accessible from the target ESXi host. Ignored if the target ESXi host has only one network.
datacenter	string	A specific data center you want to create.
cluster	string	The name of the vSAN or vLCM managed cluster.
compression_only	Boolean	Set the value to <code>true</code> to enable compression on the vSAN cluster. If this parameter is set to <code>true</code> , the <code>deduplication_and_compression</code> parameter must be set to <code>false</code> .
deduplication_and_compression	Boolean	Set the value to <code>true</code> to enable compression and deduplication on the vSAN cluster. If this parameter is set to <code>true</code> , the <code>compression_only</code> parameter must be set to <code>false</code> .
cache_disk		A list of UUIDs or canonical names of disks you want to use for cache. Specify only SSDs.

Name	Type	Description
capacity_disk		A list of UUIDs or canonical names of disks you want to use for storage. Specify either SSDs or HDDs.
enable_vlcm	Boolean	Set to <code>true</code> to create a vLCM managed cluster.
enable_vsan_esata	Boolean	Set to <code>true</code> to create a vSAN cluster with vSAN ESA enabled. vSAN Express Storage Architecture is a next-generation architecture designed to get the most out of high-performance storage devices, resulting in greater performance and efficiency.
single_tier	Array	A list of UUIDs or canonical names of disks you want to add to a vSAN storage pool. Required only if you set <code>enable_vsan_esata</code> to <code>true</code> .
vsan_hcl_database_path	String	Local vSAN HCL database path. If the vSAN HCL database present in this location is outdated, the installer will download and replace the older version with the latest version. Required only if you set <code>enable_vsan_esata</code> to <code>true</code> .
datastore	string	The name of the datastore where you want to store the configuration files and virtual disks of the appliance. Note: The datastore must be accessible from the ESXi host. If you are using the thin disk mode, the datastore size should have a minimum of 25GB space.
port	integer	The HTTPS reverse proxy port of the target ESXi host. The default port is 443. Use only if the target ESXi host uses a custom HTTPS reverse proxy port.

Table 50: Configuration Parameters in the `new_vcsa` Section, `esxi` Subsection for Deployments without a vSAN or a vLCM Managed Cluster

Name	Type	Description
hostname	string	The IP address or FQDN of the target ESXi host on which you want to deploy the appliance.
username	string	A user name with administrative privileges on the target ESXi host, for example, root.
password	string	The password of the user with administrative privileges on the target ESXi host.
deployment_network	string	The name of the network to which to connect the appliance. Note: The network must be accessible from the target ESXi host. Ignored if the target ESXi host has only one network.
datastore	string	The name of the datastore where you want to store the configuration files and virtual disks of the appliance. Note: The datastore must be accessible from the ESXi host. If you are using the thin disk mode, the datastore size should have a minimum of 25GB space.
port	integer	The HTTPS reverse proxy port of the target ESXi host. The default port is 443. Use only if the target ESXi host uses a custom HTTPS reverse proxy port.

Table 51: Configuration Parameters in the `new_vcsa` Section, `vc` Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the target vCenter Server instance on which you want to deploy the appliance.
username	string	vCenter Single Sign-On administrator user name on the target vCenter Server instance, for example, administrator@vsphere.local.
password	string	The password of the vCenter Single Sign-On administrator user on the target vCenter Server instance.
deployment_network	string	The name of the network to which to connect the appliance. Note: The network must be accessible from the target ESXi host or DRS cluster on which you want to deploy the appliance. Ignored if the target ESXi host or DRS cluster has only one network.
datacenter	array	The vCenter Server datacenter that contains the target ESXi host or DRS cluster on which you want to deploy the appliance. If the datacenter is located in a folder or a structure of folders, specify the value as a comma-separated list of strings. For example, ["parent_folder", "child_folder", "datacenter_name"] Note: The value is case-sensitive.
datastore	string	The name of the datastore that you want to store the configuration files and virtual disks of the appliance. Note: The datastore must be accessible from the target ESXi host or DRS cluster. The datastore must have at least 25 GB of free space.
port	integer	The HTTPS reverse proxy port of the target vCenter Server instance. The default port is 443. Use only if the target vCenter Server instance uses a custom HTTPS reverse proxy port.
ssl_certificate_verification	string	The CLI verifies that a server's security certificate is signed by a Certificate Authority (CA), and establishes a secure connection. If the certificate is self-signed, the CLI stops the upgrade unless you specify one of the following SSL certificate configuration options: Specify the Secure Hash Algorithm 1 (SHA-1) certificate thumbprint. A certificate thumbprint is a hexadecimal string that uniquely identifies a certificate. The thumbprint is calculated from the content of the certificate using a thumbprint algorithm. "thumbprint": "certificate SHA-1 thumbprint" Set verification_mode to NONE. "verification_mode": "NONE" If you are connecting to a server with a self-signed certificate, and fail to either specify the SHA-1 certificate thumbprint or set the verification mode to NONE, the CLI displays the server's self-signed certificate thumbprint, and prompts you to accept or reject the certificate thumbprint.
target	array	The target ESXi host or DRS cluster on which you want to deploy the appliance. Important: You must provide the name that is displayed in the vCenter Server inventory. For example, if the name of the target ESXi host is an IP address in the vCenter Server inventory, you cannot provide an FQDN. If the target ESXi host or DRS cluster is located in a folder or a structure of folders, specify the value as a comma-separated list of strings. For example, ["parent_folder", "child_folder", "esxi-host.domain.com"] If the target ESXi host is part of a cluster, specify the path as a comma-separated list of strings. For example,

Name	Type	Description
		["cluster_name", "esxi-host.domain.com"] Note: The value is case-sensitive.
vm_folder	string	Optional. The name of the VM folder where the appliance is deployed.

Table 52: Configuration Parameters in the new_vcsa Section, appliance Subsection

Name	Type	Description
thin_disk_mode	Boolean	Set to <code>true</code> to deploy the appliance with thin virtual disks.
deployment_option	string	The size of the appliance. - Set to <code>tiny</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the default storage size. Deploys an appliance with 2 CPUs, 12 GB of memory, and 315 GB of storage. - Set to <code>tiny-lstorage</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the large storage size. Deploys an appliance with 2 CPUs, 12 GB of memory, and 1390 GB of storage. - Set to <code>tiny-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the x-large storage size. Deploys an appliance with 2 CPUs, 12 GB of memory, and 3145 GB of storage. - Set to <code>small</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the default storage size. Deploys an appliance with 4 CPUs, 19 GB of memory, and 380 GB of storage. - Set to <code>small-lstorage</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the large storage size. Deploys an appliance with 4 CPUs, 19 GB of memory, and 1435 GB of storage. - Set to <code>small-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the x-large storage size. Deploys an appliance with 4 CPUs, 19 GB of memory, and 3195 GB of storage. - Set to <code>medium</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the default storage size. Deploys an appliance with 8 CPUs, 28 GB of memory, and 600 GB of storage. - Set to <code>medium-lstorage</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the large storage size. Deploys an appliance with 8 CPUs, 28 GB of memory, and 1600 GB of storage. - Set to <code>medium-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the x-large storage size. Deploys an appliance with 8 CPUs, 28 GB of memory, and 3360 GB of storage. - Set to <code>large</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the default storage size. Deploys an appliance with 16 CPUs, 37 GB of memory, and 965 GB of storage. - Set to <code>large-lstorage</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the large storage size. Deploys an appliance with 16 CPUs, 37 GB of memory, and 1665 GB of storage. - Set to <code>large-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the x-large storage size. Deploys an appliance with 16 CPUs, 37 GB of memory, and 3425 GB of storage. - Set to <code>xlarge</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the default storage size. Deploys an appliance with 24 CPUs, 56 GB of memory, and 1705 GB of storage.

Name	Type	Description
		- Set to <code>xlarge-lstorage</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the large storage size. Deploys an appliance with 24 CPUs, 56 GB of memory, and 1805 GB of storage. - Set to <code>xlarge-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the x-large storage size. Deploys an appliance with 24 CPUs, 56 GB of memory, and 3565 GB of storage.
<code>image</code>	string	Optional. A local file path or URL to the vCenter Server appliance installation package. By default the installer uses the installation package that is included in the ISO file, in the <code>vcsa</code> folder.
<code>name</code>	string	The VM name for the appliance. Must contain only ASCII characters except a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.
<code>ovftool_path</code>	string	Optional. A local file path to the OVF Tool executable file. By default the installer uses the OVF Tool instance that is included in the ISO file, in the <code>vcsa/ovftool</code> folder.

Table 53: Configuration Parameters in the `new_vcsa` Section, `install_parameters` Subsection

Name	Type	Description
<code>mac_allocation_prefix</code>	string	The MAC prefix if a prefix based MAC allocation scheme is used for the VM managed by the appliance.
<code>mac_allocation_prefix_length</code>	integer	The MAC prefix length if a prefix based MAC allocation scheme is used for for the VM managed by the appliance.

Table 54: Configuration Parameters in the `new_vcsa` Section, `network` Subsection

Name	Type	Description
<code>ip_family</code>	string	IP version for the network of the appliance. Set to <code>ipv4</code> or <code>ipv6</code> .
<code>mode</code>	string	IP assignment for the network of the appliance. Set to <code>static</code> or <code>dhcp</code> .
<code>ip</code>	string	IP address for the appliance. Required only if you use static assignment, that is, if you set the <code>mode</code> parameter to <code>static</code> . You must set an IPv4 or IPv6 address that corresponds to the network IP version, that is, to the value of the <code>ip_family</code> parameter. An IPv4 address must comply with the RFC 790 guidelines. An IPv6 address must comply with the RFC 2373 guidelines.
<code>dns_servers</code>	string or array	IP addresses of one or more DNS servers. To set more than one DNS server, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example, <pre>["x.y.z.a", "x.y.z.b"]</pre> or <pre>"x.y.z.a, x.y.z.b"</pre> Optional when the <code>mode</code> parameter is set to <code>static</code> . Not supported when mode is set to <code>DHCP</code> .
<code>prefix</code>	string	Network prefix length.

Name	Type	Description
		Use only if the <code>mode</code> parameter is set to <code>static</code> . Remove if the <code>mode</code> parameter is set to <code>dhcp</code> . The network prefix length is the number of bits that are set in the subnet mask. For example, if the subnet mask is 255.255.255.0, there are 24 bits in the binary version of the prefix length, so the network prefix length is 24. For IPv4 version, the value must be between 0 and 32. For IPv6 version, the value must be between 0 and 128.
gateway	string	IP address of the default gateway. For IPv6 version, the value can be <code>default</code> .
ports	string	Optional. Port numbers that the vCenter Server appliance uses for direct HTTP connections. By default, port 80 redirects requests to HTTPS port 443. You can customize the vCenter Server HTTP and HTTPS ports. If specifying a custom HTTP and HTTPS port number, ensure that you do not use a port already in use by vCenter Server, or the default HTTP and HTTPS ports of 80 and 443. The options to specify a custom port are: <code>"rhttpproxy.ext.port1":"port_number"</code> for the HTTP port, and <code>"rhttpproxy.ext.port2":"port_number"</code> for the HTTPS port. The following example specifies ports 81 and 444 for the HTTP and HTTPS ports: <pre>ports: {"rhttpproxy.ext.port1":"81", "rhttpproxy.ext.port2":"444"}</pre> For more information on ports in use by vCenter Server, see Required Ports for vCenter Server .
system_name	string	Primary network identity. Can be an IP address or FQDN, preferably FQDN. You cannot change the value of this parameter after the deployment. The FQDN and dotted-decimal numbers must comply with the RFC 1123 guidelines.

Table 55: Configuration Parameters in the `new_vcsa` Section, `os` Subsection

Name	Type	Description
password	string	The password for the root user of the appliance operating system. The password must contain between 8 and 20 characters, at least one uppercase letter, at least one lowercase letter, at least one number, and at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!). All characters must be lower ASCII characters without spaces.
ntp_servers	string or array	Optional. Host names or IP addresses of one or more NTP servers for time synchronization. To set more than one NTP server, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example, <pre>["x.y.z.a", "x.y.z.b"]</pre> or <pre>"x.y.z.a, x.y.z.b"</pre>
ssh_enable	Boolean	Set to <code>true</code> to enable SSH administrator login to the appliance. Note: vCenter Server appliance high availability requires remote SSH access to the appliance.
time_tools_sync	Boolean	Optional. Set to <code>true</code> to deploy the appliance with the VMware Tools time synchronization. VMware Tools synchronizes the time of the appliance with the time of the ESXi host. Ignored if you set NTP servers for time synchronization, that is, if you set the <code>ntp_servers</code> parameter.

Table 56: Configuration Parameters in the `new_vcsa` Section, `sso` Subsection

Name	Type	Description
<code>password</code>	string	Password of the vCenter Single Sign-On administrator user, <code>administrator@your_domain_name</code> . If you are deploying a vCenter Server appliance as the first instance in a new vCenter Single Sign-On domain, you must set the password for the vCenter Single Sign-On administrator user. The password must contain between 8 and 20 characters, at least one uppercase letter, at least one lowercase letter, at least one number, and at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!). All characters must be ASCII characters.
<code>domain_name</code>	string	vCenter Single Sign-On domain name, for example, <code>vsphere.local</code> . If you are deploying a vCenter Server appliance as the first instance in a new vCenter Single Sign-On domain, you must set the name for the new vCenter Single Sign-On domain.
<code>replication_partner_hostname</code>	string	The system name of the partner vCenter Server. Required only if you are deploying a replication partner in an existing vCenter Single Sign-On domain.
<code>first_instance</code>	string	Set <code>first_instance</code> to <code>true</code> if you do not want to set the replication partner for this node. If you want to set the replication partner for this node, set <code>first_instance</code> to <code>false</code> and provide the value of the <code>replication_partner_hostname</code> as the FQDN of a previously converged node which is currently in an embedded vCenter Server topology.
<code>sso_port</code>	integer	The HTTPS reverse proxy port of the partner vCenter Server . The default port is 443. Use only if the partner uses a custom HTTPS reverse proxy port.

Configuration Parameters in the `ceip` Section**Table 57: Configuration Parameters in the `ceip` Section, `settings` Subsection**

Name	Type	Description
<code>ceip_enabled</code>	Boolean	Set to <code>true</code> to join the CEIP for this appliance.

Syntax of the CLI Deployment Command

You can use command arguments to set the execution parameters of the deployment command.

You can add a space-separated list of arguments to the CLI deployment command.

```
vcsa-deploy install path_to_the_json_filelist_of_arguments
```

Argument	Description
<code>--accept-eula</code>	Accepts the end-user license agreement. Required for executing the deployment command.
<code>--acknowledge-ceip</code>	Confirms your acknowledgement of your VMware Customer Experience Improvement Program (CEIP) participation. Required if the <code>ceip.enabled</code> parameter is set to <code>true</code> in the JSON deployment template.

Argument	Description
<code>-v, --verbose</code>	Adds debug information to the console output.
<code>-t, --terse</code>	Hides the console output. Displays only warning and error messages.
<code>--log-dir</code> <i>LOG_DIR</i>	Sets the location of the log and other output files.
<code>--skip-ovftool-verification</code>	Performs basic verification of the configuration parameters in the JSON file and deploys the appliance. Does not perform verification of the OVF Tool parameters.
<code>--no-esx-ssl-verify</code>	Skips the SSL verification for ESXi connections. Important: Avoid using this option because it might cause problems during deployment or after deployment because of not validated identity of the target ESXi host.
<code>--no-ssl-certificate-verification</code>	Skips security certificate verification for all server connections.
<code>--operation-id</code> <i>OPERATION_ID</i>	Provides an operation ID to track installation activities.
<code>--pause-on-warnings</code>	Pauses and waits for acknowledgment of warnings.
<code>--verify-template-only</code>	Performs basic template verification of the configuration parameters in the JSON file. Does not deploy the appliance.
<code>--precheck-only</code>	Performs only the basic template verification and OVF Tool parameter verification. Does not deploy the appliance.
<code>--sso-ssl-thumbprint</code> <i>SSL-SHA1-THUMBPRINT</i>	Validates server certificate against the supplied SHA1 thumbprint.
<code>-h, --help</code>	Displays the help message for the <code>vcsa-deploy install</code> command.
<code>--template-help</code>	Displays the help message for the use of configuration parameters in the JSON deployment file.

After the execution finishes, you can get the exit code of the command.

Exit Code	Description
0	Command ran successfully
1	Runtime error
2	Validation error
3	Template error

Deploy a vCenter Server Appliance by Using the CLI

You can use the CLI installer to perform an unattended deployment of a vCenter Server appliance. You must run the CLI deployment from a Windows, Linux, or Mac machine that is in the network on which you want to deploy the appliance.

- See [Prerequisites for Deploying the vCenter Server Appliance](#).
- [Prepare Your JSON Configuration File for CLI Deployment](#).
- Review [Syntax of the CLI Deployment Command](#).
- Verify that the user name with which you are logged in to your client machine, the path to the vCenter Server installer, the path to your JSON configuration file, and the string values in your JSON configuration file contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

- The Visual C++ redistributable libraries version 14.0 or newer needs to be installed to run this utility on versions of Windows older than Windows 10. The Microsoft installers for these libraries are located in the `vcsa\ovftool\win32\vc_redist` directory.
1. Navigate to the `vcsa-cli-installer` subdirectory for your operating system.
 - If you are running the deployment on Windows OS, navigate to the `vcsa-cli-installer\win32` directory.
 - If you are running the deployment on Linux OS, navigate to the `vcsa-cli-installer/lin64` directory.
 - If you are running the deployment on Mac OS, navigate to the `vcsa-cli-installer/mac` directory.
 2. Optional: Run a pre-deployment check without deploying the appliance to verify that you prepared the deployment template correctly.

```
vcsa-deploy install --precheck-only path_to_the_json_file
```

3. Run the deployment command.

```
vcsa-deploy install --accept-eula --acknowledge-ceip optional_arguments path_to_the_json_file
```

Use *optional_arguments* to enter space-separated arguments to set additional execution parameters of the deployment command.

For example, you can set the location of the log and other output files that the installer generates.

```
vcsa-deploy install --accept-eula --acknowledge-ceip --log-dir=path_to_the_location path_to_the_json_file
```

Deploy Multiple vCenter Server Appliances Using the CLI

You can deploy multiple instances vCenter Server appliances concurrently (in batch mode) using the CLI installer.

To deploy multiple instances concurrently, create JSON templates for all the vCenter Server instances in your deployment. The CLI installer assesses the topology of the deployment using the JSON templates, and determines the order. For this reason, the JSON templates must use static IP addresses for all vCenter Server instances in the deployment that are dependant upon one another.

Important: The JSON templates you create for each appliance must use a static IP address to resolve the network addresses of other appliances in the deployment upon which they have a dependency.

To perform the batch deployment, place the JSON templates defining your deployment in a single directory. When invoked, the CLI installer deploys your existing deployment using the topology defined in the JSON templates.

1. In your workspace, create a folder to contain the JSON files for batch deployment. For example, *MyWorkspace/BatchDeploy*.
2. Prepare each JSON configuration file and copy the file to your batch deployment folder. See [Prepare Your JSON Configuration File for CLI Deployment](#) for instructions on configuring the JSON files.
3. Navigate to the `vcsa-cli-installer` subdirectory for your operating system.
 - If you are running the deployment on Windows OS, navigate to the `vcsa-cli-installer\win32` directory.
 - If you are running the deployment on Linux OS, navigate to the `vcsa-cli-installer/lin64` directory.
 - If you are running the deployment on Mac OS, navigate to the `vcsa-cli-installer/mac` directory.
4. Optional: Run a pre-deployment check without deploying the appliance to verify that you prepared the deployment template correctly. For example:

```
vcsa-deploy install --precheck-only MyWorkspace/BatchDeploy
```

5. Run the deployment command. For example,

```
vcsa-deploy install --accept-eula --acknowledge-ceip optional_argumentsMyWorkspace/BatchDeploy
```

Use *optional_arguments* to enter space-separated arguments to set additional execution parameters of the deployment command.

For example, you can set the location of the log and other output files that the installer generates.

```
vcsa-deploy install --accept-eula --acknowledge-ceip --log-dir=path_to_the_locationMyWorkspace/BatchDeploy
```

File-Based Backup and Restore of vCenter Server

vCenter Server supports a file-based backup and restore mechanism that helps you to recover your environment after failures.

You can use the vCenter Server Interface to create a file-based backup of the vCenter Server. After you create the backup, you can restore it by using the GUI installer of the appliance.

You use the vCenter Server Interface to perform a file-based backup of the vCenter Server core configuration, inventory, and historical data of your choice. The backed-up data is streamed over FTP, FTPS, HTTP, HTTPS, SCP, NFS, or SMB to a remote system. The backup is not stored on the vCenter Server.

You can perform a file-based restore only for a vCenter Server that you have previously backed up by using the vCenter Server Interface. You can perform such restore operation by using the GUI installer of the vCenter Server appliance. The process consists of deploying a new vCenter Server appliance and copying the data from the file-based backup to the new appliance.

Important: If you back up a vCenter Server High Availability cluster, the backup operation only backs up the primary vCenter Server instance. Before restoring a vCenter Server High Availability cluster, you must power off the active, passive, and witness nodes. The restore operation restores the vCenter Server in non- vCenter Server High Availability mode. You must reconstruct the cluster after the restore operation completes successfully.

Considerations and Limitations for File-Based Backup and Restore

When you backup or restore a vCenter Server environment, take into account these considerations and limitations.

Protocols

The following considerations apply to file-based backup and restore protocols:

- FTP and HTTP are not secure protocols.
- Backup servers must support minimum of 10 simultaneous connections for each vCenter Server.
- You can use an FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB server that is up and running with sufficient disk space to store the backup.
- You must have write permissions for upload and read permissions for download.
- Only explicit mode is supported for FTPS.
- If you use HTTP or HTTPS, you must enable WebDAV on the backup Web server.
- You can use only FTP, FTPS, HTTP, or HTTPS to transmit data through an HTTP proxy server.
- You can use IPv4 and IPv6 URLs in file-based backup and restore of a vCenter Server. Mixed mode of IP versions between the backup server and the vCenter Server is unsupported.

Configuration

After a restore, the following configurations revert to the state when the backup was taken.

- Virtual machine resource settings

- Resource pool hierarchy and setting
- Cluster-host membership
- DRS configuration and rules

Storage DRS

If the configuration changes, the following might change after a restore.

- Datastore Cluster configuration
- Datastore Cluster membership
- Datastore I/O Resource Management (Storage I/O Control) settings
- Datastore-Datacenter membership
- Host-Datastore membership

Distributed Power Management

If you put a host into standby mode after a backup, the vCenter Server might force the host to exit standby mode when you restore to the backup.

Distributed Virtual Switch

If you use a distributed virtual switch, you are advised to export separately the distributed virtual switch configuration before you restore from a backup. You can import the configuration after the restore. If you omit this consideration, you may lose the changes made to a distributed virtual switch after the backup. For detailed steps, see the VMware knowledge base article at <http://kb.vmware.com/kb/2034602>.

However, when you restore the host-membership after a backup, you will not lose any host that you have added or removed from the distributed virtual switch after the backup.

Note: You can only restore the distributed virtual switch-host membership if the host is part of the cluster.

Content Libraries

If you delete libraries or items after a backup, you cannot access or use these libraries or items after the restore. You can only delete such libraries or items. A warning message notifies you that there are missing files or folders in the storage backup.

If you create new items or item files after the backup, the Content Library Service has no record of the new items or files after the restore operation. A warning notifies you that extra folders or files were found on the storage backup.

If you create new libraries after the backup, the Content Library Service has no record of the new libraries after restore. The library content exists on the storage backing, but no warning is displayed. You must manually clean the new libraries.

Virtual Machine Life Cycle Operations

- Restoring vCenter Server from a backup that was taken during in-flight relocation operations in the vCenter Server instance.
After you restore vCenter Server, the vCenter Server view of the virtual machines might be out of sync with the ESXi view of the virtual machines. This is also true if you performed the backup during in-flight operations on vCenter Server. If virtual machines disappear after you restore vCenter Server, you can refer to the following cases.
 - a. The missing virtual machine is located on the destination ESXi host and is registered with the destination ESXi host, but it is either an orphan or not in the vCenter Server inventory. You must manually add the virtual machine to the vCenter Server inventory.
 - b. The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host and it is not in the vCenter Server inventory. You must manually register the virtual machine to the ESXi host and add the virtual machine back to the vCenter Server inventory.
 - c. The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host. In the vCenter Server instance, the missing virtual machine is marked as orphaned. You must remove the virtual machine from the vCenter Server inventory and add it again.
- Restoring vCenter Server from a backup that has an out-of-date linked clone virtual machine layout.
If you create a linked clone virtual machine after the backup and you restore vCenter Server from the old backup, then after the restore, the vCenter Server does not know about the new linked clone virtual machine until vCenter Server discovers the new linked clone virtual machine. If you remove all existing virtual machines before the new linked clone virtual machine is discovered, then the removal of existing virtual machines corrupts the new linked clone due to missing disks. To avoid this, you must wait until all linked clone virtual machines are discovered by the vCenter Server before you remove virtual machines.
- Restoring vCenter Server from a backup that was taken during virtual machine registration.
If you are registering a virtual machine during the backup and you restore vCenter Server from the old backup, then after the restore, the virtual machine is marked as orphaned in the vCenter Server instance. You must manually add the virtual machine to the vCenter Server inventory.

vCenter Enhanced Linked Mode

To restore a vCenter Server that is part of an Enhanced Linked Mode, you must have at least one running vCenter Server that has the VMware Directory Service database.

If the replication partners are temporarily unavailable or network-isolated, fix the replication partners before restoring the node. If all nodes in the federation have failed, you can ignore the warning and restore only the first node without replication partners. Restore the rest of the nodes according to the Enhanced Linked Mode topology.

vSphere High Availability

Restoring vCenter Server from a backup might cause it to rollback to older version for the vSphere HA cluster state (HostList, ClusterConfiguration, VM protection state) while the hosts in the cluster have the latest version for the cluster state. You need to make sure the vSphere HA cluster state stays the same during restore and backup operations. Otherwise, the following problems might occur.

- If hosts are added or removed to or from the vSphere HA cluster after backup and before vCenter Server restore, virtual machines could potentially failover to hosts not being managed by the vCenter Server but are still part of the HA cluster.
- Protection state for new virtual machines is not updated on the vSphere HA agents on the hosts that are part of the vSphere HA cluster. As a result, virtual machines are not protected or unprotected.
- New cluster configuration state is not updated on the vSphere HA agents on the hosts that are part of the vSphere HA cluster.

vCenter High Availability

Restoring vCenter Server requires vCenter HA to be reconfigured.

Storage Policy Based Management

Restoring vCenter Server from a backup can lead to the following inconsistencies related to storage policies, storage providers, and virtual machines.

- Registered storage providers after backup are lost.
- Unregistered storage providers after backup re-appear and might show different provider status.
- Changes, such as create, delete, or update, performed on storage policies after backup are lost.
- Changes, such as create, delete, or update, performed on storage policy components after backup are lost.
- Default policy configuration changes for datastores performed after backup are lost.
- Changes in the storage policy association of the virtual machine and its disks, and in their policy compliance might occur.

Virtual Storage Area Network

Restoring vCenter Server from a backup might cause inconsistencies in the vSAN. For information on how to check vSAN health, see *Administering VMware vSAN*.

Patching

Restoring vCenter Server from a backup might result in missing security patches. You must apply them again after the restore is complete. For information on patching the vCenter Server, see *vSphere Upgrade*.

Schedule a File-Based Backup

You can schedule file-based backups for vSphere 6.7 and later. You can set up a schedule that is used to perform periodic backups.

- You must have a FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB server up and running with sufficient disk space to store the backup.

The schedule can be set up with information about the backup location, recurrence, and retention for the backups.

You can only set up one schedule at a time.

1. In a Web browser, go to the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
2. Log in as root.
3. In the vCenter Server Management Interface, click **Backup**.
4. Click **Configure** to set up a backup schedule.
5. Enter the backup location details.

Option	Description
Backup location	Enter the backup location, including the protocol to use to connect to your backup server, the port, the server address, and backup folder to store the backup files. Use one of the following protocols: FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB. For FTP, FTPS, HTTP, or HTTPS the path is relative to the home directory configured for the service.

Option	Description
Backup server credentials	Enter a user name and password of a user with write privileges on the backup server.

6. Set the schedule recurrence and time for the backup.
The recurrence can be set daily, weekly, or you can customize the schedule to run the backup on a specific day or days of the week. You can specify the time of day to run the backup. The default time is 11:59pm.
7. Optional: Enter an Encryption Password if you want to encrypt your backup file.
If you select to encrypt the backup data, you must use the encryption password for the restore procedure.
8. Select **Retain all backups** or enter the number of backups to retain.
The retention information provides the number of backups to retain for a given vCenter Server.
Note: This policy applies to both manual and scheduled backups.
9. Optional: Select **Stats, Events, and Tasks** to back up additional historical data from the database.
10. Optional: Select **Supervisors Control Plane** to activate the backup and restore of a Supervisor in the case of any disaster scenario. This allows you to record the the state of the Supervisors Control Planes in a vCenter system.
For more information, see [Backup and Restore the Supervisor Control Plane](#).
11. Click **Create**.
The backup schedule information is populated in the Backup page.

The complete and in progress backups are listed under Activity.

You can perform an immediate backup with the existing schedule information by selecting **Use backup location and user name from backup schedule** from the backup schedule on the Backup Now dialog box.

Manually Back Up vCenter Server by Using the vCenter Server Management Interface

You can use the vCenter Server Management Interface to back up the vCenter Server instance. You can select whether to include historical data, such as stats, events, and tasks, in the backup file.

- You must have an FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB server up and running with sufficient disk space to store the backup.

Note: The backup operation for a vCenter High Availability cluster, backs up only the active node.

1. In a Web browser, go to the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
2. Log in as root.
3. In the vCenter Server Management Interface, click **Backup**.
The table under Activity displays the most current backup version taken of the vCenter Server.
4. Click **Backup Now**.
The **Backup Appliance** wizard opens.

5. Optional: Select **Use backup location and user name from backup schedule** to use the information from a scheduled backup.
6. Enter the backup location details.

Option	Description
Backup location	Enter the backup location, including the protocol to use to connect to your backup server, the port, the server address, and backup folder to store the backup files. Use one of the following protocols: FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB. For FTP, FTPS, HTTP, or HTTPS the path is relative to the home directory configured for the service.
Backup server credentials	Enter a user name and password of a user with write privileges on the backup server. Note: Username and password should only contain ASCII characters.

7. Optional: Enter an Encryption Password if you want to encrypt your backup file.
If you select to encrypt the backup data, you must use the encryption password for the restore procedure.
8. Optional: Select **Stats, Events, and Tasks** to back up additional historical data from the database.
9. Optional: In the **Description** text box, enter a description of the backup.
10. Click **Start** to begin the backup process.

The complete and in progress backups are listed under Activity.


Restore vCenter Server from a File-Based Backup

You can use the vCenter Server appliance GUI installer to restore a vCenter Server to an ESXi host or a vCenter Server instance. The restore procedure has two stages. The first stage deploys a new vCenter Server appliance. The second stage populates the newly deployed vCenter Server appliance with the data stored in the file-based backup.

- Verify that your system meets the minimum software and hardware requirements. See [System Requirements for the vCenter Server Appliance](#).

- If the vCenter Server instance is part of a vCenter High Availability cluster, you must power off the active, passive, and witness nodes of the cluster before restoring the vCenter Server.

Figure 33: vCenter Server Restore Workflow



Stage 1 - Deploy a New Appliance

In stage 1 of the restore process, you deploy the OVA file, which is included in the vCenter Server GUI installer.

- Download and mount the vCenter Server installer. See [Download and Mount the vCenter Server Installer](#).
Note: If you are restoring a backup from a product that has a vCenter Server product patch applied, you must download the full ISO of that particular patch. See [Download VMware vSphere](#).
- If you plan to restore the vCenter Server on an ESXi host, verify that the target ESXi host is not in lockdown or maintenance mode or part of a fully-automated DRS cluster.
- If you plan to restore the vCenter Server on a DRS cluster of a vCenter Server inventory, verify that the cluster contains at least one ESXi host that is not in lockdown or maintenance mode.
- If you plan to assign a static IP address to the appliance, verify that you have configured the forward and reverse DNS records for the IP address.
- If you are attempting to restore a vCenter Server instance that is still running, power off the backed up vCenter Server before you start the restore operation.

As an alternative to performing the first stage of the restore with the GUI installer, you can deploy the OVA file of the new vCenter Server appliance by using the vSphere Client. After the OVA deployment, you must log in to the vCenter Server Management Interface of the newly deployed appliance to proceed with the second stage of the restore process. See "Deploy an OVF or OVA Template" in *vSphere Virtual Machine Administration* vSphere Virtual Machine Administration for information about deploying an OVA file using the vSphere Client.

- In the vCenter Server installer, navigate to the `vcsa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
- On the Home page, click **Restore**.
- Review the Introduction page to understand the restore process and click **Next**.
- Read and accept the license agreement, and click **Next**.
- On the Enter backup details page, enter the details of the backup file that you want to restore, and click **Next**.

Option	Description
Backup location	Enter the server address and backup folder where the backup files are stored. You can also enter the IP address or hostname of the backup server and browse for the location of the backup folder. Specify the protocol to use to retrieve the backup from your backup server. You can select FTP, FTPS, HTTP, HTTPS, SFTP, NFS, or SMB. For example, <pre>ftp://<server_IP_address>/tmp/vCenter/hostname_v-center.com/<backup_filename></pre>
User name	Enter the user name of a user with read privileges on the backup server.
Password	Enter the password of the user with read privileges on the backup server.

6. Review the backup information, and click **Next**.
7. Connect to the ESXi host or vCenter Server on which you want to deploy the vCenter Server appliance to use for the restore operation.

Option	Steps
You can connect to an ESXi host on which to deploy the appliance to use for the restore operation.	1. Enter the FQDN or IP address of the ESXi host. 2. Enter the HTTPS port of the ESXi host. 3. Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4. Click Next. 5. Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target ESXi host, and click Yes to accept the certificate thumbprint.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the appliance to use for the restore operation.	1. Enter the FQDN or IP address of the vCenter Server instance. 2. Enter the HTTPS port of the vCenter Server instance. 3. Enter the user name and password of user with vCenter Single Sign-On administrative privileges on the vCenter Server instance, for example, the administrator@<i>your_domain_name</i> user. 4. Click Next. 5. Verify that the certificate warning displays the SHA1 thumbprint of the SSL certificate that is installed on the target vCenter Server instance, and click Yes to accept the certificate thumbprint. 6. Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next. Note: You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode. 7. Select the ESXi host or DRS cluster on which you want to deploy the appliance, and click Next.

8. Accept the certificate warning.
9. Enter a name for the vCenter Server appliance, set up the password for the root user, and click **Next**.
10. Select the deployment size for the new vCenter Server appliance depending on the size of your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 vCPUs and 14 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 vCPUs and 21 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 vCPUs and 30 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 vCPUs and 39 GB of memory.

Deployment Size Option	Description
	Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 vCPUs and 58 GB of memory. Suitable for environments with up to 2,500 hosts or 45,000 virtual machines

11. Select the storage size for the new vCenter Server appliance, and click **Next**.

Important: You must consider the storage size of the appliance that you are restoring.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 579 GB of storage.	Deploys an appliance with 694 GB of storage.	Deploys an appliance with 908 GB of storage.	Deploys an appliance with 1358 GB of storage.	Deploys an appliance with 2283 GB of storage.
Large	Deploys an appliance with 2019 GB of storage.	Deploys an appliance with 2044 GB of storage.	Deploys an appliance with 2208 GB of storage.	Deploys an appliance with 2258 GB of storage.	Deploys an appliance with 2383 GB of storage.
X-Large	Deploys an appliance with 4279 GB of storage.	Deploys an appliance with 4304 GB of storage.	Deploys an appliance with 4468 GB of storage.	Deploys an appliance with 4518 GB of storage.	Deploys an appliance with 4643 GB of storage.

12. From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**.

13. On the Configure network settings page review the settings populated from the backup file of the vCenter Server.

14. Optional: Edit the network configuration to match the current network environment where the vCenter Server is restored.

15. On the Ready to complete stage 1 page, review the deployment settings for the restored vCenter Server appliance and click **Finish** to start the OVA deployment process.

16. Wait for the OVA deployment to finish, and click **Continue** to proceed with stage 2 of the restore process to transfer the data to the newly deployed appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface to transfer the data.

The newly deployed vCenter Server appliance is running on the target server but the data is not copied from the backup location.

Stage 2 - Transfer Data to the Newly Deployed Appliance

After the OVA deployment finishes, you are redirected to stage 2 of the restore process in which the data from the backup location is copied to the newly deployed vCenter Server appliance.

1. Review the introduction to stage 2 of the restore process and click **Next**.
2. Review the backup details and click **Next**.
3. If you are restoring a node with enhanced linked mode support for vCenter Server appliance, you are asked to provide the Single Sign-On credentials. Enter the Single Sign-On user name and password, then click **Validate and Recover**.
4. On the Ready to complete page, review the details, click **Finish**, and click **OK** to complete stage 2 of the restore process.

The restore process restarts the vCenter Server Management Service. You cannot access the vCenter Server Management API during the restart.

Important: If a restore operation of a vCenter Server appliance VM results with a failure, you must power off and delete the partially restored VM. After that you can try to restore the VM again.

5. Optional: After the restore process finishes, click the **`https://vcenter_server_appliance_fqdn/ui`** to go to the vSphere Client and log in to the vCenter Server instance, or click the **`https://vcenter_server_appliance_fqdn:443`** to go to the vCenter Server Getting Started page.
6. Click **Close** to exit the wizard.
You are redirected to the vCenter Server Getting Started page.
7. If the backed up vCenter node is part of a vCenter High Availability cluster, the last needs to be reconfigured after the restore operation completes successfully.

For information about how to perform backup and restore operations, see *vSphere Availability*.

After successful completion of the restore operation, in the event of a total loss of all storage and servers where all the physical hardware or the data in the hardware is lost, follow the instructions in the VMware knowledge base article at <http://kb.vmware.com/kb/76585> in order to recover the system to a pre-restore state at the time of the backup.

Image-Based Backup and Restore of a vCenter Server Environment

You can use vSphere APIs with a third-party product to back up and restore a virtual machine that contains vCenter Server.

You can perform a full image backup of a virtual machine that contains vCenter Server. The virtual machine must use a fully qualified domain name (FQDN) with correct DNS resolution, or the hostname must be configured to be an IP address.

Considerations and Limitations for Image-Based Backup and Restore

When you restore a vCenter Server environment, take into account these considerations and limitations.

Note: Restoring a vCenter Server instance with DHCP network configuration results in changing its IP address. The changed IP address prevents some vCenter Server services from starting properly. To start all vCenter Server services successfully, after the restore, you must reconfigure the IP address of the restored vCenter Server instance to the IP address that the instance was set to when you performed the backup.

Synchronizing Clocks with NTP Time Synchronization Prior to Backup

Before creating a backup of your vCenter Server deployment, verify that all components on the vSphere network have their clocks synchronized using NTP time synchronization. See [Synchronizing Clocks on the vSphere Network](#).

Configuration

After a restore, the following configurations revert to the state when the backup was taken.

- Virtual machine resource settings
- Resource pool hierarchy and setting
- Cluster-host membership
- DRS configuration and rules

Storage DRS

If the configuration changes, the following might change after a restore.

- Datastore Cluster configuration
- Datastore Cluster membership
- Datastore I/O Resource Management (Storage I/O Control) settings
- Datastore-Datacenter membership
- Host-Datastore membership

Distributed Power Management

If you put a host into standby mode after a backup, the vCenter Server might force the host to exit standby mode when you restore to the backup.

Distributed Virtual Switch

If you use a distributed virtual switch, you are advised to export separately the distributed virtual switch configuration before you restore to a backup. You can import the configuration after the restore. If you omit this consideration, you may lose the changes made to a distributed virtual switch after the backup. For detailed steps, see the VMware Knowledge Base article at <http://kb.vmware.com/kb/2034602>.

Content Libraries

If you delete libraries or items after a backup, you cannot access or use these libraries or items after the restore. You can only delete such libraries or items. A warning message notifies you that there are missing files or folders in the storage backup.

If you create new items or item files after the backup, the Content Library Service has no record of the new items or files after the restore operation. A warning notifies you that extra folders or files were found on the storage backup.

If you create new libraries after the backup, the Content Library Service has no record of the new libraries after restore. The library content exists on the storage backing, but no warning is displayed. You must manually clean the new libraries.

Virtual Machine Life Cycle Operations

- Restoring vCenter Server from a backup that was taken while there are in-flight relocation operations within the vCenter Server instance.
After you restore vCenter Server, the vCenter Server view of the virtual machines may be out of sync with the ESXi view of the virtual machines. This is also true if you performed the backup while there were in-flight operations on vCenter Server. If virtual machines disappear after you restore vCenter Server, you can refer to the following cases.
 - a. The missing virtual machine is located on the destination ESXi host and is registered with the destination ESXi host, but it is not in the vCenter Server inventory. You must manually add the virtual machine to the vCenter Server inventory.
 - b. The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host and it is not in the vCenter Server inventory. You must manually register the virtual machine to the ESXi and add the virtual machine back to the vCenter Server inventory.
 - c. The missing virtual machine is located on the destination ESXi host, but it is not registered with the destination ESXi host. Within the vCenter Server instance, the missing virtual machine is marked as orphaned. You must remove the virtual machine from the vCenter Server inventory and add it again.
- Restoring vCenter Server from a backup that has an out of date linked clone virtual machine layout.
If you create a linked clone virtual machine after the backup and you restore vCenter Server from the old backup, then after the restore, vCenter Server does not know about the new linked clone virtual machine until vCenter Server discovers the new linked clone virtual machine. If you remove all existing virtual machines before the new linked clone virtual machine is discovered, then the removal of existing virtual machines corrupts the new linked clone due to missing disks. To avoid this corruption, you must wait until all linked clone virtual machines get discovered by the vCenter Server before you remove virtual machines.

vSphere High Availability

Restoring vCenter Server from a backup may cause it to roll back to older version for the vSphere HA cluster state (HostList, ClusterConfiguration, VM protection state) while the hosts in the cluster have the latest version for the cluster state. Ensure that the vSphere HA cluster state stays the same during restore and backup operations. Otherwise, the following potential problems are present.

- If hosts are added or removed to/from the vSphere HA cluster after backup and before vCenter Server restore, virtual machines could potentially fail over to hosts not managed by the vCenter Server but are still part of the HA cluster.
- Protection states for new virtual machines are not updated on the vSphere HA agents on the hosts which are part of the vSphere HA cluster. As a result, virtual machines are not protected/unprotected.
- New cluster configuration state is not updated on the vSphere HA agents on the hosts which are part of the vSphere HA cluster.

vCenter High Availability

Restoring vCenter Server requires vCenter HA to be reconfigured.

Storage Policy Based Management

Restoring vCenter Server from a backup can lead to the following inconsistencies related to storage policies, storage providers, and virtual machines.

- Registered storage providers after backup are lost.
- Unregistered storage providers after backup reappear and might show different provider status.
- Changes, such as create, delete, or update, performed on storage policies after backup are lost.
- Changes, such as create, delete, or update, performed on storage policy components after backup are lost.
- Default policy configuration changes for datastores performed after backup are lost.

- Changes in the storage policy association of the virtual machine and its disks, and in their policy compliance might occur.

Virtual Storage Area Network

Restoring vCenter Server from a backup may cause inconsistencies in the vSAN. For information how to check vSAN health, see *Administering VMware vSAN*.

Patching

Restoring vCenter Server from a backup might result in missing security patches. You must apply them again after the restore is complete. For information on patching the vCenter Server appliance, see *vSphere Upgrade*.

Restore a vCenter Server Image-based Environment

You can use a third-party product that uses vSphere APIs to restore a virtual machine that contains vCenter Server.


You can perform an image-based restore of a virtual machine that contains vCenter Server. The virtual machine must use a fully qualified domain name (FQDN) with correct DNS resolution, or the host name of the machine must be configured to be an IP address.

You can restore a virtual machine to the original location by either overwriting the backed up virtual machine or by creating a new virtual machine that contains the restored vCenter Server on the same ESXi host. You can also restore the virtual machine on a new ESXi host.

You can restore a virtual machine that contains vCenter Server directly on the ESXi host that is running the third-party appliance when the vCenter Server service becomes unavailable or when you cannot access the third-party user interface by using the vSphere Client.

Important: Restoring virtual machines that have snapshots or that are configured with Fault Tolerance is unsupported.

Figure 34: vCenter Server Restore Workflow



Restore a vCenter Server Instance

You can use a third-party product to restore a vCenter Server environment.

Important: You can back up and restore only virtual machines that contain vCenter Server instances. You cannot back up and restore physical machines that are running vCenter Server by using a third-party product.

1. Restore the vCenter Server virtual machine onto the ESXi host using a third-party solution.
2. Use the `service-control --status --all` command to verify that the services have started.
Wait for all the vCenter Server services to start, which can take several minutes.
3. Launch the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
4. When you launch the vCenter Server Management Interface, based on the change in the vCenter Server appliance data you will see one of the following:
 - a) If there is a change in the vCenter Server appliance data that needs to be synchronized, reconciliation window will be displayed. Follow the prompts to complete the reconciliation and then provide the Single Sign-On credentials to log into the vCenter Server Management Interface.
 - b) If there is no change in the vCenter Server appliance data, the vCenter Server Management Interface login screen will be displayed.

Power on the restored virtual machine that contains the restored vCenter Server instance.

Restore a vCenter Enhanced Linked Mode Environment

You can restore a vCenter Enhanced Linked Mode environment using a third-party solution.

Important: You can back up and restore only virtual machines that contain vCenter Server instances. You cannot back up and restore physical machines that are running vCenter Server by using a third-party product.

1. Restore the vCenter Server virtual machine onto the ESXi host using a third-party solution.
The services are masked and are not yet running.
2. Use the `systemctl status applmgmt` command to verify that the `systemd` instance of the `applmgmt` service has started.
Wait for all the vCenter Server services to start, which can take several minutes.
3. Log into the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.

Power on the restored virtual machine that contains the restored vCenter Server instance.

After You Deploy the vCenter Server Appliance

After you deploy the vCenter Server appliance, consider these post installation options before adding inventory for the vCenter Server to manage.

If you plan to install Windows 11 as a guest operating system on a virtual machine, you must configure a key provider. Installing Windows 11 requires a Trusted Platform Module (TPM) 2.0. When installing Windows 11 as a guest operating system on a virtual machine, instead of using a physical TPM, you can use a virtual Trusted Platform Module (vTPM). A vTPM is a software-based representation of a physical TPM 2.0 chip. A vTPM depends on virtual machine encryption to secure vital TPM data, and so requires that you configure a key provider. For information about the key providers supported by vSphere, see the Virtual Machine Encryption chapter in the *vSphere Security* documentation. The easiest way is to configure a VMware vSphere® Native Key Provider™. vSphere Native Key Provider is included in all vSphere editions and does not require an external key server. For information about configuring a vSphere Native Key Provider, see Configuring and Managing vSphere Native Key Provider chapter in the *vSphere Security* documentation. As with all security solutions, consider the system design, implementation considerations, and tradeoffs of using vSphere Native Key Provider.

For information about configuring the vSphere Authentication Proxy service, see *vSphere Security*.

For information about upgrading vCenter Server, see *vCenter Server Upgrade*.

How to Log in to Your vCenter Server by Using the vSphere Client

Learn how to use the vSphere Client to log in to vCenter Server and manage your vSphere inventory.

The vSphere Client is automatically installed as part of the vCenter Server appliance deployment. This way, the vSphere Client always points to the same vCenter Single Sign-On instance.

1. Open a Web browser and enter the URL for your vCenter Server instance:

`https://vcenter_server_ip_address_or_fqdn`

As an alternative, you can open a Web browser and enter the URL for the vSphere Client:

`https://vcenter_server_ip_address_or_fqdn/ui.`

2. If a warning message about a potential security risk appears, select to continue to the website.

Browser	Action
Microsoft Edge	1. Click Details. 2. Under the additional message that appears, click Go on to the webpage.
Mozilla Firefox	1. Click Advanced. 2. Under the additional message that appears, click Accept the Risk and Continue.
Google Chrome	1. Click Advanced. 2. Under the additional message that appears, click Proceed to vcenter_server_ip_address_or_fqdn.

3. On the vSphere Welcome page, select **Launch vSphere Client (HTML5)**.
4. If the warning message about a potential security risk appears again, repeat Step 2.
5. Enter the credentials of a user who has permissions on vCenter Server and click **Login**.

The vSphere Client connects to all the vCenter Server systems on which the specified user has permissions, and you can view and manage the vSphere inventory.

Repoint vCenter Server to Another vCenter Server in a Different Domain

You can move a vCenter Server from one vSphere domain to another vSphere domain. Services such as tagging and licensing are retained and migrated to the new domain.

The following use cases are supported:


- You can migrate a vCenter Server from an existing domain to an another existing domain with or without replication. The migrated vCenter Server moves from its current Single Sign-On domain and joins the other existing domain as another vCenter Server connected via enhanced linked mode.
 - See [Repoint a Single vCenter Server Node to an Existing Domain without a Replication Partner](#) for instructions on repointing a single embedded node from one domain to an existing domain without a replication partner.
 - See [Repoint a vCenter Server Node to an Existing Domain with a Replication Partner](#) for instructions on repointing an embedded node from one domain to an existing domain using replication.
- You can migrate a vCenter Server from an existing domain to a newly created domain (where the migrated vCenter Server is the first instance). See [Repoint a vCenter Server Node to a New Domain](#) for instructions of this type of repointing. In this case, there is no replication partner.

Repoint a Single vCenter Server Node to an Existing Domain without a Replication Partner

You can repoint a single vCenter Server from one Single Sign-On domain to an existing Single Sign-On domain without a replication partner. Each Single Sign-On domain contains a single vCenter Server.

- Repointing is only supported with vCenter Server 6.7 Update 1 and later.
- You must repoint to a vCenter Server that is of the same version and nodes that are of the same version and build number.
- To ensure no loss of data, take a file-based backup of each node before proceeding with repointing the vCenter Server.

See [Figure 35: Repointing a Single vCenter Server from One Domain to an Existing Domain](#) for an example of repointing a single vCenter Server from one domain to another existing domain. This is one of many ways to create an Enhanced Linked Mode node. In this case, there is no replication.

Figure 35: Repointing a Single vCenter Server from One Domain to an Existing Domain

1. Make sure that both vCenter Server nodes are powered on before beginning the repointing process.
2. Optional: Run the pre-check mode command. The pre-check mode fetches the tagging (tags and categories) and authorization (roles and privileges) data from the vCenter Server. Pre-check does not migrate any data, but checks for conflicts between the source and destination vCenter Server. For example, run the pre-check with the following CLI:

```
cmsso-util domain-repoint -m pre-check --src-emb-admin Administrator --replication-partner-fqdn FQD-
N_of_destination_node --replication-partner-admin PSC_Admin_of_destination_node --dest-domain-name desti-
nation_PSC_domain
```

Note: Pre-check is not required if a replication partner does not exist (repointing to a newly created domain).

See [Syntax of the Domain Repoint Command](#) for argument definitions for the `cmsso-util domain-repoint` command.

The pre-check writes the conflicts to the `/storage/domain-data` directory.

3. Optional: Review the conflicts and apply resolutions for all conflicts, or apply a separate resolution for each conflict. The conflict resolutions are:
 - Copy: Create a duplicate copy of the data in the target domain.
 - Skip: Skips copying the data in the target domain.
 - Merge: Merges the conflict without creating duplicates.

Note: The default resolution mode for Tags and Authorization conflicts is Copy, unless overridden in the conflict files generated during pre-check.

4. Run the `execute` command. In execute mode, the data generated during the pre-check mode is read and imported to the target node. Then, the vCenter Server is repointed to the target domain. For example, repointing without a replication partner, run the `execute` command with the following:

```
cmsso-util domain-repoint -m execute --src-emb-admin Administrator --replication-partner-fqdn FQDN_of_destination_node --replication-partner-admin PSC_Admin_of_destination_node --dest-domain-name destination_PSC_domain
```


See [Syntax of the Domain Repoint Command](#) for argument definitions for the `cmsso-util domain-repoint` command.

Repoint a vCenter Server Node to an Existing Domain with a Replication Partner

You can repoint a vCenter Server from one Single Sign-On domain to an existing domain using a replication partner.

- Repointing is only supported with vCenter Server 6.7 Update 1 and later.
- You must repoint to a vCenter Server that is of the same version and nodes that are of the same version and build number.
- To ensure no loss of data, take a file-based backup of each node before proceeding with repointing the vCenter Server.

See [Figure 36: Repointing a vCenter Server from One Domain to an Existing Domain](#) for an example of repointing to an existing domain. In this case, there is replication.

Figure 36: Repointing a vCenter Server from One Domain to an Existing Domain

1. Shut down the node (for example, Node C) that is being repointed (moved to a different domain).
2. Decommission the vCenter Server node that is being repointed. For example, to decommission Node C, log into Node B (on the original domain) and run the following command:

```
cmsso-util unregister --node-pnid Node_C_FQDN --username Node_B_sso_administrator@sso_domain.com --password Node_B_sso_adminuser_password
```

After unregistering Node C, services are restarted. References to Node C are deleted from Node B and any other nodes that were linked with Node C on the original domain.

3. Power on Node C to begin the repointing process.
4. Optional: Run the pre-check mode command. The pre-check mode fetches the tagging (tags and categories) and authorization (roles and privileges) data from the vCenter Server. Pre-check does not migrate any data, but checks the conflicts between the source and destination vCenter Server. For example, run the pre-check with the following CLI:

```
cmsso-util domain-repoint -m pre-check --src-emb-admin Administrator --replication-partner-fqdn FQDN_of_destination_node --replication-partner-admin PSC_Admin_of_destination_node --dest-domain-name destination_PSC_domain
```

Note: Pre-check is not required if a replication partner does not exist (repointing to a newly created domain).

See [Syntax of the Domain Repoint Command](#) for argument definitions for the `cmsso-util domain-repoint` command.

The pre-check writes the conflicts to the `/storage/domain-data` directory.

5. Optional: Check conflicts and apply resolutions for all conflicts or apply a separate resolution for each conflict.

The conflict resolutions are:

- Copy: Create a duplicate copy of the data in the target domain.
- Skip: Skips copying the data in the target domain.
- Merge: Merges the conflict without creating duplicates.

Note: The default resolution mode for Tags and Authorization conflicts is Copy, unless overridden in the conflict files generated during pre-check.

6. Run the execute command. In execute mode, the data generated during the pre-check mode is read and imported to the target node. Then, the vCenter Server is repointed to the target domain. For example, run the execute command with the following:

```
cmsso-util domain-repoint -m execute --src-emb-admin Administrator --replication-partner-fqdn FQDN
_of_destination_node --replication-partner-admin destination_node_PSC_Admin_user_name --dest-do-
main-name destination_PSC_domain
```


See [Syntax of the Domain Repoint Command](#) for argument definitions for the `cmsso-util domain-repoint` command.

Repoint a vCenter Server Node to a New Domain

You can repoint a vCenter Server from an existing domain to a newly created domain.

- Repointing is only supported with vCenter Server 6.7 Update 1 and later.
- You must repoint to a vCenter Server that is of the same version and nodes that are of the same version and build number.
- To ensure no loss of data, take a file-based backup and a snapshot in powered-off state of each node before proceeding with repointing the vCenter Server.

See [Figure 37: Repointing a vCenter Server from One Domain to a New Domain](#) for an example of repointing to a new domain. In this case, there is no replication partner.

Figure 37: Repointing a vCenter Server from One Domain to a New Domain

1. Shut down the node (for example, Node C) that is being repointed (moved to a different domain).

Note:

If the vCenter Server is part of a vCenter Server HA cluster, remove the vCenter Server HA configuration before proceeding with the domain repoint. For more information, see *vSphere Availability*.

2. Decommission the vCenter Server node that is being repointed. For example, to decommission Node C, log into Node B (on the original domain) and run the following command:

```
cmsso-util unregister --node-pnid Node_C_FQDN --username Node_B_sso_administrator@sso_domain.com --password Node_B_sso_adminuser_password
```

After unregistering Node C, services are restarted. References to Node C are deleted from Node B and any other nodes that were linked with Node C on the original domain.

3. Power on Node C to begin the repointing process.
4. Run the execute command. In execute mode, the data generated during the pre-check mode is read and imported to the target node. Then, the vCenter Server is repointed to the target domain. For example, repointing with no replication partner (repointing to a new domain), run the execute command with the following:

```
cmsso-util domain-repoint -m execute --src-emb-admin Administrator --dest-domain-name destination_PSC_domain
```

See [Syntax of the Domain Repoint Command](#) for argument definitions for the `cmsso-util domain-repoint` command.

Syntax of the Domain Repoint Command

You can use command arguments to set the execution parameters of the domain-repoint command.

The `cmsso-util domain-repoint` CLI repoints vCenter Server from one domain to another.

You can add a space-separated list of arguments to the CLI repoint command

Use the following command to repoint a vCenter Server to another vCenter Server node:

```
cmsso-util domain-repoint -m execute --src-emb-admin Administrator --replication-partner-fqdn FQDN_of_destination_node --replication-partner-admin destination_node_PSC_Admin_user_name --dest-domain-name destination_PSC_domain
```

Argument	Description
<code>-m, --mode</code>	<i>mode</i> can be <code>pre-check</code> or <code>execute</code> . The <code>pre-check</code> argument runs the command in pre-check mode. The <code>execute</code> argument runs the command in execute mode.
<code>-spa, --src-psc-admin</code>	SSO administrator user name for the source vCenter Server. Do not append the <code>@domain</code> .
<code>-dpf, --dest-psc-fqdn</code>	The FQDN of the vCenter Server to repoint.
<code>-dpa, --dest-psc-admin</code>	SSO administrator user name for the destination vCenter Server. Do not append <code>@domain</code> .
<code>-ddn, --dest-domain-name</code>	SSO domain name of the destination vCenter Server.
<code>-dpr, --dest-psc-rhttps</code>	(Optional) HTTPS port for the destination vCenter Server. If not set, the default 443 is used.
<code>-dvf, --dest-vc-fqdn</code>	The FQDN of the vCenter Server pointing to a destination vCenter Server. The vCenter Server is used to check for component data conflicts in the pre-check mode. If not provided, conflict checks are skipped and the default resolution (COPY) is applied for any conflicts found during the import process. Note: This argument is optional only if the destination domain does not have a vCenter Server. If a vCenter Server exists in the destination domain, this argument is mandatory.
<code>-sea, --src-emb-admin</code>	Administrator for the vCenter Server with embedded vCenter Server. Do not append <code>@domain</code> to the administrator id.
<code>-rpf, --replication-partner-fqdn</code>	(Optional) The FQDN of the replication partner node to which the vCenter Server is replicated.
<code>-rpr, --replication-partner-rhttps</code>	(Optional) The HTTPS port for the replication node. If not set, the default is 443.
<code>-rpa, --replication-partner-admin</code>	(Optional) SSO administrator user name of the replication partner vCenter Server.
<code>-dvr, --dest-vc-rhttps</code>	(Optional) The HTTPS port for the vCenter Server pointing to the destination vCenter Server. If not set, the default 443 is used.
<code>--ignore-snapshot</code>	(Optional) Ignore snapshot warnings.
<code>--no-check-certs</code>	(Optional) Ignore certification validations.
<code>--debug</code>	(Optional) Retrieves command execution detail.
<code>-h, --help</code>	(Optional) Displays the help message for the <code>cmsso-util domain repoint</code> command.

Understanding Tagging and Authorization Conflicts

When you run the domain repoint command in pre-check mode, data from the vCenter Server is exported, examined, and conflicts are written to a file.

The following data is exported to the `/storage/domain-data/` or `ProgramData/VMWare/vCenterServerdata/domain-data` folder:

- `All_Privileges.json`
- `All_Roles.json`
- `All_TagCategories.json`
- `All_Tags.json`

These files contain the all the data (Authorization and Tagging) from the vCenter Server on which this command was run.

If a secondary vCenter Server is provided using the `-dvf` or `--dest-vc-fqdn` option, any conflicts are also exported to the same folder:

- `Conflicts_Roles.json`
- `Conflicts_TagCategories.json`
- `Conflicts_Tags.json`

The following is a sample conflicts file:

```
<---- Sample Conflict file code block ---->
{
  "global" : {
    "resolution" : "MERGE|SKIP|COPY",
    "description" : "Default resolution option used to resolve Role Conflicts is COPY. The
conflicts list describes the differences between Role entities on source and target vCenter Server. If
the source information represents an empty JSON array, it simply means that all the entity
attributes from source and target are identical. If the source lists few entries, it means
that only these entity attributes are missing from the target. If the target lists few entries,
it means that only these entity attributes are missing from the source. Though a global resolution
can be set, it can also be overridden at each conflict level by providing individual resolution
mode."
  },
  "conflicts-count" : 1,
  "conflicts-list" : {
    "NoCryptoAdmin" : {
      "source" : {
        "privileges" : "[]"
      },
      "target" : {
        "privileges" : "[Group-1.SamplePriv-1, Group-1.SamplePriv-4, Group-2.SamplePriv-10,
Group-2.SamplePriv-3, Group-2.SamplePriv-7, Group-3.SamplePriv-2, Group-3.SamplePriv-9]"
      },
      "resolution" : ""
    }
  }
}
<----- End of code block ---->
```

The parts of the sample conflict files are:

- `description`. Provides the details on how the respective conflicts file is read and understood.

- `source` and `target` . JSON objects that list only the differences between the source and target vCenter Server objects.
- `resolution` . User supplies one valid resolution. Valid resolutions are `MERGE` , `COPY` , and `SKIP` .

To specify the resolution for handling conflicts, you can provide a default resolution option all conflicts in the "global" : "resolution" = "MERGE|SKIP|COPY" section. If you do not provide a valid global resolution type for `resolution` or leave it unedited, the system uses `COPY` as the default resolution option.

You can also provide a valid resolution option for each of the conflicts by editing the `resolution` property at each conflict level which overrides the global resolution option.

The types of conflicts listed in [Table 58: Conflict Types](#).

Table 58: Conflict Types

Conflict	Properties used to compare Category Objects	Conflict Types	Conflicting Properties	Conflict Resolution Options
Role conflict	• <code>name</code> : Name of the category. • <code>privilegeId</code> : List of privileges for the role.	RoleName conflict occurs while importing roles and a role with the same name exists in the target vCenter Server but with different privileges.	Properties that can be conflicting for RoleName conflict type can be <code>Privileges</code> .	• <code>COPY</code> . A copy of the conflicting role is created in the target vCenter Server, with <code>--copy</code> appended to the role name. The new role is created with a new role ID with the same set of privilege IDs. The new role ID is updated in the <code>VPX_ACCESS</code> table. The new role ID is applicable for both role name conflict and role ID conflict. Note: The default resolution option to resolve Role conflicts is <code>COPY</code> . • <code>MERGE</code> .The <code>MERGE</code> option is resolved in the following sequence: a. If the source vCenter Server has a role with the same name and privilege list as a role in the target vCenter Server, but the role IDs are different, the role ID from the target vCenter Server is used and updated in the <code>VPX_ACCESS</code> table. b. If the source vCenter Server has a role with the same name as a role in the target vCenter Server, but with a different privilege list, then the privilege lists for both roles are merged. • <code>SKIP</code> . Do nothing. The specific role is skipped.

Conflict	Properties used to compare Category Objects	Conflict Types	Conflicting Properties	Conflict Resolution Options
Tag Category conflict: A category name must be unique in a vCenter Server.	- name : Name of the category. - cardinality: Cardinality of Category, either Single or Multiple. - associableEntityType: List of vCenter Server object that can be associated with a tag from this category. A value of All indicates all vCenter Server objects.	Only one type of conflict can be seen while importing Tag Categories, CategoryName conflict. This conflict indicates that a category with the same name exists in the target vCenter Server but with different properties (cardinality or associableEntityType).	Properties that can be conflicting for conflict type CategoryName can be at least one of two types: Cardinality or AssociableTypes.	- COPY . A copy of the conflicting category is created in the target vCenter Server, with --copy appended to the category name. The new category is created with the same property name as in the source vCenter Server. All the tags that were present under this category is imported under the newly created CategoryCopy. Note: The default resolution option to resolve CategoryName conflicts is COPY. - MERGE . Conflicting properties are merged with the category that is already present in the SSO. Properties are merged as follows: - Description . The description that is already present is used. - Cardinality . Cardinality cannot shrink. If there is a cardinality conflict, the cardinality is set to multiple . It cannot be reduced to single. - AssociableTypes . If either the associableEntityType values are null, it is set to null. Otherwise, Objects types are merged. - SKIP . Do nothing. All tags are imported under the category that exists.
Tags Conflict: A tag object always belongs to a category Object. A tag Name must be unique only inside a category.	- name - description	Only one type of conflict can be seen while importing tags: TagName conflict. This conflict indicates that a Tag with the same name exists under the same category and in the target vCenter Server but with different properties.	Properties that can be conflicting for a conflict of type: TagName can be Description.	- COPY . A copy of the conflicting tag is created in the target vCenter Server, with --copy appended to the tag name. Take the MoRef (Internal tag ID) of the newly created tag and update the tag association if necessary. Note: The default resolution option to resolve CategoryName conflicts is COPY.

Conflict	Properties used to compare Category Objects	Conflict Types	Conflicting Properties	Conflict Resolution Options
				MERGE .Keep the existing description. Take the <code>MoRef</code> (Internal Tag ID) and update one or more Tag Associations if necessary. SKIP . Do nothing. Do not create this tag. Clean up any Tag Associations.

vCenter Server Domain Repoint License Considerations

Domain repointing copies license keys to a new domain. Copying the license keys ensures that valid licensing of all assets is maintained after repointing.

vCenter Server tracks license usage on a per domain basis. If a key is used in more than one domain, you must ensure that the aggregate use of the key does not exceed its capacity. To simplify your license management, remove each license copied to a second domain and assign a new license to assets.

Consider the following two cases:

- License keys that are no longer in use (that is, assigned to assets) in the original domain post repointing.
- License keys that are in use (that is, assigned to assets) in multiple domains.

License Keys Not in Use in a Domain

License Keys in Use in Multiple Domains

If after completing repointing, a license key is in use (that is, assigned to assets) in more than one domain, to remove the license key from all but one domain, first a different license key must be assigned to each asset in domains from which the license key will be removed. Two common approaches:

Troubleshooting vCenter Server Installation or Deployment

The vCenter Server deployment troubleshooting topics provide solutions to problems that you might encounter during the vCenter Server appliance deployment process.

Collecting Logs for Troubleshooting a vCenter Server Installation or Upgrade

You can collect installation or upgrade log files for vCenter Server. If an installation or upgrade fails, checking the log files can help you identify the source of the failure.

You can also collect deployment log files for vCenter Server.

Retrieve Installation Logs Manually

You can retrieve the installation log files manually for examination.

1. Navigate to the installation log file locations.

- `%PROGRAMDATA%\VMware\vCenterServer\logs` directory, usually `C:\ProgramData\VMware\vCenterServer\logs`
- `%TEMP%` directory, usually `C:\Users\username\AppData\Local\Temp`

The files in the `%TEMP%` directory include `vc-install.txt`, `vminst.log`, `pkgmgr.log`, `pkgmgr-comp-msi.log`, and `vim-vcs-msi.log`.

Note:

- In Linux, the installation logs are in `/tmp/vcsaUiInstaller`
- In MacOS, the installation logs are in `/var/folders/**/vcsaUiInstaller`
The installer randomly picks a temporary folder to place the log files. Locate the `vcsaUiInstaller` folder within `/var/folders/` for the log files.

2. Open the installation log files in a text editor for examination.

Collect Deployment Log Files for the vCenter Server Appliance

If the vCenter Server appliance deployment fails, you can retrieve the log files and examine them for the reason of the failure.

The full path to the log files is displayed in the vCenter Server appliance deployment wizard.

In case of firstboot failure, you can download the support bundle on a Windows host machine and examine the log files to determine which firstboot script failed. See [Export a vCenter Server Support Bundle for Troubleshooting](#).

1. On the Windows machine that you use for deploying the vCenter Server appliance, navigate to the log files folder.
You must have administrator privilege to access the log files and the log files are available at:
 - Windows: `C:\Users\{user name}\AppData\Local\Temp\vcsaUiInstaller`
 - Linux: `/tmp/vcsaUiInstaller`
 - Mac: `/var/folders/**/**/**/vcsaUiInstaller`
2. Open the installation log files in a text editor for examination.

Export a vCenter Server Support Bundle for Troubleshooting

You can export the support bundle of the vCenter Server instance in the appliance for troubleshooting using the URL displayed on the DCUI home screen.

You can also collect the support bundle from the vCenter Server appliance Bash shell by running the `vc-support.sh` script.

The support bundle is exported in `.tgz` format.

1. Log in to the Windows host machine on which you want to download the bundle.
2. Open a Web browser and enter the URL to the support bundle displayed in the DCUI.
`https://appliance-fully-qualified-domain-name:443/appliance/support-bundle`
3. Enter the user name and password of the root user.
4. Click **Enter**.
The support bundle is downloaded as `.tgz` file on your Windows machine.

vCenter Server Upgrade

vCenter Server Upgrade describes how to upgrade VMware vCenter Server™ to the current version.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

vCenter Server Upgrade is for anyone who must upgrade from earlier versions of vSphere. These topics are for experienced Microsoft Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

vSphere Upgrade Options


vSphere 8.0 provides many options for upgrading your vSphere deployment. For a successful vSphere upgrade, you must understand the upgrade options, the configuration details that impact the upgrade process, and the sequence of tasks.

The two core components of vSphere are VMware ESXi™ and VMware vCenter Server™. ESXi is the virtualization platform on which you can create and run virtual machines and virtual appliances. vCenter Server is a management plane that acts as a central administrator for ESXi hosts connected in a network. You use the vCenter Server system to pool and manage the resources of multiple hosts. vCenter Server appliance is a preconfigured virtual machine optimized to run vCenter Server.

You can upgrade existing vCenter Server deployments that include either an embedded or an external Platform Services Controller to a deployment consisting of a vCenter Server appliance.

How to Upgrade vSphere

vSphere is a sophisticated product with multiple components to upgrade. Understanding the required sequence of tasks is vital for a successful vSphere upgrade.

Figure 38: Overview of vSphere Upgrade Tasks

Upgrading vSphere includes the following tasks:


1. Read the vSphere release notes.
2. Verify that you have backed up your configuration.
3. If your vSphere system includes VMware solutions or plug-ins, verify that they are compatible with the vCenter Server appliance version to which you are upgrading. See *VMware Product Interoperability Matrix* at <https://compatibilityguide.broadcom.com/>.
4. Upgrade vCenter Server.
For detailed instructions, see [Upgrading the vCenter Server Appliance](#).

5. Upgrade your ESXi hosts.
See [Upgrading ESXi Hosts](#).
6. To ensure sufficient disk storage for log files, consider setting up a syslog server for remote logging. Setting up logging on a remote host is especially important for hosts with a limited amount of local storage.
For detailed instructions, see *ESXi Upgrade*.
See [Required Free Space for System Logging](#) and [Configure Syslog on Hosts](#).
7. Upgrade your VMs manually or by using vSphere Lifecycle Manager to perform an orchestrated upgrade.
See [Upgrading Virtual Machines and VMware Tools](#)

Overview of the vCenter Server Upgrade Process

VMware provides many options to upgrade to vCenter Server 8.0.

You can upgrade or migrate your vCenter Server version 6.7 or version 7.0 installation to version 8.0 .

Figure 39: Overview of vCenter Server Upgrade Tasks

High-level steps for upgrading or migrating vCenter Server:

1. Select your upgrade goal.
 - [Upgrading the vCenter Server Appliance](#)
 - [Migrating vCenter Server for Windows to vCenter Server Appliance](#)
2. Verify that your system meets the hardware and software requirements. See [System Requirements for the vCenter Server Appliance](#).
3. Prepare your environment for the upgrade or migration.
4. Upgrade or migrate your vCenter Server for Windows or vCenter Server appliance deployment.
5. Finish any required post-upgrade or post-migration tasks.

Concurrent upgrades are not supported and upgrade order matters. For information on the order in which to upgrade transitional environments, see [Example Upgrade Paths from vCenter Server 6.7 to vCenter Server 8.0](#).

vCenter Server Supported Upgrade Methods

Graphical User Interface (GUI) Installer

The GUI installer provides a two-step upgrade method using an OVA file that you deploy, and the vCenter Server appliance Management GUI. The first step deploys vCenter Server appliance as an OVA file. The second step uses the vCenter Server Management GUI to configure the new appliance using the source deployment data.

Command Line Interface (CLI) Installer

The CLI installer provides advanced users with a CLI method for upgrading the vCenter Server appliance or migrating vCenter Server for Windows to an appliance. You can upgrade or migrate to a vCenter Server appliance using customized CLI templates.

Migration Assistant Interface for Migrating vCenter Server for Windows to vCenter Server appliance

When you migrate a legacy Platform Services Controller or vCenter Server for Windows to an appliance using the Migration Assistant interface. You can use either the GUI method or the CLI method to migrate the legacy Windows installation data to a target appliance. See [How to Migrate from vCenter Server on Windows to vCenter Server Appliance](#).

Deprecated vCenter Server Deployment Models

When upgrading or migrating from deprecated deployment models, you must first migrate your deployment to a currently supported deployment model before attempting to upgrade or migrate it to a vCenter Server 8.0 deployment. For more information, see [Moving from a Deprecated to a Supported vCenter Server Deployment Topology Before Upgrade or Migration](#).

Patching and Updating vCenter Server

A patch or update brings the vCenter Server 8.0 software up to the current, minor version. You can use the patching process to make minor upgrades to your 8.0 deployment. See [Differences Between vSphere Upgrades, Patches, Updates, and Migrations](#) and [Patching and Updating vCenter Server 8.0 Deployments](#).

vCenter Server Upgrade Compatibility

Upgrading to vCenter Server 8.0 affects other software components of the data center.

[Table 59: Upgrading vCenter Server and Related VMware Products and Components](#) summarizes how upgrading vCenter Server can affect your data center components.

vCenter Server 8.0 can manage ESXi version 6.7 hosts in the same cluster with ESXi 8.0 hosts. vCenter Server 8.0 cannot manage ESXi 6.5 or earlier hosts.

vSphere supports upgrades from vCenter Server 6.7 and later to vCenter Server 8.0. To upgrade from vCenter Server 5.0, 5.1, 5.5, 6.0, or 6.5 you must first upgrade the vCenter Server instance to version 6.7 or later releases, and then upgrade to vCenter Server 8.0. For information about upgrading vCenter Server 5.0, 5.1, 5.5, 6.0, or 6.5 to version 6.7 or 7.0, see the *VMware vSphere 6.7 Documentation* or *VMware vSphere 7.0 Documentation*.

Table 59: Upgrading vCenter Server and Related VMware Products and Components

Product or Component	Compatibility
vCenter Server	Verify support for the upgrade path from your current version of vCenter Server to your planned upgrade version. See the VMware Product Interoperability Matrix at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php .

Product or Component	Compatibility
ESX and ESXi hosts	Verify that your ESX or ESXi host works with the vCenter Server version that you are upgrading to. vCenter Server 8.0 requires ESXi host version 6.7 or later. Upgrade if necessary. See the VMware Product Interoperability Matrix at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php .
VMware Host Profiles	Host Profiles is a vCenter Server tool for designing and deploying ESX and ESXi hosts. Ensure that you are using Host Profiles version 6.0 or later. See Knowledge Base article KB 52932 . For more information on upgrade problems related to Host Profiles, see Upgrade Issues with vCenter Server Containing Host Profiles and the sections on Host Profiles upgrade workflows in the <i>vSphere Host Profiles</i> documentation.
VMFS volumes	You cannot upgrade a VMFS5 datastore to VMFS6. If you have a VMFS5 datastore in your environment, create a VMFS6 datastore and migrate virtual machines from the VMFS5 datastore to VMFS6. For more information on VMFS datastores, see the <i>vSphere Storage</i> documentation.
Virtual machines	Upgrade options depend on your current version. See the information about upgrading virtual machines in the <i>ESXi Upgrade</i> documentation.
VMware Tools	Upgrade options depend on your current version. See the information about upgrading VMware Tools in the <i>ESXi Upgrade</i> documentation.
Auto Deploy	To ensure compatibility and best performance, when you upgrade to vCenter Server 8.0, use Auto Deploy to upgrade ESXi hosts to the same version.
vSphere Distributed Virtual Switch (DVS)	You must upgrade to DVS version 6.6 or later before you upgrade to vCenter Server 8.0. See Knowledge Base article KB 52826 .
vSphere Network I/O Control	DVS version 6.0 or later supports only Network I/O Control version 3. If you are using an earlier version of Network I/O Control you must upgrade to Network I/O Control version 3. For more information, see the <i>vSphere Networking</i> documentation.
vSAN	To avoid potential faults because of differences in the vSAN support in vCenter Server and ESXi, synchronize versions of vCenter Server and ESXi. For the best integration between vSAN components on vCenter Server and ESXi, deploy the latest version of these two vSphere components. For more information, see the <i>ESXi Installation and Setup</i> , <i>vCenter Server Installation and Setup</i> , <i>ESXi Upgrade</i> , and <i>vCenter Server Upgrade</i> documentation.
vSAN disk version	vSAN has several different on-disk format versions available depending on the version and upgrade history of the cluster. Some on-disk format versions are transient while others are intended for long-term production. As certain vSAN features are tied to the on-disk format version, the format version must be accounted for when determining interoperability. See Knowledge Base article KB 2148493 .
Legacy Fault Tolerance	If the vCenter Server inventory contains a VM on which legacy VMware Fault Tolerance (FT) is in use, upgrade or migration will be blocked until you turn off this feature. For information about legacy FT, see Knowledge Base article KB 2143127 . For information on disabling or turning off FT, see Knowledge Base article KB 1008026 .

Differences Between vSphere Upgrades, Patches, Updates, and Migrations

vSphere products distinguish between upgrades, which make major changes to the software, patches and updates, which make smaller changes to the software, and migrations, which make changes to the software platform.

VMware product versions are numbered with two digits, for example, vSphere 8.0. A release that changes either digit, for example, from 6.5 to 6.7, or from 6.7 to 7.0, involves major changes in the software, and requires an upgrade from the previous version. A release that makes a smaller change, requiring only a patch or update, is indicated by an update number, for example, vSphere 6.7 Update 1.

For information about upgrading vCenter Server installations, see [Upgrading the vCenter Server Appliance](#).

For information about patching or updating vCenter Server, see [Patching and Updating vCenter Server 8.0 Deployments](#)

When you upgrade an ESXi host, some host configuration information is preserved in the upgraded version, and the upgraded host, after rebooting, can join a vCenter Server instance that has been upgraded to the same level. Because updates and patches do not involve major changes to the software, host configuration is not affected. For more information, see the *ESXi Upgrade* documentation.

When you upgrade a vCenter Server for Windows instance and convert it to a vCenter Server appliance instance, it is a migration.

For information about migrating a vCenter Server installation to a, see [Migrating vCenter Server for Windows to vCenter Server Appliance](#).

Component Behaviour Changes That Affect vCenter Server Upgrade

When upgrading to vCenter Server 8.0, it is important to understand changes in component behavior for version 8.0 that can affect the upgrade process.

Understanding changes from previous versions of vSphere can help in your upgrade planning. For a complete list of new features in vSphere 8.0, see the Release Notes for version 8.0 releases.

Upgrade Methods

There are multiple methods for upgrading vCenter Server to version 8.0.

Supported Migration Path from vCenter Server for Windows to vCenter Server

You can migrate from an existing vCenter Server for Windows to a vCenter Server 8.0 deployment using a graphical user interface-based installer or a command-line interface-based installer. See [Migrating vCenter Server for Windows to vCenter Server Appliance](#).

Support for Graphical User Interface (GUI) Deployments of vCenter Server

You can upgrade an existing vCenter Server appliance version 6.7 or 7.0 deployment to vCenter Server 8.0 using the GUI. See [GUI Upgrade of the vCenter Server Appliance](#).

Support for Command Line Interface (CLI) Deployments of vCenter Server Appliance

You can upgrade an existing vCenter Server appliance version 6.7 or 7.0 deployment to vCenter Server 8.0 using the CLI. See [CLI Upgrade of the vCenter Server Appliance](#).

Support for Reduced Downtime Upgrade through vSphere Client

You can use a migration-based approach to upgrade vCenter Server between major and minor versions that reduces downtime to several minutes. See [Reduced Downtime Upgrade](#).

Changes in Supported Deployment Types

Changes from previous versions of vSphere can affect your deployment type.

Mixed IPv4 and IPv6 Upgrade and Migration

- Upgrade and migration from vCenter Server 6.7 or 7.0 to 8.0 is supported for pure IPv4 or pure IPv6 management networks only.
- Using dual IP stacks (IPv4 and IPv6) is not supported.

Removal of Platform Services Controller

Beginning in vSphere 7.0, deploying or upgrading vCenter Server in vSphere 7.0 requires the use of vCenter Server appliance, a preconfigured virtual machine optimized for running vCenter Server.

The new vCenter Server contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, and licensing. It is no longer necessary nor possible to deploy and use

an external Platform Services Controller. All Platform Services Controller services are consolidated into vCenter Server, and deployment and administration are simplified.

As these services are now part of vCenter Server, they are no longer described as a part of Platform Services Controller. In vSphere 7.0, the *vSphere Authentication* publication replaces the *Platform Services Controller Administration* publication. The new publication contains complete information about authentication and certificate management.

Upgrade or Migration for vCenter Server Instances with an External Platform Services Controller

When you upgrade or migrate a vCenter Server deployment using an external Platform Services Controller, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the upgrade or migration.

When you converge the Platform Services Controller into a vCenter Server appliance, you must identify the management node to use for your deployment's vSphere SSO domain. In domains with multiple vCenter Server instances, you must specify the SSO replication partner to use for each subsequent vCenter Server.

The method you use to upgrade or migrate determines how you specify the management node to use for your deployment's vSphere SSO domain.

- If you upgrade or migrate using the GUI-based installer, the Upgrade wizard prompts you to specify the replication topology for the vCenter Server deployment.
To learn about upgrading using the GUI-based installer, see [Upgrade a vCenter Server Appliance 7.0 or 7.0 with an External Platform Services Controller Instance by Using the GUI](#). To learn about migrating using the GUI-based installer, see [GUI Migration of vCenter Server with an External Platform Services Controller to an Appliance](#).
- If you upgrade or migrate using the CLI-based installer, you specify the replication topology for the vCenter Server deployment using the JSON templates. The JSON files contain the configuration parameters and their values for your upgrade or migration specification.
To learn about upgrading using the CLI-based installer, see [About Upgrading a vCenter Server with an External Platform Services Controller Using the CLI](#). To learn about migrating using the CLI-based installer, see [About Migrating a vCenter Server for Windows with an External Platform Services Controller Using the CLI](#).
- You can use a migration-based approach to upgrade vCenter Server between major versions that reduces downtime to several minutes. See [Reduced Downtime Upgrade](#).

When the upgrade or migration process finishes, the newly deployed vCenter Server 8.0 appliance incorporates the former Platform Services Controller services. You can then decommission the external Platform Services Controller in your environment. See [Decommission the Platform Services Controller](#).

To learn more about converging a vCenter Server deployment using an external Platform Services Controller into a vCenter Server appliance, see Knowledge Base article [KB 68137](#).

Upgrading or Migrating to vSphere License Service

The License Service provides common license inventory and management capabilities to the vCenter Server systems. When upgrading or migrating your vCenter Server deployment, licensing data is transferred to the License Service in the vCenter Server appliance.

When upgrading vCenter Server versions 7.0 or 7.0 that are connected to a Platform Services Controller, their licensing data is transferred to the License Service in the vCenter Server appliance. The licensing data includes the available licenses and license assignments for hosts, vCenter Server systems, vSAN clusters, and other products that you use with vSphere.

After the upgrade or migration of the vCenter Server systems finishes, the License Services stores the available licenses and manages the license assignments for the entire vSphere environment.

For more information about the License Service and managing licenses in vSphere, see *vCenter Server and Host Management*.

Upgrading the Trust Authority vCenter Server

Upgrade the Trust Authority vCenter Server and Trust Authority Hosts.

Starting with vSphere 8.0, you can take advantage of VMware® vSphere Trust Authority™. vSphere Trust Authority is a foundational technology that enhances workload security. vSphere Trust Authority establishes a greater level of trust in your organization by associating an ESXi host's hardware root of trust to the workload itself.

The best practice for upgrading a vSphere Trust Authority infrastructure is to upgrade the Trust Authority vCenter Server and Trust Authority Hosts first. In this way, you get the most benefit from the latest vSphere Trust Authority features. However, you can perform separate, standalone upgrades of vCenter Server and ESXi hosts to fit specific business reasons. To learn about upgrading the Trust Authority Cluster vCenter Server see "vSphere Trust Authority Life Cycle" in the *vSphere Security* guide.

Support for Upgrading a vCenter Server with Multi-Homing

When upgrading a vCenter Server appliance with two or more network interface cards (NICs), the upgrade transfers the IP addresses from each network interface from the source vCenter Server to the target vCenter Server.

During the upgrade process, the Upgrade installer performs a pre-check of the source vCenter Server and records the network settings. When this information is transferred to the target vCenter Server appliance, the network settings are transferred as part of the configuration. When the upgrade process finishes, you can log in to the vCenter Server appliance using the vSphere Client vSphere Client, and verify that the IP addresses have been successfully transferred to the newly upgraded vCenter Server appliance.

With two or more NICs per vCenter Server appliance, you can better manage network traffic in your environment. For example, with multiple network interfaces you can:

- Configure backup traffic to use a different network than management traffic.
- Keep ESXi hosts on a physically different LAN segment from your management traffic. You can connect one network interface to the network with your ESXi hosts, and connect the other to a network from which administrative clients connect to vCenter Server.

Support for Federal Information Processing Standard 140-2

vCenter Server 8.0 supports Federal Information Processing Standard (FIPS) 140-2.

FIPS 140-2 is a U.S. and Canadian government standard that specifies security requirements for cryptographic modules. By default, FIPS 140-2 is always activated after installation or upgrade of vCenter Server 8.0.

To learn more about support for FIPS 140-2 in VMware products, see <https://www.vmware.com/security/certifications/fips.html>.

To learn how to activate or deactivate FIPS 140-2 support, see the *vSphere Security* documentation.

Support for Transport Security Layer 1.3

Starting in 8.0 Update 3, vSphere supports TLS 1.3 and 1.2 through the use of TLS profiles. By default, vSphere supports the Transport Security Layer (TLS) 1.3 encryption protocol.

Upgrading to vSphere 8.0 Update 3 activates the default TLS profile, named COMPATIBLE, on ESXi and vCenter Server hosts. The COMPATIBLE profile supports TLS 1.3 and some TLS 1.2 connections.

To learn how to manage TLS configuration, see the *vSphere Security* guide.

Moving from a Deprecated to a Supported vCenter Server Deployment Topology Before Upgrade or Migration


Before you upgrade or migrate your environment to vSphere 8.0, you must move any deprecated deployment topology to a supported deployment topology.

When you first install vCenter Server 6.7, your deployment includes either an embedded Platform Services Controller or an external Platform Services Controller.

The installer does not validate whether the Platform Services Controller is external or embedded with vCenter Server. Although many types of join operations are possible, not all resulting topologies are supported. Before you upgrade or migrate your environment to vSphere 8.0, you must move any deprecated deployment topology to a supported deployment topology.


Moving to a Supported Topology from a vCenter Server Pointing to an Embedded Platform Services Controller

Figure 40: Deprecated topology of a vCenter Server Pointing to an Embedded Platform Services Controller



Moving to a Supported Topology from an Embedded Platform Services Controller and an External Platform Services Controller in Replication

Figure 41: Deprecated Topology of an Embedded Platform Services Controller and an External Platform Services Controller in Replication




Example Upgrade Paths from vCenter Server6.7 to vCenter Server8.0

vCenter Server6.7 deployment upgrade to 8.0.


The vCenter Server example upgrade paths demonstrate vCenter Server6.7 upgrade outcomes.

The installer upgrades vCenter Server6.7 with an embedded Platform Services Controller instance to vCenter Server8.0.

Figure 42: vCenter Server 6.7 with Embedded Platform Services Controller Before and After Upgrade



Before upgrading a vCenter Server deployment with an external Platform Services Controller, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the upgrade.

Figure 43: vCenter Server 6.7 with External Platform Services Controller Before and After Upgrade



Example Migration Paths from vCenter Server for Windows to vCenter Server 8.0

You can migrate a vCenter Server for Windows instance to a vCenter Server 8.0 instance.

You can migrate a vCenter Server for Windows version 6.7 to a vCenter Server 8.0 appliance.

The vCenter Server example migration paths demonstrate supported migration outcomes.

You can migrate a vCenter Server 6.7 instance with an embedded Platform Services Controller to a 8.0. In this case the software migrates the vCenter Server instance and the embedded Platform Services Controller instance at the same time.

Figure 44: vCenter Server 6.7 with Embedded Platform Services Controller Installation Before and After Migration**Figure 45: vCenter Server 6.7 with External Platform Services Controller Installation Before and After Migration**

Before migrating a vCenter Server 6.5 or 6.7 instance with an external Platform Services Controller, you must converge the external Platform Services Controller to an embedded Platform Services Controller and then migrate to vCenter Server 8.0.

Upgrading the vCenter Server Appliance

You can upgrade vCenter Server appliance 7.0 to version 8.0. All the installation files that are necessary for the upgrade are included in the vCenter Server installer, which you can download from the VMware website.

The upgrade of the vCenter Server appliance is a migration of the old version to the new version, which includes deploying a new vCenter Server appliance of version 8.0. You can deploy the new appliance on an ESXi host 7.0 or later, or on the inventory of a vCenter Server instance 7.0 or later. You assign a temporary IP address to the new appliance to facilitate the configuration and services data migration from the old appliance to the newly deployed appliance. After the migration, the IP address and host name of the old appliance are applied to the new upgraded appliance of version 8.0. At the end of the upgrade, the temporary IP address is released and the old appliance is powered off.

Version 8.0 of vCenter Server uses an embedded PostgreSQL database. During the upgrade, you must select a storage size for the new appliance that is suitable for the database size.

vCenter Server 8.0 uses an embedded vSphere Lifecycle Manager service that allows you to perform centralized and simplified lifecycle management of clusters with ESXi 7.0 hosts. vSphere Lifecycle Manager in vSphere 8.0 includes the functionality that Update Manager provided in earlier vSphere releases for host upgrade and patching operations, and upgrade of virtual machine hardware and VMware Tools.

If you are upgrading a vCenter Server appliance or migrating a vCenter Server that use an external Update Manager instance that runs on Windows, in vSphere 8.0 the external Update Manager instance is migrated to the embedded vSphere Lifecycle Manager extension service of the new, upgraded vCenter Server appliance.

If you are upgrading a vCenter Server appliance that uses an embedded VMware Update Manager instance, in vSphere 7.0 the embedded VMware Update Manager instance is upgraded to the embedded vSphere Lifecycle Manager extension service of the new upgraded vCenter Server appliance. The embedded VMware vSphere Update Manager Extension uses the embedded PostgreSQL database. Before the upgrade, you must run the Migration Assistant on the source Update Manager instance.

For information about the software included in the vCenter Server 8.0, see *vCenter Server Installation and Setup*.

Note: vCenter Server deployments using an external Platform Services Controller will not be supported in a future vSphere release. For more information, see [Decommission the Platform Services Controller](#) and Knowledge Base article [KB 60229](#). For topologies with external Platform Services Controller instances, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the upgrade. To learn more about converging a vCenter Server deployment using an external Platform Services Controller into a vCenter Server appliance, see Knowledge Base article [KB 68137](#). After a successful upgrade, the external Platform Services Controller is powered off can be removed from your vSphere inventory.

The vCenter Server installer contains executable files for both GUI and CLI upgrades which you can use alternatively.

- The GUI upgrade is a two stage process. The first stage is a Deployment wizard that deploys the OVA file of the new appliance on the target ESXi host or vCenter Server instance. After the OVA deployment finishes, you are redirected to the second stage of the process that sets up and transfers the services and configuration data from the old appliance to the newly deployed appliance.
- The CLI upgrade method involves running a CLI command against a JSON file that you previously prepared. The CLI installer parses the configuration parameters and their values from the JSON file and generates an OVF Tool command that deploys the new appliance. The OVF Tool command also transfers services and configuration data and from the old appliance to the new appliance.

For information about the vCenter Server and Platform Services Controller appliance upgrade requirements, see [System Requirements for the vCenter Server Appliance](#).

Important: If the appliance that you are upgrading is configured in a mixed IPv4 and IPv6 environment, only the IPv4 settings are preserved. For information on the transfer of networking configuration settings for mixed mode IPv4 and IPv6 deployments, see [Mixed IPv4 and IPv6 Upgrade and Migration](#).

If you are deploying the vCenter Server appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not shown. After the upgrade, you can manually connect the appliance to the original non-ephemeral distributed virtual port group. This is not a limitation when deploying the appliance through a vCenter Server, and you can deploy to ephemeral or non-ephemeral distributed virtual port groups.

To upgrade vCenter Server appliance version 6.7 or earlier, you must first upgrade to version 7.0 and then upgrade to version 8.0. For information about upgrading vCenter Server appliance 6.7 to version 7.0, see the *VMware vSphere 7.0* documentation. For information on the upgrade compatibility of vCenter Server, see the [VMware Compatibility Guide](#).

For information about deploying the vCenter Server, see *vCenter Server Installation and Setup*.

For information about configuring the vCenter Server, see *vCenter Server Configuration*.

About the Upgrade Process of the vCenter Server Appliance

You can upgrade the vCenter Server Appliance from version 6.7 or 7.0 to version 8.0.

Important:

If your current vCenter Server Appliance is version 6.5 or earlier, you must upgrade to version 6.7 or 7.0 before upgrading to version 8.0.

Important: In environments with multiple instances of vCenter Server Appliance, you cannot perform parallel upgrades. Each vCenter Server instance must be upgraded separately. The reason for this is that VMware Directory Services (vmdir) encounters replication problems with single-sign on and certificate information.

When you run the GUI or CLI upgrade, the process includes:

1. Deploying a new vCenter Server 8.0 with a temporary network configuration. If you are upgrading a vCenter Server, you must select a deployment size for the new appliance that is suitable for your vSphere environment size. You must also select a storage size for the new appliance that is suitable for the vCenter Server database. If the source vCenter Server uses an external database, see [Determine the Oracle Database Size and the Storage Size for the New Appliance](#).
 2. Exporting the services and configuration data from the source appliance of version 6.7 or 7.0 that you want to upgrade. You must select the data types that you want to transfer to the new appliance.
If you are upgrading a vCenter Server Appliance that uses an external Update Manager instance, you must ensure that the Migration Assistant is running on the Update Manager computer. Migration Assistant facilitates the export of the Update Manager configuration and database.
 3. Transferring the exported data to the newly deployed appliance.
If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and do not appear as options during the upgrade. After the upgrade, you can manually connect the appliance to the original non-ephemeral distributed virtual port group. This process is not a limitation when deploying the appliance through a vCenter Server, and you can deploy to ephemeral or non-ephemeral distributed virtual port groups.
If you are upgrading a vCenter Server Appliance that uses a Update Manager instance, the Update Manager instance is migrated to the vSphere Lifecycle Manager of the new upgraded appliance. To learn more about vSphere Lifecycle Manager, see the *Managing Host and Cluster Lifecycle* documentation.
 4. If you are upgrading a vCenter Server Appliance that uses an external Platform Services Controller, it is converted to a vCenter Server instance with those services embedded in the appliance.
The new vCenter Server Appliance contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, and licensing. It is no longer necessary nor possible to deploy and use an external Platform Services Controller. All Platform Services Controller services are consolidated into vCenter Server during the upgrade.
 5. Powering off the source appliance. The new upgraded vCenter Server Appliance assumes the network configuration of the source appliance.
- For the new appliance requirements, see [System Requirements for the vCenter Server Appliance](#).
 - For the appliance upgrade preparation, see [Preparing to Upgrade the vCenter Server Appliance](#).
 - For the appliance upgrade procedures, see [Upgrading the vCenter Server Appliance](#).
 - For the appliance post-upgrade procedures, see [After Upgrading or Migrating vCenter Server](#).

System Requirements for the vCenter Server Appliance

You can deploy the vCenter Server appliance on an ESXi host 6.7 or later, or on a vCenter Server instance 6.7 or later. Your system must also meet specific software and hardware requirements.

When you use Fully Qualified Domain Names, verify that the client machine from which you are deploying the appliance and the network on which you are deploying the appliance use the same DNS server.

Before you deploy the appliance, synchronize the clocks of the target server and all vCenter Server instances on the vSphere network. Unsynchronized clocks might result in authentication problems and can cause the installation to fail or prevent the appliance services from starting. See [Synchronizing Clocks on the vSphere Network](#).

Hardware Requirements for the vCenter Server Appliance

When you deploy the vCenter Server appliance, you can select to deploy an appliance that is suitable for the size of your vSphere environment. The option that you select determines the number of CPUs and the amount of memory for the appliance.

The hardware requirements for a vCenter Server appliance depend on the size of your vSphere inventory.

	Number of vCPUs	Memory
Tiny environment (up to 10 hosts or 100 virtual machines)	2	14 GB
Small environment (up to 100 hosts or 1,000 virtual machines)	4	21 GB
Medium environment (up to 400 hosts or 4,000 virtual machine)	8	30 GB
Large environment (up to 1,000 hosts or 10,000 virtual machines)	16	39 GB
X-Large environment (up to 2,000 hosts or 35,000 virtual machines)	24	58 GB

Note: If you want to add an ESXi host with more than 512 LUNs and 2,048 paths to the vCenter Server inventory, you must deploy a vCenter Server appliance for a large or x-large environment.

Storage Requirements for the vCenter Server Appliance

When you deploy the vCenter Server appliance, the ESXi host or DRS cluster on which you deploy the appliance must meet minimum storage requirements. The required storage depends not only on the size of the vSphere environment and the storage size, but also on the disk provisioning mode.

The storage requirements are different for each vSphere environment size and depend on your database size requirements.

	Default Storage Size	Large Storage Size	X-Large Storage Size
Tiny environment (up to 10 hosts or 100 virtual machines)	579 GB	2019 GB	4279 GB
Small environment (up to 100 hosts or 1,000 virtual machines)	694 GB	2044 GB	4304 GB

	Default Storage Size	Large Storage Size	X-Large Storage Size
Medium environment (up to 400 hosts or 4,000 virtual machine)	908 GB	2208 GB	4468 GB
Large environment (up to 1,000 hosts or 10,000 virtual machines)	1358 GB	2258 GB	4518 GB
X-Large environment (up to 2,000 hosts or 35,000 virtual machines)	2283 GB	2383 GB	4643 GB

Note: The storage requirements include the requirements for the vSphere Lifecycle Manager that runs as a service in the vCenter Server appliance.

Related Links

[Determine the Oracle Database Size and the Storage Size for the New Appliance on page 1293](#)

Before upgrading a vCenter Server appliance or migrating a vCenter Server on Windows that uses an external Oracle database, you must determine the size of the existing database. Based on the size of the existing database, you can calculate the minimum storage size for the new vCenter Server appliance database using an embedded PostgreSQL database.

Software Requirements for the vCenter Server Appliance

The VMware vCenter Server appliance can be deployed on ESXi6.7 hosts or later, or on vCenter Server instances 6.7 or later.

You can deploy the vCenter Server appliance using the GUI or CLI installer. You run the installer from a network client machine that you use to connect to the target server and deploy the appliance on the server. You can connect directly to an ESXi6.7 host on which to deploy the appliance. You can also connect to a vCenter Server6.7 instance to deploy the appliance on an ESXi host or DRS cluster that resides in the vCenter Server inventory.

For information about the requirements for network client machine, see [System Requirements for the vCenter Server Installer](#).

Required Ports for vCenter Server

The vCenter Server system must be able to send data to every managed host and receive data from the vSphere Client. To enable migration and provisioning activities between managed hosts, the source and destination hosts must be able to receive data from each other through predetermined TCP and UDP ports.

vCenter Server is accessed through predetermined TCP and UDP ports. If you manage network components from outside a firewall, you might be required to reconfigure the firewall to allow access on the appropriate ports. For the list of all supported ports and protocols in vSphere, see the VMware Ports and Protocols Tool™ at <https://ports.broadcom.com>.

During installation, if a port is in use or is blocked using a denylist, the vCenter Server installer displays an error message. You must use another port number to proceed with the installation. There are internal ports that are used only for inter-process communication.

VMware uses designated ports for communication. Additionally, the managed hosts monitor designated ports for data from vCenter Server. If a built-in firewall exists between any of these elements, the installer opens the ports during the installation or upgrade process. For custom firewalls, you must manually open the required ports. If you have a firewall between two managed hosts and you want to perform source or target activities, such as migration or cloning, you must configure a means for the managed hosts to receive data.

To configure the vCenter Server system to use a different port to receive vSphere Client data, see the *vCenter Server and Host Management* documentation.

DNS Requirements for the vCenter Server Appliance

When you deploy the vCenter Server appliance with a static IP address, you ensure that in case of system restart, the IP address of the appliance remains the same.

Before you deploy the vCenter Server appliance with a static IP address, you must verify that this IP address has a valid internal domain name system (DNS) registration.

When you deploy the vCenter Server appliance, the installation of the web server component that supports the vSphere Client fails if the installer cannot look up the fully qualified domain name (FQDN) for the appliance from its IP address. Reverse lookup is implemented using PTR records.

If you plan to use an FQDN for the appliance system name, you must verify that the FQDN is resolvable by a DNS server, by adding forward and reverse DNS A records.

You can use the `nslookup` command to verify that the DNS reverse lookup service returns an FQDN when queried with the IP address and to verify that the FQDN is resolvable.

```
nslookup -nosearch -nodefname FQDN_or_IP_address
```

If you use DHCP instead of a static IP address for the vCenter Server appliance, verify that the appliance name is updated in the domain name service (DNS). If you can ping the appliance name, the name is updated in DNS.

Ensure that the ESXi host management interface has a valid DNS resolution from the vCenter Server and all vSphere Client instances. Ensure that the vCenter Server has a valid DNS resolution from all ESXi hosts and vSphere Client.

vSphere Client Software Requirements

Use of the vSphere Client requires a supported web browser.

VMware has tested and supports the following guest operating systems and browser versions for the vSphere Client.

Supported Guest Operating Systems

- Windows 32-bit and 64-bit
- Mac OS

Supported Browser Versions

- Google Chrome 90 or later
- Mozilla Firefox 80 or later
- Microsoft Edge 89 or later

Note: Later versions of these browsers are likely to work, but have not been tested.

Preparing to Upgrade the vCenter Server Appliance

Before you upgrade the vCenter Server appliance, you must download the vCenter Server installer ISO file and mount it to a network virtual machine or physical server from which you want to perform the upgrade.

The computer from which you upgrade the appliance must run on a Windows, Linux, or Mac operating system that meets the operating system requirements. See [System Requirements for the vCenter Server Installer](#).

Before upgrading a vCenter Server appliance, you must prepare the ESXi hosts in the inventory.

Before upgrading a vCenter Server deployment using an external Platform Services Controller, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the upgrade. See Knowledge Base article [KB 68137](#).

If the vCenter Server appliance uses an external Oracle database, you must determine the size of the existing database.

If vCenter Server appliance uses an external Update Manager instance, you must run the Migration Assistant on the computer running Update Manager.

System Requirements for the vCenter Server Installer

You can run the vCenter Server GUI or CLI installer from a network client machine that is running on a Windows, Linux, or Mac operating system of a supported version.

To ensure optimal performance of the GUI and CLI installers, use a client machine that meets the minimum hardware requirements.

Table 60: System Requirements for the GUI and CLI Installers

Operating System	Supported Versions	Minimum Hardware Configuration for Optimal Performance
Windows	- Windows 10, 11 - Windows 2016 x64 bit - Windows 2019 x64 bit - Windows 2022 x64 bit	4 GB RAM, 2 CPU having 4 cores with 2.3 GHz, 32 GB hard disk, 1 NIC
Linux	- SUSE 15 - Ubuntu 18.04, 20.04, 21.10	4 GB RAM, 1 CPU having 2 cores with 2.3 GHz, 16 GB hard disk, 1 NIC Note: The CLI installer requires 64-bit OS.
Mac	- macOS 10.15, 11, 12 - macOS Catalina, Big Sur, Monterey	8 GB RAM, 1 CPU having 4 cores with 2.4 GHz, 150 GB hard disk, 1 NIC

Note: For client machines that run on Mac 10.15 or later, concurrent GUI deployments of multiple appliances are unsupported. You must deploy the appliances in a sequence.

Note: Visual C++ redistributable libraries need to be installed to run the CLI installer on versions of Windows older than Windows 10. The Microsoft installers for these libraries are located in the `vcsa-cli-installer/win32/vcredist` directory.

Note: Deploying the vCenter Server appliance with the GUI requires a minimum resolution of 1024x768 to properly display. Lower resolutions can truncate the UI elements.

Download and Mount the vCenter Server Installer

The vCenter Server appliance ISO image contains GUI and CLI installers for the vCenter Server appliance.

- Create an account at <https://support.broadcom.com/web/ecx>.
- Verify that your client machine meets the system requirements for the vCenter Server installer. See [System Requirements for the vCenter Server Installer](#).

With the GUI and CLI executable files that are included in the vCenter Server installer, you can:

- Deploy the vCenter Server appliance.
- Upgrade the vCenter Server appliance.
- Converge older versions of vCenter Server with an external Platform Services Controller to the current version of vCenter Server.

- Restore a vCenter Server appliance from a file-based backup.
1. Log in to Broadcom Support Portal.
 2. Download the vCenter Server appliance ISO image. See the Knowledge Base article <https://knowledge.broadcom.com/external/article/366685/vmware-vsphere-downloads-vmware-converte.html>.
 3. Confirm that the md5sum is correct by using an MD5 checksum tool.
 4. Mount the ISO image to the client machine from which you want to deploy, upgrade, migrate, or restore the appliance.

Note: ISO mounting software that does not allow more than eight directory levels, for example, MagicISO Maker on Windows, is unsupported.

For Linux OS and Mac OS, Archive Manager is unsupported.

For Mac OS, you can use DiskImageMounter.

For Ubuntu 14.04, you can use Disk Image Mounter.

For SUSE 12 OS, you can use the terminal.

```
$ sudo mkdir mount_dir
$ sudo mount -o loop VMware-vCSA-all-version_number-build_number.iso mount_dir
```

Important: Due to a security change in MacOS Catalina, you must modify the security settings on your computer until the vCenter Server deployment completes. If you attempt to run the installer under MacOS Catalina without modifying the security settings, the vCenter Server installer reports the error: `ovftool cannot be opened because the developer cannot be verified`. For more information, see [KB 79416](#).

Open the `readme.txt` file and review the information about the other files and directories in the vCenter Server appliance ISO image.

Synchronizing Clocks on the vSphere Network

Verify that all components on the vSphere network have their clocks synchronized. If the clocks on the physical machines in your vSphere network are not synchronized, SSL certificates and SAML tokens, which are time-sensitive, might not be recognized as valid in communications between network machines.

Unsynchronized clocks can result in authentication problems, which can cause the installation to fail or prevent the vCenter Server `vmware-vpxd` service from starting.

Time inconsistencies in vSphere can cause the first boot of a component in your environment to fail at different services depending on where in the environment time is not accurate and when the time is synchronized. Problems most commonly occur when the target ESXi host for the destination vCenter Server is not synchronized with NTP or PTP. Similarly, issues can arise if the destination vCenter Server migrates to an ESXi host set to a different time due to fully automated DRS.

To avoid time synchronization issues, ensure that the following is correct before installing, migrating, or upgrading a vCenter Server instance.

- The target ESXi host where the destination vCenter Server is to be deployed is synchronized to NTP or PTP.
- The ESXi host running the source vCenter Server is synchronized to NTP or PTP.
- When upgrading or migrating from vSphere 6.7 to vSphere 8.0, if the vCenter Server appliance is connected to an external Platform Services Controller, ensure the ESXi host running the external Platform Services Controller is synchronized to NTP or PTP.
- If you are upgrading or migrating from vSphere 6.7 to vSphere 8.0, verify that the source vCenter Server or vCenter Server appliance and external Platform Services Controller have the correct time.

Verify that any Windows host machine on which vCenter Server runs is synchronized with the Network Time Server (NTP) server.

To synchronize ESXi clocks with an NTP or a PTP server, you can use the VMware Host Client. For information about editing the time configuration of an ESXi host, see topic *Edit the Time Configuration of an ESXi Host in the VMware Host Client* in the *vSphere Single Host Management - VMware Host Client* documentation.

To learn how to change time synchronization settings for vCenter Server, see topic *Configure the System Time Zone and Time Synchronization Settings* in the *vCenter Server Configuration* documentation.

To learn how to edit the time configuration for a host by using the vSphere Client, see topic *Editing the Time Configuration Settings of a Host* in the *vCenter Server and Host Management* documentation.

Transferring Data from an Existing vCenter Server Appliance

Typically, data transfer happens as part of the upgrade downtime. However, in some scenarios you can choose to transfer a part of the data after upgrade or migration to minimize the downtime.

Transfer of data after upgrade or migration is supported in the following scenarios:

- Migrating vCenter Server for Windows to vCenter Server Appliance
- Upgrading vCenter Server appliance with external database

Note: You can only transfer data from external database to the embedded PostgreSQL database in the background after upgrade or migration. Transferring data from an embedded PostgreSQL database to another embedded PostgreSQL database after upgrade or migration is not supported.

You can transfer the minimum amount of data necessary to be operational during the upgrade. Later, you can transfer the remaining data from the external database to the embedded PostgreSQL database in the background, preserving your deployment's historical event and performance data.

You can transfer historical data after the upgrade or migration from the following vCenter Server versions using external databases to the embedded PostgreSQL database within the vCenter Server appliance.

Table 61: vCenter Server Database Support for Data Transfer After Upgrade or Migration

vCenter Server Versions	Databases	Support for Data Transfer After Upgrade or Migration
vCenter Server appliance 6.7	Embedded PostgreSQL database	Not Supported
vCenter Server for Windows 6.7	External Oracle or MS SQL database	Supported
	Embedded PostgreSQL database	Not supported
vCenter Server appliance 7.0	Embedded PostgreSQL database	Not supported

You can migrate the following types of data from an external database to the embedded PostgreSQL database in use by version 8.0 of the vCenter Server appliance.

Configuration data

The minimum type of data you must transfer from your existing vCenter Server appliance external database to upgrade or migrate to a functioning vCenter Server appliance. This data can be transferred in a reasonably short amount of time, minimizing downtime during the upgrade to vCenter Server 8.0.

Configuration and historical data

In addition to the configuration data, you can choose to transfer historical data such as usage statistics, events, and tasks. The vCenter Server installer gives you the option to copy both the configuration and historical data at one time, or to copy only the configuration data during the upgrade process. You can then choose to copy the historical data in the background after you start your new vCenter Server.

Configuration, historical, and performance metrics data

Attention: If you choose to copy all data at one time, you cannot start vCenter Server until all the data has been transferred to the embedded PostgreSQL database within vCenter Server8.0.

The largest amount of data you can choose to transfer to vCenter Server. You can copy only the configuration data during the upgrade process, and then copy the historical and performance data in the background after the upgrade or migration to vCenter Server8.0 finishes.

Monitor and Manage Historical Data Migration

You can monitor and manage the background migration of historical data using the vCenter Server Management Interface. You can perform the following data management tasks:

- Monitor the progress of the data migration.
- Pause the data import.
- Cancel the data import.

Important: If you cancel the data import, the historical data is not imported to the embedded PostgreSQL database. If you cancel the data import operation, you cannot recover the data. If you do cancel the data import, and want to import the historical data later, you must restart the upgrade or migration process from Stage 1 of the GUI installer.

See [Monitor and Manage Historical Data Migration](#).

Prepare ESXi Hosts for vCenter Server Appliance Upgrade

Before upgrading to vCenter Server8.0, you must prepare your ESXi hosts.

To upgrade vCenter Server appliance, your ESXi hosts must meet the requirements for upgrade.

- ESXi hosts must be at version 6.7 or later. For information on ESXi compatibility, see the [VMware Compatibility Guide](#).
 - Your source and target ESXi hosts must not be in lockdown or maintenance mode, and not part of fully automated DRS clusters.
 - Review your environment's Certificate Authority (CA) signed SSL certificates. To learn about manually reviewing CA signed SSL certificates in a vSphere environment, see Knowledge Base article [KB 2111411](#).
1. If you have Custom or Thumbprint certificates, see [ESXi Host Upgrades and Certificates](#) to determine your preparatory steps.
 2. If you have vSphere HA clusters, SSL certificate checking must be enabled.
If certificate checking is not enabled when you upgrade ESXi hosts, vSphere HA fails to configure on the hosts.
 - a) Select the vCenter Server appliance instance in the vSphere Client inventory tree.
 - b) Select the **Configure** tab, and then under **Settings** select **General**.
 - c) Click **Edit**.
 - d) Select **SSL settings**.

Verify that **vCenter Server requires verified host SSL certificates** is selected.

Your ESXi hosts are ready for the vCenter Server appliance upgrade.

ESXi Host Upgrades and Certificates

If you upgrade an ESXi host to ESXi6.7 or later, the upgrade process replaces the self-signed (thumbprint) certificates with VMCA-signed certificates. If the ESXi host uses custom certificates, the upgrade process retains those certificates even if those certificates are expired or invalid.

The recommended upgrade workflow depends on the current certificates.

Host Provisioned with Thumbprint Certificates

If your host is currently using thumbprint certificates, it is automatically assigned VMCA certificates as part of the upgrade process.

Note: You cannot provision legacy hosts with VMCA certificates. You must upgrade those hosts to ESXi6.7 or later.

Host Provisioned with Custom Certificates

If your host is provisioned with custom certificates, usually third-party CA-signed certificates, those certificates remain in place during upgrade. Change the certificate mode to **Custom** to ensure that the certificates are not replaced accidentally during a certificate refresh later.

Note: If your environment is in VMCA mode, and you refresh the certificates from the vSphere Client, any existing certificates are replaced with certificates that are signed by VMCA.

Going forward, vCenter Server monitors the certificates and displays information, for example, about certificate expiration, in the vSphere Client.

Hosts Provisioned with Auto Deploy

Hosts that are being provisioned by Auto Deploy are always assigned new certificates when they are first booted with ESXi6.7 or later software. When you upgrade a host that is provisioned by Auto Deploy, the Auto Deploy server generates a certificate signing request (CSR) for the host and submits it to VMCA. VMCA stores the signed certificate for the host. When the Auto Deploy server provisions the host, it retrieves the certificate from VMCA and includes it as part of the provisioning process.

You can use Auto Deploy with custom certificates.

Change the ESXi Certificate Mode

Use VMware Certificate Authority (VMCA) to provision the certificates of the ESXi hosts in your environment unless your corporate policy requires that you use custom certificates. To use custom certificates with a different root CA, edit the advanced vCenter Server setting, `vpzd.certmgmt.mode`. After the change, the hosts are no longer automatically provisioned with VMCA certificates when you refresh the certificates. You are responsible for the certificate management in your environment.

You can use the advanced vCenter Server settings to change to thumbprint mode or to custom CA mode. Use thumbprint mode only as a fallback option.

1. In the vSphere Client, select the vCenter Server system that manages the hosts.
2. Click **Configure**, and under Settings, click **Advanced Settings**.
3. Click **Edit Settings**.
4. Click the **Filter** icon in the Name column, and in the Filter box, enter `vpzd.certmgmt` to display only certificate management parameters.

Note: The available options are **vmca**, **custom**, and **thumbprint**.

5. Change the value of `vpd.certmgmt.mode` to **custom** if you intend to manage your own certificates, and to **thumbprint** if you temporarily want to use thumbprint mode, and click **Save**.

Determine the Oracle Database Size and the Storage Size for the New Appliance

Before upgrading a vCenter Server appliance or migrating a vCenter Server on Windows that uses an external Oracle database, you must determine the size of the existing database. Based on the size of the existing database, you can calculate the minimum storage size for the new vCenter Server appliance database using an embedded PostgreSQL database.

You must have the vCenter Server database login credentials.

You run scripts to determine the Oracle core table size, the events and tasks table size, and the statistics table size. The Oracle core table corresponds to the database (`/storage/db`) partition of the PostgreSQL database. The Oracle events and tasks and statistics tables correspond to the statistics, events, alarms, and tasks (`/storage/seat`) partition of the PostgreSQL database.

During the upgrade of the appliance, you must select a storage size for the new appliance that is at least twice the size of the Oracle tables size.

During the upgrade of the appliance, you can select the types of data to transfer to the new appliance. To minimize upgrade time and storage requirements for the new appliance, you can choose to transfer only the configuration data.

1. Log in to a SQL*Plus session with the vCenter Server database user.
2. Determine the core table size by running the following script.

```
SELECT ROUND(SUM(s.bytes)/(1024*1024)) SIZE_MB
FROM   user_segments s
WHERE  (s.segment_name,s.segment_type)
        IN (SELECT seg_name, seg_type FROM
              (SELECT t.table_name seg_name, t.table_name tname,
                'TABLE' seg_type
               FROM   user_tables t
              UNION
               SELECT i.index_name, i.table_name,
                'INDEX'
               FROM   user_indexes i
              ) ti
        WHERE (ti.tname LIKE 'VPX_%'
              OR ti.tname LIKE 'CL_%'
              OR ti.tname LIKE 'VDC_%')
        AND ti.tname NOT LIKE 'VPX_SAMPLE_TIME%'
        AND ti.tname NOT LIKE 'VPX_HIST_STAT%'
        AND ti.tname NOT LIKE 'VPX_TOPN%'
        AND ti.tname NOT LIKE 'VPX_SDRS_STATS_VM%'
        AND ti.tname NOT LIKE 'VPX_SDRS_STATS_DATASTORE%'
        AND ti.tname NOT LIKE 'VPX_TASK%'
        AND ti.tname NOT LIKE 'VPX_EVENT%'
        AND ti.tname NOT LIKE 'VPX_PROPERTY_BULLETIN%');
```

The script returns the database storage size in MB.

3. Determine the events and tasks table size by running the following script.

```
SELECT ROUND(SUM(s.bytes)/(1024*1024)) SIZE_MB
FROM   user_segments s
WHERE  (s.segment_name,s.segment_type)
```

```

IN (SELECT seg_name, seg_type FROM
    (SELECT t.table_name seg_name, t.table_name tname,
        'TABLE' seg_type
    FROM    user_tables t
    UNION
    SELECT i.index_name, i.table_name,
        'INDEX'
    FROM    user_indexes i
    ) ti
WHERE
    ti.tname LIKE 'VPX_TASK%'
    OR ti.tname LIKE 'VPX_EVENT%');

```

The script returns the events and tasks storage size in MB.

4. Determine the statistics table size by running the following script.

```

SELECT ROUND(SUM(s.bytes)/(1024*1024)) SIZE_MB
FROM    user_segments s
WHERE   (s.segment_name,s.segment_type)
    IN (SELECT seg_name, seg_type FROM
        (SELECT t.table_name seg_name, t.table_name tname,
            'TABLE' seg_type
        FROM    user_tables t
        UNION
        SELECT i.index_name, i.table_name,
            'INDEX'
        FROM    user_indexes i
        ) ti
    WHERE
        ti.tname LIKE 'VPX_SAMPLE_TIME%'
        OR ti.tname LIKE 'VPX_TOPN%'
        OR ti.tname LIKE 'VPX_TASK%'
        OR ti.tname LIKE 'VPX_EVENT%'
        OR ti.tname LIKE 'VPX_HIST_STAT%');

```

The script returns the statistics storage size in MB.

5. Calculate the minimum storage size for the new appliance that you are going to deploy during the upgrade.

- a) The size of the database (/storage/db) partition of the embedded PostgreSQL database must be at least twice the size of the Oracle core table returned in [2](#).
- b) The size of the statistics, events, alarms, and tasks (/storage/seat) partition of the embedded PostgreSQL database must be at least twice the sum of the sizes of the Oracle events and tasks and statistics tables returned in [3](#) and [4](#).

For example, if the Oracle core table is 100 MB, the events and tasks table is 1,000 MB, and the statistics table is 2,000 MB, then the Postgres /storage/db partition must be at least 200 MB and the /storage/seat partition must be at least 6,000 MB.

Related Links

[Storage Requirements for the vCenter Server Appliance on page 1215](#)

When you deploy the vCenter Server appliance, the ESXi host or DRS cluster on which you deploy the appliance must meet minimum storage requirements. The required storage depends not only on the size of the vSphere environment and the storage size, but also on the disk provisioning mode.

Download and Run VMware Migration Assistant on the Source Update Manager Machine

During the upgrade of a vCenter Server appliance that uses an external Update Manager, the Migration Assistant must be running on the source Update Manager machine. This procedure describes how to download and run the Migration Assistant manually before the upgrade.

- [Download and Mount the vCenter Server Installer](#).
- Log in to the source Update Manager machine as an administrator.

The Migration Assistant facilitates the migration of the Update Manager server and database to the new upgraded vCenter Server appliance. The Migration Assistant uses port 9123 by default. If port 9123 is used by another service on your Update Manager machine, the Migration Assistant automatically finds a different free port to use.

Alternatively, if you plan to upgrade the vCenter Server appliance by using the CLI installer, you can add the `source.vum` section and `run.migration.assistant` subsection to your JSON template. For information about the CLI upgrade configuration parameters, see [Upgrade Configuration Parameters for CLI Upgrade of vCenter Server Appliance](#).

1. From the vCenter Server appliance installer package, copy the `migration-assistant` directory to the source Update Manager machine.
2. From the `migration-assistant` directory, double-click `VMware-Migration-Assistant.exe` and provide the vCenter Single Sign-On administrator password.
3. Leave the Migration Assistant window open until the upgrade of the vCenter Server appliance finishes.

When the pre-checks are finished and any errors are addressed, your source Update Manager system is ready for the upgrade.

CAUTION: Closing the Migration Assistant window causes the upgrade process to stop.

Prerequisites for Upgrading the vCenter Server Appliance

To ensure a successful upgrade of the vCenter Server appliance, you must perform the following required tasks and pre-checks before beginning the upgrade process.

General Prerequisites

- [Download and Mount the vCenter Server Installer](#).
- Verify that the clocks of all computers on the vSphere network are synchronized. See [Synchronizing Clocks on the vSphere Network](#).

Target System Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [System Requirements for the vCenter Server Appliance](#).
- If you plan to deploy the new appliance on an ESXi host, verify that the target ESXi host is not in lockdown or maintenance mode.
- If you plan to deploy the new appliance on an ESXi host managed by vCenter Server, review the Distributed Resource Scheduler (DRS) settings on the cluster. Ensure that the DRS settings for the cluster on which the target ESXi host resides are not configured for Full Automation. Set the Automation Level to Manual or Partially Automated. This Automation Level ensures that the target ESXi host does not reboot during the upgrade process.
- If you plan to deploy the new appliance on a DRS cluster of the inventory of a vCenter Server instance, verify that the cluster contains at least one ESXi host that is not in lockdown or maintenance mode.
- If you plan to upgrade a vCenter Server instance in an environment that has vSphere with Tanzu enabled, verify that you deploy the new appliance on a cluster where the Supervisor is not running. Ensure that the DRS settings for the cluster on which the target ESXi host resides are not configured for Full Automation.

Source System Prerequisites

- If the source appliance is on a DRS cluster of the inventory of a vCenter Server instance and during upgrade:
 - If you plan to use the ESXi host where the appliance resides, review the Distributed Resource Scheduler (DRS) settings on the cluster. Ensure that the DRS settings of the cluster is set to manual or partially automated so that the VMs of the source system do not move during the upgrade.
 - If you plan to use the vCenter Server instance where the appliance resides, you need not modify the DRS settings of the cluster.
- Verify that port 22 is open on the vCenter Server appliance that you want to upgrade. The upgrade process establishes an inbound SSH connection to download the exported data from the source vCenter Server appliance.
- If you are upgrading a vCenter Server appliance that is configured with an external Update Manager, run the Migration Assistant on the source Update Manager computer.
 For a GUI upgrade, you must run Migration Assistant manually. See [Download and Run VMware Migration Assistant on the Source Windows Machine](#).
 For a CLI upgrade, you can run Migration Assistant manually or automatically. To run Migration Assistant automatically, add the `source.vum` section and `run.migration.assistant` subsection to your JSON template. See [Upgrade Configuration Parameters for CLI Upgrade of vCenter Server Appliance](#).
- Verify that port 443 is open on the source ESXi host on which the appliance that you want to upgrade resides. The upgrade process establishes an HTTPS connection to the source ESXi host to verify that the source appliance is ready for upgrade and to set up an SSH connection between the new and the existing appliance.
Note: When upgrading, the temporary vCenter Server instance requires the same access rights as the permanent vCenter Server instance to port 443. Ensure that any firewalls in your environment allow both the temporary and permanent vCenter Server instances to access port 443.
- Verify that you have sufficient free disk space on the appliance that you want to upgrade so that you can accommodate the data for the upgrade.
- Create an image-based backup (snapshot) of the vCenter Server appliance you are upgrading as a precaution in case there is a failure during the upgrade process. If you are upgrading a vCenter Server appliance with an external Platform Services Controller, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the upgrade. See Knowledge Base article [KB 68137](#).
Important: To take a pre-upgrade image-based backup in an Enhanced Linked Mode environment, power off all the vCenter Server and Platform Services Controller nodes in your environment, and back up each node. After you have taken backups of all the nodes, you can restart them and proceed with the upgrade procedure.
 If the upgrade fails, delete the newly deployed vCenter Server appliance, and restore the vCenter Server and Platform Services Controller nodes from their respective backups. You must restore all the nodes in the environment from their backups. Failing to do so will cause the replication partners to be out of synchronization with the restored node.
 To learn about image-based back, see "Image-Based Backup and Restore of a vCenter Server Environment" in *vCenter Server Installation and Setup*.
- If you use an external database, determine the database size and the minimum storage size for the new appliance. See [Determine the Oracle Database Size and the Storage Size for the New Appliance](#).
- If your deployment uses an external database, you must back up the external vCenter Server appliance database.
- If you plan to upgrade a vCenter Server instance in an environment that has vSphere with Tanzu enabled, ensure that the DRS settings of the cluster where the vCenter Server appliance resides is set to manual or partially automated so that the VMs of the source system do not move during the upgrade.

Network Prerequisites

- Verify that the new appliance can connect to the source ESXi host and vCenter Server instance on which resides the appliance that you want to upgrade.

- If you plan to assign a static IP address and an FQDN as a system name in the temporary network settings of the appliance, verify that you have configured the forward and reverse DNS records for the IP address.
- If you plan to assign a DHCP IP address in the temporary network settings of the new vCenter Server appliance, verify that the ESXi host on which you want to deploy the new appliance is in the same subnet as the ESXi host on which the existing vCenter Server appliance runs.
- If you plan to assign a DHCP IPv4 address in the temporary network settings of the new vCenter Server appliance, verify that the ESXi host on which you want to deploy the new appliance is connected to at least one network that is associated with a port group which accepts MAC address changes. Consider the default security policy of a distributed virtual switch, which is to reject MAC address changes. For information about how to configure the security policy for a switch or port group, see *vSphere Networking*.

GUI Upgrade of the vCenter Server Appliance

You can use the GUI installer to perform an interactive upgrade of a vCenter Server appliance.

To perform the GUI upgrade, download the vCenter Server appliance installer on a network client machine, run the Upgrade Wizard from the client machine, and provide the inputs for the deployment and setup of the new upgraded appliance.

Note: When you upgrade vCenter Server 6.7 or 7.0 with an external Platform Services Controller to vCenter Server 8.0, the upgrade process converges the external Platform Services Controller into the new vCenter Server appliance. The new vCenter Server contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, and licensing. It is no longer necessary nor possible to upgrade and use an external Platform Services Controller.

After successfully upgrading your environment to vCenter Server 8.0, the pre-existing Platform Services Controller will be powered off, and you can remove it from your vSphere inventory. See [Decommission the Platform Services Controller](#).

The GUI upgrade process includes a series of two stages.

Figure 46: Stage 1 - OVA Deployment



The first stage walks you through the deployment wizard to get the deployment type of the source appliance that you want to upgrade and configure the new appliance settings. During this stage, you deploy the new appliance with temporary network settings. This stage finishes the deployment of the OVA file on the target server with the same deployment type as the source appliance and the appliance settings that you provide.

As an alternative to performing the first stage of the upgrade with the GUI installer, you can deploy the OVA file of the new vCenter Server appliance using the vSphere Client or VMware Host Client. To deploy the OVA file on an ESXi host or vCenter Server instance 6.7 or later, you can also use the vSphere Client. After the OVA deployment, you must log in to the appliance management interface of the newly deployed appliance to proceed with the second stage of the upgrade process.

Figure 47: Stage 2 - Appliance Setup



The second stage walks you through the setup wizard to select the data types to transfer from the old to the new appliance. The new appliance uses the temporary network settings until the data transfer finishes. After the data transfer finishes, the new appliance assumes the network settings of the old appliance. This stage finishes the data transfer, starts the services of the new upgraded appliance, and powers off the old appliance.

As an alternative to performing the second stage of the upgrade with the GUI installer, you can log in to the vCenter Server Management Interface of the newly deployed appliance, https://FQDN_or_IP_address:5480.

Required Information for Upgrading a vCenter Server Appliance 7.0 or 8.0

The GUI upgrade wizard prompts you for information about the vCenter Server appliance that you want to upgrade, deployment information for the new vCenter Server appliance 8.0. It is a best practice to keep a record of the values that you enter.

You can use this worksheet to record the information that you need for upgrading a vCenter Server appliance with an embedded Platform Services Controller or a vCenter Server appliance with an external Platform Services Controller.

Table 62: Required Information During Stage 1 of the Upgrade

Required for Upgrade of	Required Information	Default	Your Entry
All deployment types	FQDN or IP address of the source appliance that you want to upgrade.	-	
	HTTPS port of the source appliance.	443	
	vCenter Single Sign-On administrator user name of the source appliance. Important: The user must be administrator@ <i>your_domain_name</i> .	administrator@vsphere.local	
	Password of the vCenter Single Sign-On administrator user.	-	
	Password of the root user of the source appliance	-	
All deployment types	FQDN or IP address of the source server on which resides that appliance that you want to upgrade. The source server can be either an ESXi host or a vCenter Server instance. Note: The source server cannot be the vCenter Server appliance that you want to upgrade. In such cases, use the source ESXi host.	-	
	HTTPS port of the source server.	443	
	User name with administrative privileges on the source server. - If your source server is an ESXi host, use root. - If your source server is a vCenter Server instance, use <i>user_name@your_domain_name</i>, for example, administrator@vsphere.local.	-	
	Password of the user with administrative privileges on the source server.	-	
All deployment types	FQDN or IP address of the target server on which you want to deploy the new appliance. The target server can be either an ESXi host or a vCenter Server instance.	-	

Required for Upgrade of	Required Information	Default	Your Entry
	Note: The target server cannot be the vCenter Server appliance that you want to upgrade. In such cases, use an ESXi host as a target server.		
	HTTPS port of the target server.	443	
	User name with administrative privileges on the target server - If your target server is an ESXi host, use root. - If your target server is a vCenter Server instance, use <code>user_name@your_domain_name</code>, for example, <code>administrator@vsphere.local</code>.	-	
	Password of the user with administrative privileges on the target server.	-	
All deployment types Only if your target server is a vCenter Server instance.	Data center from the vCenter Server inventory on which you want to deploy the new appliance. Optionally you can provide a data center folder.	-	
	ESXi host or DRS cluster from the data center inventory on which you want to deploy the new appliance .	-	
All deployment types	The virtual machine name for the new appliance. - Must not contain a percent sign (%), backslash (\), or forward slash (/). - Must be no more than 80 characters in length.	VMware vCenter Server Appliance	
All deployment types	Password for the root user of the appliance operating system. - Must contain only the lower ASCII character set without spaces. - Must be at least 8 characters, but no more than 20 characters in length. - Must contain at least one uppercase letter. - Must contain at least one lowercase letter. - Must contain at least one number. - Must contain at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!).	-	
- vCenter Server appliance 6.7 with an external Platform Services Controller - vCenter Server appliance 6.7 with an embedded Platform Services Controller - vCenter Server appliance 7.0	Deployment size of the new vCenter Server appliance for your vSphere environment. Tiny Deploys an appliance with 2 CPUs and 12 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines. Small Deploys an appliance with 4 CPUs and 19 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines.	The size of the source vCenter Server appliance and Platform Services Controller determines the default deployment size. Using this information, the upgrade installer calculates the best size to use for your environment, which	

Required for Upgrade of	Required Information	Default	Your Entry
	- Medium Deploys an appliance with 8 CPUs and 28 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines. - Large Deploys an appliance with 16 CPUs and 37 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines. - X-Large Deploys an appliance with 24 CPUs and 56 GB of memory. Suitable for environments with up to 2,000 hosts or 35,000 virtual machines.	might be the same as your existing deployment, or calculated to the next largest size.	
vCenter Server appliance 6.7 with an external Platform Services Controller	Storage size of the new vCenter Server appliance for your vSphere environment. Increase the default storage size if you want larger volume for SEAT data (stats, events, alarms, and tasks). Note: Consider the database size of the appliance that you want to upgrade and the types of the data that you want transfer to the new appliance. For an external database, see Determine the Oracle Database Size and the Storage Size for the New Appliance. - Default For tiny deployment size, deploys the appliance with 415 GB of storage. For small deployment size, deploys the appliance with 480 GB of storage. For medium deployment size, deploys the appliance with 700 GB of storage. For large deployment size, deploys the appliance with 1065 GB of storage. For x-large deployment size, deploys the appliance with 1805 GB of storage. - Large For tiny deployment size, deploys the appliance with 1490 GB of storage. For small deployment size, deploys the appliance with 1535 GB of storage. For medium deployment size, deploys the appliance with 1700 GB of storage. For large deployment size, deploys the appliance with 1765 GB of storage. For x-large deployment size, deploys the appliance with 1905 GB of storage.	Default Note: The sizing algorithm in use by the upgrade installer might select a larger storage size for your environment. Items that might affect the storage size selected by the installer include modifications to the vCenter Server appliance disks (for example, changing the size of the logging partition), or databases having a database table that the installer determines to be exceptionally large and requiring additional hard disk space.	

Required for Upgrade of	Required Information	Default	Your Entry
	X-Large For tiny deployment size, deploys the appliance with 3245 GB of storage. For small deployment size, deploys the appliance with 3295 GB of storage. For medium deployment size, deploys the appliance with 3460 GB of storage. For large deployment size, deploys the appliance with 3525 GB of storage. For x-large deployment size, deploys the appliance with 3665 GB of storage.		
All deployment types	Name of the datastore on which you want to store the configuration files and virtual disks of the new appliance. Note: The installer displays a list of datastores that are accessible from your target server.	-	
	Enable or deactivate Thin Disk Mode .	Disabled	
All deployment types	Name of the network to which to connect the new appliance. Note: The installer displays a drop-down menu with networks that depend on the network settings of your target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are not supported and are not displayed in the drop-down menu. The network must be accessible from the source server on which resides that appliance that you want to upgrade. The network must be accessible from the physical client machine from which you perform the deployment.	-	
	IP version for the appliance temporary address Can be either IPv4 or IPv6.	IPv4	
	IP assignment for the appliance temporary address Can be either static or DHCP.	static	
All deployment types Only if you use a static assignment for the temporary IP address.	Temporary system name (FQDN or IP address) The system name is used for managing the local system. The system name must be FQDN. If a DNS server is not available, provide a static IP address.	-	
	Temporary IP address	-	
	For IPv4 version, a subnet mask as a dot decimal notation or a network prefix as an integer between 0 and 32. For IPv6 version, a network prefix as an integer between 0 and 128.	-	
	Default gateway.	-	
	DNS servers separated by commas.	-	
All deployment types	Temporary system name (FQDN).	-	

Required for Upgrade of	Required Information	Default	Your Entry
Only if you use a DHCP assignment with IPv4 version for the temporary IP address and a DDNS server is available in your environment.			

Table 63: Required Information During Stage 2 of the Upgrade

Required for	Required Information	Default	Your Entry
- vCenter Server appliance 7.0 - vCenter Server appliance 6.7 with an embedded or external Platform Services Controller	Data types to transfer from the old appliance to the new appliance. In addition to the configuration data, you can transfer the events, tasks, and, performance metrics. Note: For minimum upgrade time and storage requirement for the new appliance, select to transfer only the configuration data.	-	
- vCenter Server appliance 7.0 - vCenter Server appliance 6.7 with an embedded Platform Services Controller	Join or do not participate in the VMware Customer Experience Improvement Program (CEIP). For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i>.	Join the CEIP.	

Upgrade a vCenter Server Appliance 6.7 with an Embedded Platform Services Controller or 7.0 by Using the GUI

You can use the GUI installer to perform an interactive upgrade of a vCenter Server appliance 6.7 that uses an embedded Platform Services Controller or 7.0 to vCenter Server appliance 8.0. You must run the GUI upgrade from a Windows, Linux, or Mac OS X machine that is in the same network as the appliance that you want to upgrade.

- See [Prerequisites for Upgrading the vCenter Server Appliance](#).
- See [Required Information for Upgrading a vCenter Server Appliance 7.0 or 8.0](#).

Stage 1 - Deploy the OVA File of the New vCenter Server Appliance

With stage 1 of the upgrade process, you deploy the OVA file, which is included in the vCenter Server appliance installer, for the new vCenter Server appliance.

- In the vCenter Server installer, navigate to the `vcasa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.

2. On the Home page, click **Upgrade**.
3. Review the Introduction page to understand the upgrade process and click **Next**.
4. Read and accept the license agreement, and click **Next**.
5. Connect to the source appliance that you want to upgrade.
 - a) Enter the information about the source vCenter Server appliance that you want to upgrade, and click **Connect to Source**.

Option	Action
Appliance FQDN or IP address	Enter the IP address or FQDN of the vCenter Server appliance that you want to upgrade.
Appliance HTTPS port	If the source appliance uses a custom HTTPS port, change the default value to that of the custom port. The default port value is 443. Custom port values are supported beginning with vCenter Server appliance version 6.5 Update 2. If you are upgrading from earlier versions, you cannot specify a custom port.

- b) Enter the information about the vCenter Single Sign-On administrator and root user.

Option	Action
SSO user name	Enter the vCenter Single Sign-On administrator user name. Important: The user must be administrator@ <i>your_domain_name</i> .
SSO password	Enter the password of the vCenter Single Sign-On administrator.
Appliance (OS) root password	Enter the password of the root user.

- c) Enter the information about the source ESXi host or vCenter Server instance on which resides the vCenter Server appliance that you want to upgrade and click **Next**.

Option	Description
Source server or host name	IP address or FQDN of the source ESXi host or vCenter Server instance on which the vCenter Server appliance that you want to upgrade resides. Note: The source vCenter Server instance cannot be the vCenter Server appliance that you want to upgrade. In such cases, use the source ESXi host.
HTTPS port	If the ESXi host or vCenter Server instance uses a custom HTTPS port, change the default value. The default value is 443.
User name	User name of a user with administrative privileges on the ESXi host or vCenter Server instance.
Password	Password of the user with administrative privileges on the ESXi host or vCenter Server instance.

6. Verify that the certificate warning displays the SHA1 thumbprints of the SSL certificates that are installed on the source appliance and its source server, and click **Yes** to accept the certificate thumbprints.
7. Connect to the target server on which you want to deploy the new vCenter Server appliance.

Option	Steps
You can connect to an ESXi host on which to deploy the new appliance.	1. Enter the FQDN or IP address of the ESXi host. 2. Enter the HTTPS port of the ESXi host. 3. Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the new appliance. Note: The target server cannot be the vCenter Server appliance that you want to upgrade. In such cases, use an ESXi host as a target server.	1. Enter the FQDN or IP address of the vCenter Server instance. 2. Enter the HTTPS port of the vCenter Server instance. 3. Enter the user name and password of a vCenter Single Sign-On user with administrative privileges on the vCenter Server instance, for example, the administrator@your_domain_name user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes. 6. Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next. Note: You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode. 7. Select the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next.

8. On the Set up target appliance VM page, enter a name for the new vCenter Server appliance, set the password for the root user, and click **Next**.

The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.

The password must contain only lower ASCII characters without spaces, at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

Note: The root password of the old appliance is not transferred to the new upgraded appliance.

9. Select the deployment size for the new vCenter Server appliance for your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 vCPUs and 14 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 vCPUs and 21 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 vCPUs and 30 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines

Deployment Size Option	Description
Large	Deploys an appliance with 16 vCPUs and 39 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 vCPUs and 58 GB of memory. Suitable for environments with up to 2,500 hosts or 45,000 virtual machines

10. Select the storage size for the new vCenter Server appliance and click **Next**.

Important: You must consider the storage size of the appliance that you are upgrading and the database size if external.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 579 GB of storage.	Deploys an appliance with 694 GB of storage.	Deploys an appliance with 908 GB of storage.	Deploys an appliance with 1358 GB of storage.	Deploys an appliance with 2283 GB of storage.
Large	Deploys an appliance with 2019 GB of storage.	Deploys an appliance with 2044 GB of storage.	Deploys an appliance with 2208 GB of storage.	Deploys an appliance with 2258 GB of storage.	Deploys an appliance with 2383 GB of storage.
X-Large	Deploys an appliance with 4279 GB of storage.	Deploys an appliance with 4304 GB of storage.	Deploys an appliance with 4468 GB of storage.	Deploys an appliance with 4518 GB of storage.	Deploys an appliance with 4643 GB of storage.

11. From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**. NFS datastores are thin provisioned by default.

12. Configure the temporary network for communication between the vCenter Server appliance that you want to upgrade and the new vCenter Server appliance, and click **Next**.

Option	Action
Choose a network	Select the network to which to connect the new appliance temporarily. The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are unsupported and are not displayed in the drop-down menu. Important: If you want to assign a temporary IPv4 address with DHCP allocation, you must select a network that is associated with a port group which accepts MAC address changes.
IP Address family	Select the version for the temporary IP address of the new appliance. Can be either IPv4 or IPv6.
Network type	Select the allocation method for the temporary IP address of the appliance.

Option	Action
	• Static The wizard prompts you to enter the temporary IP address, subnet mask or prefix length, default gateway, and DNS servers. • DHCP A DHCP server is used to allocate the temporary IP address. Select this option only if a DHCP server is available in your environment. Optionally, you can provide a temporary system name (FQDN) if a DDNS server is available in your environment.

13. On the Ready to complete stage 1 page, review the deployment settings for the new vCenter Server appliance and click **Finish** to start the OVA deployment process.

14. Wait for the OVA deployment process to finish and click **Continue** to proceed with stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the new appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface of the newly deployed vCenter Server appliance to transfer the data from the old appliance and set up the services.

The newly deployed vCenter Server appliance 8.0 is running on the target server but is not configured.

Important: The data from the old appliance is not transferred and the services of the new appliance are not started.

Stage 2 - Transfer the Data and Set up the Newly Deployed vCenter Server Appliance With an Embedded Platform Services Controller

When the OVA deployment finishes, you are redirected to stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the newly deployed vCenter Server appliance 8.0.

Understand the data migration options available to you when upgrading or migrating to a vCenter Server appliance with an embedded PostgreSQL database. See [Transferring Data from an Existing vCenter Server Appliance](#).

1. Review the introduction to stage 2 of the upgrade process and click **Next**.

2. Wait for the pre-upgrade check to finish and read the pre-upgrade check result if any.

- If the pre-upgrade check result contains error messages, read the messages and click **Logs** to export and download a support bundle for troubleshooting.

You cannot proceed with the upgrade until you have corrected the errors.

Important: If you have provided incorrect vCenter Single Sign-On user name and password of the source appliance during stage 1, the pre-upgrade check fails with an authentication error.

- If the pre-upgrade check result contains warning messages, read the messages and click **Close**.

After you have verified that your system meets the requirements from the warning message, you can proceed with the upgrade.

3. Connect to the source appliance that you want to upgrade.

- a) Enter the information about the source vCenter Server appliance that you want to upgrade, and click **Connect to Source**.

Option	Action
Appliance FQDN or IP address	Enter the IP address or FQDN of the vCenter Server appliance that you want to upgrade.

Option	Action
Appliance HTTPS port	If the source appliance uses a custom HTTPS port, change the default value to that of the custom port. The default port value is 443. Custom port values are supported beginning with vCenter Server appliance version 6.5 Update 2. If you are upgrading from earlier versions, you cannot specify a custom port.

- b) Enter the information about the vCenter Single Sign-On administrator and root user.

Option	Action
SSO user name	Enter the vCenter Single Sign-On administrator user name. Important: The user must be administrator@your_domain_name.
SSO password	Enter the password of the vCenter Single Sign-On administrator.
Appliance (OS) root password	Enter the password of the root user.

- c) Enter the information about the source ESXi host or vCenter Server instance on which resides the vCenter Server appliance that you want to upgrade and click **Next**.

Option	Description
Source server or host name	IP address or FQDN of the source ESXi host or vCenter Server instance on which the vCenter Server appliance that you want to upgrade resides. Note: The source vCenter Server instance cannot be the vCenter Server appliance that you want to upgrade. In such cases, use the source ESXi host.
HTTPS port	If the ESXi host or vCenter Server instance uses a custom HTTPS port, change the default value. The default value is 443.
User name	User name of a user with administrative privileges on the ESXi host or vCenter Server instance.
Password	Password of the user with administrative privileges on the ESXi host or vCenter Server instance.

4. On the **Select migration data** page, choose the types of data that you want to transfer from the old appliance to the new, upgraded appliance.

A large amount of data requires more time to be transferred to the new appliance. For the minimum upgrade time and storage requirements for the new appliance, select to transfer only the configuration data. If you are using an external Oracle database, you can also choose to migrate historical and performance metrics data in the background after you deploy and start the new vCenter Server appliance.

5. Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.

For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.

6. On the Ready to complete page, review the upgrade settings, accept the backup acknowledgment, and click **Finish**.
7. Read the shutdown warning message and click **OK**.
8. Wait for the data transfer and setup process to finish and click **OK** to go to the vCenter Server Getting Started page.

The vCenter Server appliance is upgraded. The old vCenter Server appliance is powered off and the new appliance starts.

- [Verify Your vCenter Server Upgrade or Migration Is Successful](#).
- If the old vCenter Server appliance uses a non-ephemeral distributed virtual port group, to preserve the port group setting, you can manually connect the new appliance to the original non-ephemeral distributed virtual port group. For information about configuring virtual machine networking on a vSphere distributed switch, see *vSphere Networking*.
- You can configure high availability for the vCenter Server appliance. For information about providing vCenter Server appliance high availability, see *vSphere Availability*.

Upgrade a vCenter Server Appliance 7.0 or 7.0 with an External Platform Services Controller Instance by Using the GUI

You can use the GUI installer to perform an interactive upgrade of a vCenter Server appliance 7.0 or 7.0 that uses an external Platform Services Controller instance to vCenter Server appliance 8.0. You must run the GUI upgrade from a Windows, Linux, or Mac OS X computer that is in the same network as the appliance that you want to upgrade.

- The new vCenter Server appliance contains all Platform Services Controller services. It is no longer necessary nor possible to deploy and use an external Platform Services Controller, as all Platform Services Controller services are consolidated into vCenter Server. To learn more about this change to vCenter Server, see [Removal of Platform Services Controller](#).
- See [Prerequisites for Upgrading the vCenter Server Appliance](#).
- See [Required Information for Upgrading a vCenter Server Appliance 7.0 or 8.0](#).

Stage 1 - Deploy the OVA File of the New vCenter Server Appliance 8.0

With stage 1 of the upgrade process, you deploy the OVA file of the new vCenter Server appliance 8.0.

1. In the vCenter Server installer, navigate to the `vcsa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
2. On the Home page, click **Upgrade**.
3. Review the Introduction page to understand the upgrade process and click **Next**.
4. Read and accept the license agreement, and click **Next**.
5. Connect to the source appliance that you want to upgrade.
 - a) Enter the information about the source vCenter Server appliance that you want to upgrade, and click **Connect to Source**.

Option	Action
Appliance FQDN or IP address	Enter the IP address or FQDN of the vCenter Server appliance that you want to upgrade.
Appliance HTTPS port	If the source appliance uses a custom HTTPS port, change the default value to that of the custom port. The default port value is 443.

Option	Action
	Custom port values are supported beginning with vCenter Server appliance version 6.5 Update 2. If you are upgrading from earlier versions, you cannot specify a custom port.

- b) Enter the information about the vCenter Single Sign-On administrator and root user.

Option	Action
SSO user name	Enter the vCenter Single Sign-On administrator user name. Important: The user must be administrator@ <i>your_domain_name</i> .
SSO password	Enter the password of the vCenter Single Sign-On administrator.
Appliance (OS) root password	Enter the password of the root user.

- c) Enter the information about the source ESXi host or vCenter Server instance on which resides the vCenter Server appliance that you want to upgrade and click **Next**.

Option	Description
Source server or host name	IP address or FQDN of the source ESXi host or vCenter Server instance on which the vCenter Server appliance that you want to upgrade resides. Note: The source vCenter Server instance cannot be the vCenter Server appliance that you want to upgrade. In such cases, use the source ESXi host.
HTTPS port	If the ESXi host or vCenter Server instance uses a custom HTTPS port, change the default value. The default value is 443.
User name	User name of a user with administrative privileges on the ESXi host or vCenter Server instance.
Password	Password of the user with administrative privileges on the ESXi host or vCenter Server instance.

- Verify that the certificate warning displays the SHA1 thumbprints of the SSL certificates that are installed on the source appliance and its source server, and click **Yes** to accept the certificate thumbprints.
- Confirm that the vCenter Server appliance using an external Platform Services Controller will be converged to a vCenter Server appliance with the Platform Services Controller services consolidated within vCenter Server, and click **Yes** to accept and continue with the upgrade.
- Connect to the target server on which you want to deploy the new vCenter Server appliance.

Option	Steps
You can connect to an ESXi host on which to deploy the new appliance.	- Enter the FQDN or IP address of the ESXi host. - Enter the HTTPS port of the ESXi host. - Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. - Click Next. - Accept the certificate warning, if any, by clicking Yes.

Option	Steps
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the new appliance. Note: The target server cannot be the vCenter Server appliance that you want to upgrade. In such cases, use an ESXi host as a target server.	1. Enter the FQDN or IP address of the vCenter Server instance. 2. Enter the HTTPS port of the vCenter Server instance. 3. Enter the user name and password of a vCenter Single Sign-On user with administrative privileges on the vCenter Server instance, for example, the administrator@<i>your_domain_name</i> user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes. 6. Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next. Note: You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode. 7. Select the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next.

9. On the Set up target appliance VM page, enter a name for the new vCenter Server appliance, set the password for the root user, and click **Next**.

The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.

The password must contain only lower ASCII characters without spaces, at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

Note: The root password of the old appliance is not transferred to the new upgraded appliance.

10. Select the deployment size for the new vCenter Server appliance for your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 vCPUs and 14 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 vCPUs and 21 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 vCPUs and 30 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 vCPUs and 39 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 vCPUs and 58 GB of memory. Suitable for environments with up to 2,500 hosts or 45,000 virtual machines

11. Select the storage size for the new vCenter Server appliance and click **Next**.

Important: You must consider the storage size of the appliance that you are upgrading and the database size if external.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 579 GB of storage.	Deploys an appliance with 694 GB of storage.	Deploys an appliance with 908 GB of storage.	Deploys an appliance with 1358 GB of storage.	Deploys an appliance with 2283 GB of storage.
Large	Deploys an appliance with 2019 GB of storage.	Deploys an appliance with 2044 GB of storage.	Deploys an appliance with 2208 GB of storage.	Deploys an appliance with 2258 GB of storage.	Deploys an appliance with 2383 GB of storage.
X-Large	Deploys an appliance with 4279 GB of storage.	Deploys an appliance with 4304 GB of storage.	Deploys an appliance with 4468 GB of storage.	Deploys an appliance with 4518 GB of storage.	Deploys an appliance with 4643 GB of storage.

12. From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**. NFS datastores are thin provisioned by default.

13. Configure the temporary network for communication between the vCenter Server appliance that you want to upgrade and the new vCenter Server appliance, and click **Next**.

Option	Action
Choose a network	Select the network to which to connect the new appliance temporarily. The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are unsupported and are not displayed in the drop-down menu. Important: If you want to assign a temporary IPv4 address with DHCP allocation, you must select a network that is associated with a port group which accepts MAC address changes.
IP Address family	Select the version for the temporary IP address of the new appliance. Can be either IPv4 or IPv6.
Network type	Select the allocation method for the temporary IP address of the appliance. • Static The wizard prompts you to enter the temporary IP address, subnet mask or prefix length, default gateway, and DNS servers. • DHCP A DHCP server is used to allocate the temporary IP address. Select this option only if a DHCP server is available in your environment. Optionally, you can provide a temporary system name (FQDN) if a DDNS server is available in your environment.

14. On the Ready to complete stage 1 page, review the deployment settings for the new vCenter Server appliance and click **Finish** to start the OVA deployment process.
15. Wait for the OVA deployment process to finish and click **Continue** to proceed with stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the new appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface of the newly deployed vCenter Server appliance to transfer the data from the old appliance and set up the services.

The newly deployed target vCenter Server appliance 8.0 is running on the target server but is not configured.

Proceed to stage 2 of the upgrade process to transfer the data from the source vCenter Server appliance and start the services of the newly deployed vCenter Server appliance. See [Stage 2 - Transfer the Data and Set up the Newly Deployed vCenter Server Appliance](#).

Stage 2 - Transfer the Data and Set up the Newly Deployed vCenter Server Appliance

When the OVA deployment finishes, you are redirected to stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the newly deployed vCenter Server appliance 8.0.

Understand the data migration options available to you when upgrading or migrating to a vCenter Server appliance with an embedded PostgreSQL database. See [Transferring Data from an Existing vCenter Server Appliance](#).

1. Review the introduction to stage 2 of the upgrade process and click **Next**.
2. Wait for the pre-upgrade check to finish and read the pre-upgrade check result if any.
 - If the pre-upgrade check result contains error messages, read the messages and click **Logs** to export and download a support bundle for troubleshooting.

You cannot proceed with the upgrade until you have corrected the errors.

Important: If you have provided incorrect vCenter Single Sign-On user name and password of the source appliance during stage 1, the pre-upgrade check fails with an authentication error.

- If the pre-upgrade check result contains warning messages, read the messages and click **Close**.

After you have verified that your system meets the requirements from the warning message, you can proceed with the upgrade.

3. Specify the replication topology for the vCenter Server. When converging vCenter Server instance with an external Platform Services Controller you must specify the replication topology.

Convergence is the process of converting a vCenter Server instance with an external Platform Services Controller to a vCenter Server instance with those services embedded in the appliance.

- This is the first vCenter Server in the topology that I want to converge.
- This is a subsequent vCenter Server.

If this is a subsequent vCenter Server, provide the IP address of its partner vCenter Server and its HTTPS port.

4. On the **Select upgrade data** page, choose the types of data that you want to transfer from the old vCenter Server appliance to the new, upgraded vCenter Server appliance.

A large amount of data requires more time to be transferred to the new appliance. For the minimum upgrade time and storage requirements for the new appliance, select to transfer only the configuration data.

Note: If you are using an external Oracle database, you can also choose to migrate historical and performance metrics data in the background after you deploy and start the new vCenter Server appliance.

5. Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.
For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.
6. On the Ready to complete page, review the upgrade settings, accept the backup acknowledgment, and click **Finish**.
7. Read the shutdown warning message and click **OK**.
8. Wait for the data transfer and setup process to finish and click **OK** to go to the vCenter Server Getting Started page.

The vCenter Server appliance is upgraded. The old vCenter Server appliance is powered off and the new appliance starts.

- [Verify Your vCenter Server Upgrade or Migration Is Successful](#).
- If the old vCenter Server appliance uses a non-ephemeral distributed virtual port group, to preserve the port group setting, you can manually connect the new appliance to the original non-ephemeral distributed virtual port group. For information about configuring virtual machine networking on a vSphere distributed switch, see *vSphere Networking*.
- Upgrade all vCenter Server instances in the vCenter Single Sign-On domain.
- After converging a vCenter Server with external Platform Services Controller to a vCenter Server appliance, you must decommission the original external Platform Services Controller. Decommissioning a Platform Services Controller shuts it down and removes it from the single sign-on domain. See [Decommission the Platform Services Controller](#).
- You can configure high availability for the vCenter Server appliance. For information about providing vCenter Server appliance high availability, see *vSphere Availability*.

Upgrading vCenter Server in High Availability Environments

You can use the GUI installer to perform an interactive upgrade of a vCenter Server appliance in a high availability (HA) environment.

Prerequisites for Upgrading vCenter Server High Availability Environments

To ensure a successful upgrade of vCenter Server in a high availability (HA) environment, your environment must meet certain prerequisites before running the upgrade.

General Prerequisites

Ensure that your environment meets the prerequisites for upgrading vCenter Server appliance. See [Prerequisites for Upgrading the vCenter Server Appliance](#).

High Availability Prerequisites

In addition to the prerequisites for upgrading a standard vCenter Server appliance, the vCenter High Availability (vCenter HA) upgrade has the following prerequisites.

- A vCenter HA cluster consists of three vCenter Server appliances that act as Active, Passive, and Witness nodes. The Active node must be configured as the vCenter HA node.
- The Active node is part of the vCenter HA cluster.
- All nodes must be present in the cluster.
- The vCenter HA cluster must be in a healthy state.
- The vCenter HA cluster must be in enabled mode.
- The host that has the vCenter Server virtual machine (VM) must be managed by a container vCenter Server. It must not be standalone.
- The target vCenter Server must have the same placement as the source vCenter Server.
- vCenter HA cannot be set up on a vCenter Server that has disks on more than one datastore.

To learn more about vCenter HA configuration, see *vSphere Availability*.

Upgrade a vCenter Server Appliance 7.0 or 6.7 HA Cluster with an Embedded Platform Services Controller by Using the GUI

You can use the GUI installer to perform an interactive upgrade of a vCenter Server appliance 6.7 or 7.0 in a High Availability (HA) cluster that uses an embedded vCenter Single Sign-On or Platform Services Controller to vCenter Server appliance 8.0. You must run the GUI upgrade from a Windows, Linux, or Macintosh computer that is in the same network as the appliance that you want to upgrade.

- See [Prerequisites for Upgrading vCenter Server High Availability Environments](#).
- See [Required Information for Upgrading a vCenter Server Appliance 7.0 or 8.0](#).

You can deploy version 8.0 of the vCenter Server appliance on hosts that are running ESXi6.7 or later and on vCenter Server instances 6.7 or later.

A vCenter HA cluster consists of three vCenter Server appliances that act as Active, Passive, and Witness nodes. The Active vCenter HA node is upgraded using a migration-based upgrade while preserving the existing configuration.

Stage 1 - Deploy the OVA File of the New vCenter High Availability Cluster

In stage 1 of the upgrade process, you deploy the OVA file for the new vCenter Server appliance.

Ensure that your environment meets the prerequisites for upgrading a vCenter HA cluster. See [Prerequisites for Upgrading vCenter Server High Availability Environments](#).

1. In the vCenter Server installer, navigate to the `vcsa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
2. On the Home page, click **Upgrade**.
3. Review the Introduction page to understand the upgrade process and click **Next**.
4. Connect to the source appliance that you want to upgrade. This appliance is the Active vCenter HA node.
 - a) Enter the information about the source vCenter Server appliance that you want to upgrade, and click **Connect to Source**.

Option	Action
Appliance FQDN or IP address	Enter the IP address or FQDN of the Active vCenter HA node that you want to upgrade.
Appliance HTTPS port	The default value (443) is displayed and cannot be edited.

- b) Enter the information about the vCenter Single Sign-On administrator and root user.

Option	Action
SSO user name	Enter the vCenter Single Sign-On administrator user name. Important: The user must be <code>administrator@your_domain_name</code> .
SSO password	Enter the password of the vCenter Single Sign-On administrator.

Option	Action
Appliance (OS) root password	Enter the password of the root user.

- c) Enter the information about the source vCenter Server instance on which resides the vCenter Server appliance that you want to upgrade and click **Next**.

Option	Description
Source server or host name	IP address or FQDN of the Active node. The Active node must be configured as the vCenter HA node.
HTTPS port	If the vCenter Server instance uses a custom HTTPS port, change the default value. The default value is 443.
User name	User name of a user with administrative privileges on the ESXi host or vCenter Server instance.
Password	Password of the user with administrative privileges on the ESXi host or vCenter Server instance.

- Verify that the certificate warning displays the SHA1 thumbprints of the SSL certificates that are installed on the source appliance and its source server, and click **Yes** to accept the certificate thumbprints.
- If vCenter HA is successfully detected, the target appliance is set to the manager of the source appliance. Click **OK**. The appliance deployment target information is populated.
- On the Set up target appliance VM page, enter a name for the target vCenter Server appliance, set the password for the root user, and click **Next**.

The password must contain at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

Note: The root password of the source is not transferred to the target appliance.

- Select the deployment size for the new vCenter Server appliance for your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 vCPUs and 14 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 vCPUs and 21 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 vCPUs and 30 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 vCPUs and 39 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 vCPUs and 58 GB of memory. Suitable for environments with up to 2,500 hosts or 45,000 virtual machines

9. Select the storage size for the new vCenter Server appliance and click **Next**.

Important: You must consider the storage size of the appliance that you are upgrading and the database size if external.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 579 GB of storage.	Deploys an appliance with 694 GB of storage.	Deploys an appliance with 908 GB of storage.	Deploys an appliance with 1358 GB of storage.	Deploys an appliance with 2283 GB of storage.
Large	Deploys an appliance with 2019 GB of storage.	Deploys an appliance with 2044 GB of storage.	Deploys an appliance with 2208 GB of storage.	Deploys an appliance with 2258 GB of storage.	Deploys an appliance with 2383 GB of storage.
X-Large	Deploys an appliance with 4279 GB of storage.	Deploys an appliance with 4304 GB of storage.	Deploys an appliance with 4468 GB of storage.	Deploys an appliance with 4518 GB of storage.	Deploys an appliance with 4643 GB of storage.

10. From the list of available datastores, select the location to store the virtual machine (VM) configuration files and virtual disks, and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**. NFS datastores are thin provisioned by default.

Note: vCenter HA cannot be set up on a vCenter Server that has disks on more than one datastore.

11. Configure the temporary network for communication between the vCenter Server appliance that you want to upgrade and the new vCenter Server appliance, and click **Next**.

Option	Action
Choose a network	Select the network to which to connect the new appliance temporarily. The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are unsupported and are not displayed in the drop-down menu. Important: If you want to assign a temporary IPv4 address with DHCP allocation, you must select a network that is associated with a port group which accepts MAC address changes.
IP Address family	Select the version for the temporary IP address of the new appliance. Can be either IPv4 or IPv6.
Network type	Select the allocation method for the temporary IP address of the appliance. • Static The wizard prompts you to enter the temporary IP address, subnet mask or prefix length, default gateway, and DNS servers. • DHCP A DHCP server is used to allocate the temporary IP address. Select this option only if a DHCP server is available in your environment. Optionally, you can provide a temporary system name (FQDN) if a DDNS server is available in your environment.

12. On the Ready to complete stage 1 page, review the deployment settings for the new vCenter Server appliance and click **Finish** to start the OVA deployment process.
13. Wait for the OVA deployment process to finish and click **Continue** to proceed with stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the new appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface of the newly deployed vCenter Server appliance to transfer the data from the old appliance and set up the services.

The newly deployed vCenter Server appliance 8.0 is running on the target server but is not configured.

Important: The data from the old appliance is not transferred and the services of the new appliance are not started.

Stage 2 - Transfer Data and Set-Up the Newly Deployed vCenter High Availability Cluster

When the OVA deployment finishes, you are redirected to stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the newly deployed vCenter Server appliance. When the deployment finishes, vCenter Server has high availability protection.

Understand the data migration options available to you when upgrading or migrating to a vCenter Server appliance with an embedded PostgreSQL database. You can select to migrate historical and other types of data in the background after deploying and starting vCenter Server. See [Transferring Data from an Existing vCenter Server Appliance](#).

1. Review the introduction to stage 2 of the deployment process and click **Next**.
2. Wait for the pre-upgrade check to finish and read the pre-upgrade check result if any.
 - If the pre-upgrade check result contains error messages, read the messages and click **Logs** to export and download a support bundle for troubleshooting.
You cannot proceed with the upgrade until you have corrected the errors.
Important: If you have provided incorrect vCenter Single Sign-On user name and password of the source appliance during stage 1, the pre-upgrade check fails with an authentication error.
 - If the pre-upgrade check result contains warning messages, read the messages and click **Close**.
After you have verified that your system meets the requirements from the warning message, you can proceed with the upgrade.
3. On the **Select migration data** page, choose the types of data that you want to transfer from the old appliance to the new, upgraded appliance.
A large amount of data requires more time to be transferred to the new appliance. For the minimum upgrade time and storage requirements for the new appliance, select to transfer only the configuration data. If you are using an external Oracle database, you can also choose to migrate historical and performance metrics data in the background after you deploy and start the new vCenter Server appliance.
4. Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.
For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.
5. On the Ready to complete page, review the upgrade settings, accept the backup acknowledgment, and click **Finish**.
6. Read the shutdown warning message and click **OK**.
7. Wait for the data transfer and setup process to finish and click **OK** to go to the vCenter Server Getting Started page.

The vCenter Server appliance is upgraded. The old vCenter Server appliance is powered off and the new appliance starts.

After the Active node is upgraded, Auto Deployment automatically creates new Passive and Witness nodes using a clone operation. For Manual Deployment, nodes are not created automatically. You must clone the Passive and Witness virtual machines, and set the cluster mode to **Enabled**.

When the deployment finishes, vCenter Server has high availability protection. You can click **Edit** to enter Maintenance Mode, Deactivate, or Remove vCenter HA. You can also Initiate vCenter HA failover.

For information about configuring and managing vCenter HA, see *vSphere Availability*.

Upgrade a vCenter Server Appliance 6.7 HA Cluster with an External Platform Services Controller by Using the GUI

You can use the GUI installer to perform an interactive upgrade of a vCenter HA cluster 6.7 that uses an external Platform Services Controller instance to vCenter Server appliance 8.0. You must run the GUI upgrade from a Windows, Linux, or Macintosh computer that is in the same network as the appliance that you want to upgrade.

- The new vCenter Server appliance contains all Platform Services Controller services. It is no longer necessary nor possible to deploy and use an external Platform Services Controller, as all Platform Services Controller services are consolidated into vCenter Server. To learn more about this change to vCenter Server, see [Removal of Platform Services Controller](#).
- Ensure that your environment meets the prerequisites for upgrading a vCenter HA cluster. See [Prerequisites for Upgrading vCenter Server High Availability Environments](#).
- See [Required Information for Upgrading a vCenter Server Appliance 7.0 or 8.0](#).

A vCenter HA cluster consists of three vCenter Server appliances that act as Active, Passive, and Witness nodes. The Active vCenter HA node is upgraded using a migration-based upgrade while preserving the existing configuration.

Stage 1 - Deploy the OVA File of the New vCenter High Availability Cluster

With stage 1 of the upgrade process, you deploy the OVA file of the new vCenter Server appliance 8.0.

Ensure that your environment meets the prerequisites for upgrading a vCenter HA cluster. See [Prerequisites for Upgrading vCenter Server High Availability Environments](#).

1. In the vCenter Server installer, navigate to the `vcasa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
2. On the Home page, click **Upgrade**.
3. Review the Introduction page to understand the upgrade process and click **Next**.
4. Read and accept the license agreement, and click **Next**.
5. Connect to the source appliance that you want to upgrade. This appliance is the Active vCenter HA node.
 - a) Enter the information about the source vCenter Server appliance that you want to upgrade, and click **Connect to Source**.

Option	Action
Appliance FQDN or IP address	Enter the IP address or FQDN of the Active vCenter HA node that you want to upgrade.

Option	Action
Appliance HTTPS port	The default value (443) is displayed and cannot be edited.

b) Enter the information about the vCenter Single Sign-On administrator and root user.

Option	Action
SSO user name	Enter the vCenter Single Sign-On administrator user name. Important: The user must be administrator@your_domain_name.
SSO password	Enter the password of the vCenter Single Sign-On administrator.
Appliance (OS) root password	Enter the password of the root user.

c) Enter the information about the source vCenter Server instance on which resides the vCenter Server appliance that you want to upgrade and click **Next**.

Option	Description
Source server or host name	IP address or FQDN of the Active node. The Active node must be configured as the vCenter HA node.
HTTPS port	If the vCenter Server instance uses a custom HTTPS port, change the default value. The default value is 443.
User name	User name of a user with administrative privileges on the ESXi host or vCenter Server instance.
Password	Password of the user with administrative privileges on the ESXi host or vCenter Server instance.

6. Verify that the certificate warning displays the SHA1 thumbprints of the SSL certificates that are installed on the source appliance and its source server, and click **Yes** to accept the certificate thumbprints.

7. Confirm that the vCenter Server appliance using an external Platform Services Controller will be converged to a vCenter Server appliance with the Platform Services Controller services consolidated within vCenter Server, and click **Yes** to accept and continue with the upgrade.

8. If vCenter Server HA is successfully detected, the target appliance is set to the manager of the source appliance. Click **OK**.

The appliance deployment target information is populated.

9. Connect to the target server on which you want to deploy the new vCenter Server appliance.

Option	Steps
You can connect to an ESXi host on which to deploy the new appliance.	1. Enter the FQDN or IP address of the ESXi host. 2. Enter the HTTPS port of the ESXi host. 3. Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the new appliance.	1. Enter the FQDN or IP address of the vCenter Server instance. 2. Enter the HTTPS port of the vCenter Server instance.

Option	Steps
Note: The target server cannot be the vCenter Server appliance that you want to upgrade. In such cases, use an ESXi host as a target server.	3. Enter the user name and password of a vCenter Single Sign-On user with administrative privileges on the vCenter Server instance, for example, the administrator@<i>your_domain_name</i> user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes. 6. Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next. Note: You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode. 7. Select the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next.

10. On the Set up target appliance VM page, enter a name for the new vCenter Server appliance, set the password for the root user, and click **Next**.

The appliance name must not contain a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.

The password must contain only lower ASCII characters without spaces, at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

Note: The root password of the old appliance is not transferred to the new upgraded appliance.

11. Select the deployment size for the new vCenter Server appliance for your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 vCPUs and 14 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 vCPUs and 21 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 vCPUs and 30 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 vCPUs and 39 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 vCPUs and 58 GB of memory. Suitable for environments with up to 2,500 hosts or 45,000 virtual machines

12. Select the storage size for the new vCenter Server appliance, and click **Next**.

Important: You must consider the storage size of the appliance that you are upgrading and the database size if external.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 579 GB of storage.	Deploys an appliance with 694 GB of storage.	Deploys an appliance with 908 GB of storage.	Deploys an appliance with 1358 GB of storage.	Deploys an appliance with 2283 GB of storage.
Large	Deploys an appliance with 2019 GB of storage.	Deploys an appliance with 2044 GB of storage.	Deploys an appliance with 2208 GB of storage.	Deploys an appliance with 2258 GB of storage.	Deploys an appliance with 2383 GB of storage.
X-Large	Deploys an appliance with 4279 GB of storage.	Deploys an appliance with 4304 GB of storage.	Deploys an appliance with 4468 GB of storage.	Deploys an appliance with 4518 GB of storage.	Deploys an appliance with 4643 GB of storage.

13. From the list of available datastores, select the location to store the virtual machine (VM) configuration files and virtual disks, and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**. NFS datastores are thin provisioned by default.

Note: vCenter HA cannot be set up on a vCenter Server that has disks on more than one datastore.

14. Configure the temporary network for communication between the vCenter Server appliance that you want to upgrade and the new vCenter Server appliance, and click **Next**.

Option	Action
Choose a network	Select the network to which to connect the new appliance temporarily. The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are unsupported and are not displayed in the drop-down menu. Important: If you want to assign a temporary IPv4 address with DHCP allocation, you must select a network that is associated with a port group which accepts MAC address changes.
IP Address family	Select the version for the temporary IP address of the new appliance. Can be either IPv4 or IPv6.
Network type	Select the allocation method for the temporary IP address of the appliance. • Static The wizard prompts you to enter the temporary IP address, subnet mask or prefix length, default gateway, and DNS servers. • DHCP A DHCP server is used to allocate the temporary IP address. Select this option only if a DHCP server is available in your environment. Optionally, you can provide a temporary system name (FQDN) if a DDNS server is available in your environment.

15. On the Ready to complete stage 1 page, review the deployment settings for the new vCenter Server appliance and click **Finish** to start the OVA deployment process.
16. Wait for the OVA deployment process to finish and click **Continue** to proceed with stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the new appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface of the newly deployed vCenter Server appliance to transfer the data from the old appliance and set up the services.

The newly deployed target vCenter Server appliance 8.0 is running on the target server but is not configured.

Important: The data from the source vCenter Server is not transferred and the services of the target appliance are not started.

Transfer the data from the old appliance and start the services of the newly deployed vCenter Server Appliance. See [Stage 2 - Transfer Data and Set-Up the Newly Deployed vCenter High Availability Cluster](#).

Stage 2 - Transfer Data and Set-Up the Newly Deployed vCenter High Availability Cluster

When the OVA deployment finishes, you are redirected to stage 2 of the upgrade process to transfer the data from the old appliance and start the services of the newly deployed vCenter Server appliance. When the deployment finishes, vCenter Server has high availability protection.

Understand the data migration options available to you when upgrading or migrating to a vCenter Server appliance with an embedded PostgreSQL database. You can select to migrate historical and other types of data in the background after deploying and starting vCenter Server. See [Transferring Data from an Existing vCenter Server Appliance](#).

1. Review the introduction to stage 2 of the upgrade process and click **Next**.
2. Wait for the pre-upgrade check to finish and read the pre-upgrade check result if any.
 - If the pre-upgrade check result contains error messages, read the messages and click **Logs** to export and download a support bundle for troubleshooting.

You cannot proceed with the upgrade until you have corrected the errors.

Important: If you have provided incorrect vCenter Single Sign-On user name and password of the source appliance during stage 1, the pre-upgrade check fails with an authentication error.

- If the pre-upgrade check result contains warning messages, read the messages and click **Close**.

After you have verified that your system meets the requirements from the warning message, you can proceed with the upgrade.

3. Specify the replication topology for the vCenter Server. When converging vCenter Server instance with an external Platform Services Controller you must specify the replication topology.

Convergence is the process of converting a vCenter Server instance with an external Platform Services Controller to a vCenter Server instance with those services embedded in the appliance.

- This is the first vCenter Server in the topology that I want to converge.
- This is a subsequent vCenter Server.

If this is a subsequent vCenter Server, provide the IP address of its partner vCenter Server and its HTTPS port.

4. On the **Select migration data** page, choose the types of data that you want to transfer from the old appliance to the new, upgraded appliance.

A large amount of data requires more time to be transferred to the new appliance. For the minimum upgrade time and storage requirements for the new appliance, select to transfer only the configuration data. If you are using an external Oracle database, you can also choose to migrate historical and performance metrics data in the background after you deploy and start the new vCenter Server appliance.

5. On the Ready to complete page, review the upgrade settings, accept the backup acknowledgment, and click **Finish**.
6. Read the shutdown warning message and click **OK**.
7. Wait for the data transfer and setup process to finish and click **OK** to go to the vCenter Server Getting Started page.

The vCenter Server appliance is upgraded. The old vCenter Server appliance is powered off and the new appliance starts.

After the Active node is upgraded, a new Passive and Witness node is created using a clone operation. These nodes are created automatically for Auto Deployment. For Manual Deployment, nodes are not created automatically. You must clone the Passive and Witness VMs and set the cluster mode to **enabled**.

When the deployment finishes, vCenter Server has high availability protection. You can click **Edit** to enter Maintenance Mode, Deactivate, or Remove vCenter HA. You can also Initiate vCenter HA failover.

- [Verify Your vCenter Server Upgrade or Migration Is Successful](#).
- If the old vCenter Server appliance uses a non-ephemeral distributed virtual port group, to preserve the port group setting, you can manually connect the new appliance to the original non-ephemeral distributed virtual port group. For information about configuring virtual machine networking on a vSphere distributed switch, see *vSphere Networking*.
- Upgrade all vCenter Server instances in the vCenter Single Sign-On domain.
- After converging a vCenter Server with external Platform Services Controller node to vCenter Server Appliance, you must decommission the original external Platform Services Controller. Decommissioning a Platform Services Controller shuts it down and removes it from the single sign-on domain. See [Decommission the Platform Services Controller](#).
- You can configure high availability for the vCenter Server appliance. For information about providing vCenter Server appliance high availability, see *vSphere Availability*.

CLI Upgrade of the vCenter Server Appliance

You can use the CLI installer to perform an unattended upgrade of a vCenter Server appliance on an ESXi host or vCenter Server instance.

The CLI upgrade process includes downloading the vCenter Server appliance installer on a network virtual machine or physical server from which you want to perform the upgrade, preparing a JSON configuration file with the upgrade information, and running the upgrade command.

Important: The user name that you use to log in to the machine from which you want to run the CLI upgrade, the path to the vCenter Server appliance ISO file, the path to your JSON configuration file, and the string values in your JSON configuration file, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

The vCenter Server appliance ISO file contains templates of JSON files that contain the minimum configuration parameters that are required for upgrading a vCenter Server appliance. For information about preparing JSON templates for CLI upgrade of the vCenter Server appliance, see [Prepare Your JSON Configuration File for CLI Upgrade](#).

Note: When you upgrade vCenter Server 6.7 or 7.0 with an external Platform Services Controller to vCenter Server 8.0, the upgrade process converges the external Platform Services Controller into the new vCenter Server appliance. The new vCenter Server contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, and licensing. It is no longer necessary nor possible to upgrade and use an external Platform Services Controller.

After successfully upgrading your environment to vCenter Server 8.0, the pre-existing Platform Services Controller will be powered off, and you can remove it from your vSphere inventory. See [Decommission the Platform Services Controller](#).

Prepare Your JSON Configuration File for CLI Upgrade

Before you run the CLI command to upgrade a vCenter Server appliance, you must prepare a JSON file with configuration parameters and their values for your upgrade specification.

- You must be familiar with the JSON syntax.
- [Download and Mount the vCenter Server Installer](#).

The vCenter Server installer contains JSON templates for all upgrade types. For information about the templates, see [JSON Templates for CLI Upgrade of the vCenter Server Appliance](#).

You can upgrade an appliance with minimum configurations by setting values to the configuration parameters in the JSON template for your specification. You can edit the preset values, remove configuration parameters, and add configuration parameters for custom configurations.

For a complete list of the configuration parameters and their descriptions, navigate to the installer subdirectory for your operating system and run the `vcasa-deploy upgrade --template-help` command or see [Upgrade Configuration Parameters for CLI Upgrade of vCenter Server Appliance](#).

1. In the vCenter Server Appliance installer, navigate to the `vcasa-cli-installer` directory, and open the `templates` subfolder.
2. Copy the upgrade templates from the `upgrade` subfolder to your workspace.

Important: The path to the JSON configuration files must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

3. Open the template file for your use case in a text editor.

To ensure the correct syntax of your JSON configuration file, use a JSON editor.

4. Fill in the values for the required configuration parameters and, optionally, enter additional parameters and their values.

For example, if you want to use an IPv4 DHCP assignment for the temporary network of the new appliance, in the `temporary_network` subsection of the template, change the value of the `mode` parameter to `dhcp` and remove the default configuration parameters that are for a static assignment.

```
"temporary_network": {
  "ip_family": "ipv4",
  "mode": "dhcp"
},
```

Important: The string values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To set a value that contains a backslash (\) or quotation mark (") character, you must precede the character with the backslash (\) character. For example, `"password": "my\"password"` sets the password `my"password`, `"image": "G:\\vcasa\\VMware-vCenter-Server-Appliance-6.7.0.XXXX-YYYYYYY_OVF10.ova"` sets the path `G:\vcasa\VMware-vCenter-Server-Appliance-6.7.0.XXXX-YYYYYYY_OVF10.ova`.

The Boolean values must contain only lowercase characters, that is, a value can be either `true` or `false`. For example, `"ssh_enable": false`.

5. Optional: Use a JSON editor of your choice to validate the JSON file.
6. Save in UTF-8 format and close the file.

You can create and save additional templates if needed for your upgrade specification.

JSON Templates for CLI Upgrade of the vCenter Server Appliance

The vCenter Server appliance installer contains JSON templates that are located in the `vcasa-cli-installer/templates` directory. In the `upgrade` subfolder, you can find the JSON templates with the minimum configuration parameters for all upgrade types.

For each upgrade type, there is one template for deploying the new appliance on an ESXi host and another template for deploying the new appliance on a vCenter Server instance.

Table 64: Upgrade JSON Templates Included in the vCenter Server Appliance Installer

Location	Template	Description
vcsa-cli-installer\templates\ upgrade\vcsa7.0	embedded_vCSA_on_ESXi.json	Contains the minimum configuration parameters that are required for upgrade of a vCenter Server appliance 6.7 with an embedded Platform Services Controller to vCenter Server appliance 7.0 on an ESXi host.
	embedded_vCSA_on_VC.json	Contains the minimum configuration parameters that are required for upgrade of a vCenter Server appliance 6.7 with an embedded Platform Services Controller to vCenter Server Appliance 7.0 on a vCenter Server instance.
vcsa-cli-installer\templates\ upgrade\vcsa8.0	embedded_vCSA_on_ESXi.json	Contains the minimum configuration parameters that are required for upgrade of a vCenter Server appliance 7.0 with an embedded Platform Services Controller to vCenter Server appliance 8.0 on an ESXi host.
	embedded_vCSA_on_VC.json	Contains the minimum configuration parameters that are required for upgrade of a vCenter Server appliance 7.0 with an embedded Platform Services Controller to vCenter Server Appliance 8.0 on a vCenter Server instance.

About Upgrading a vCenter Server with an External Platform Services Controller Using the CLI

When upgrading a vCenter Server appliance that uses an external Platform Services Controller, it is converted to a vCenter Server instance with those services embedded in the appliance. For this reason, you must specify replication partner parameters in the JSON upgrade templates.

When upgrading a vCenter Server appliance that uses an external Platform Services Controller, you must specify replication partner parameters in the JSON upgrade template. These parameters indicate if the upgrade is for one of the following vCenter Server and Platform Services Controller upgrade scenarios.

- A single vCenter Server and a single Platform Services Controller instance within an SSO domain.
- The first vCenter Server and Platform Services Controller instance within an SSO domain containing more than one vCenter Server.
- Replicating vCenter Server and Platform Services Controller instances within an SSO domain, which you must point to an existing vCenter Server 7.0 instance as the replication partner.

Important: It is not possible to upgrade an external Platform Services Controller. The upgrade process converges the Platform Services Controller services into the vCenter Server appliance. For more information, see [Removal of Platform Services Controller](#).

Deployments using an external Platform Services Controller are often called an $M \times N$ deployment, where M represents the vCenter Server, and N represents the external Platform Services Controller. When representing multiple instances of vCenter Server and Platform Services Controller in a domain, this is expressed as $M \times N \times n$, where n is the number of

instances. For example, M2xN2 indicates two vCenter Server instances and two Platform Services Controller instances within the same domain. When upgrading such a deployment, you first upgrade one instance to version 8.0, and then upgrade the second instance to 8.0, specifying the first of the two instances as the replicating partner.

The replication parameters in the JSON upgrade template are located in the `sso` section of the template, and are `first_instance` and `replication_partner_hostname`.

Parameter Name	Description
<code>first_instance</code>	The <code>first_instance</code> parameter identifies the vCenter Server as either the first instance in your vCenter Server SSO domain, or as a replicating partner. If it is the first vCenter Server instance, set <code>first_instance</code> to <code>true</code> . For additional instances of vCenter Server in the same SSO domain, set <code>first_instance</code> to <code>false</code> .
<code>replication_partner_hostname</code> *	If you set <code>first_instance</code> to <code>false</code> , you must specify the FQDN or IP address of a vCenter Server that has already been upgraded to version 8.0, and whose external PSC has been converged into the vCenter Server appliance.

When you perform the CLI upgrade of a single vCenter Server appliance that uses an external Platform Services Controller, the process is as follows:

1. Upgrade the vCenter Server instance in your topology using the `vCSA_on_ESXi.json` template when upgrading on ESXi, or the `vCSA_on_VC.json` template when upgrading on a vCenter Server deployment.
2. For a deployment topology consisting of a single vCenter Server and a Platform Services Controller (an M1xN1 deployment), edit the `sso` section of the template such that the `first_instance` parameter is set to `true`, and remove the `replication_partner_hostname` parameter.

```
"sso": {
  "__comments": [ This is the first instance (M1) vCenter Server management node, with "first_instance" set
    to "true" and "replication_partner_hostname" removed.
  ],
  "first_instance": true
}
```
3. Specify the remaining upgrade configuration parameters in the JSON template, and upgrade vCenter Server to version 8.0 using the CLI. See [Upgrade a vCenter Server Appliance by Using the CLI](#).
4. After successfully upgrading vCenter Server, decommission the Platform Services Controller. See [Decommission the Platform Services Controller](#).

Important: Decommission the Platform Services Controller instances only after upgrading and converging all vCenter Server instances in your topology.

For a deployment consisting of two or more vCenter Server and Platform Services Controller instances (an M2xN2 deployment), you first upgrade the M1 vCenter Server instance to version 8.0, and then upgrade the second, replicating vCenter Server instance.

1. Upgrade the vCenter Server instances in your deployment using the `vCSA_on_ESXi.json` template when upgrading on ESXi, or the `vCSA_on_VC.json` template when upgrading on a vCenter Server deployment.
2. For the first vCenter Server instance (M1), edit the `sso` section of the template such that the `first_instance` parameter is set to `true`, and remove the `replication_partner_hostname` parameter.

```
"sso": {
  "__comments": [ This is the first instance (M1) vCenter Server management node, with "first_instance" set
    to "true" and "replication_partner_hostname" removed.
  ],
  "first_instance": true
}
```

```
}

```

- For the remaining vCenter Server instance (M2), edit the `sso` section of the template such that the `first_instance` parameter is set to `false`, and provide the FQDN or IP address of the vCenter Server instance to which you want to replicate the second node. The replication partner is the first vCenter Server instance (M1) you upgraded to version 8.0, which has been converged with its Platform Services Controller instance.

```
"sso": {
  "__comments": [ This is the second instance (M2) vCenter Server management node, with "first_instance" set
    to "false" and "replication_partner_hostname" set to the hostname of the replication partner.],
  "first_instance": false,
  "replication_partner_hostname": "FQDN_or_IP_address"
}
```

- Specify the remaining upgrade configuration parameters in the JSON templates, and finish the upgrade using the CLI. See [Upgrade a vCenter Server Appliance by Using the CLI](#).
- After successfully upgrading all vCenter Server instances in the topology, decommission the Platform Services Controller instances. See [Decommission the Platform Services Controller](#).

Important: Decommission the Platform Services Controller instances only after upgrading and converging all vCenter Server instances in your topology.

Upgrade Configuration Parameters for CLI Upgrade of vCenter Server Appliance

When you prepare your JSON configuration files for CLI upgrade, you must set parameters and values to provide input data for the upgrade of a vCenter Server appliance.

Sections and Subsections of Configuration Parameters in the JSON Upgrade Files

The configuration parameters in the JSON configuration files for CLI upgrade are organized in sections and subsections.

Table 65: Sections and Subsections of Configuration Parameters in the JSON Upgrade Files

Section	Subsection	Description
<code>new_vcsa</code> - describes the new appliance that you want to deploy.	<code>esxi</code>	Use only if you want to deploy the new appliance directly on an ESXi host. Contains the configuration parameters that describe the target ESXi host. See Table 66: Configuration Parameters in the new_vcsa Section, esxi Subsection . Note: You must fill in either this subsection or the <code>vc</code> subsection.
	<code>vc</code>	Use only if you want to deploy the new appliance on the inventory of a vCenter Server instance. Contains the configuration parameters that describe the target ESXi host or DRS cluster from the vCenter Server inventory. See Table 67: Configuration Parameters in the new_vcsa Section, vc Subsection . Note: You must fill in either this subsection or the <code>esxi</code> subsection. The target vCenter Server instance cannot be the vCenter Server Appliance that you want to upgrade. In such cases, use the <code>esxi</code> subsection.
	<code>appliance</code>	Contains the configuration parameters that describe the new appliance. See Table 68: Configuration Parameters in the new_vcsa Section, appliance Subsection
	<code>os</code>	Contains only the <code>ssh_enable</code> configuration parameter to set the SSH administrator login to the new appliance.
	<code>ovftool_arguments</code>	Optional. Use this subsection for adding arbitrary arguments and their values to the OVF Tool command that the installer generates.

Section	Subsection	Description
		Important: The vCenter Server appliance installer does not validate the configuration parameters in the <code>ovftool_arguments</code> subsection. If you set arguments that the OVF Tool does not recognize, the deployment might fail.
	<code>temporary_network</code>	Contains the configuration parameters that describe the temporary network settings for the new appliance. See Table 69: Configuration Parameters in the new_vcsa Section, temporary_network Subsection
	<code>user_options</code>	Contains only the <code>vcdb_migrateSet</code> configuration parameter to set the types of data that you want to transfer from the old appliance to the new appliance. See Table 70: Configuration Parameters in the new_vcsa Section, user_options Subsection
<code>source_vc</code> - describes the existing appliance that you want to upgrade.	<code>managing_esxi_or_vc</code>	Contains the configuration parameters that describe the source ESXi host or vCenter Server appliance instance on which resides the appliance that you want to upgrade. See Table 72: Configuration Parameters in the source_vc Section, managing_esxi_or_vc Subsection .
	<code>vc_vcsa</code>	Contains the configuration parameters that describe the source appliance that you want to upgrade. See Table 73: Configuration Parameters in the source_vc Section, vc_vcsa Subsection .
<code>source_vum</code> - describes the source VMware Update Manager instance. Use if you want to automatically run Migration Assistant on the VMware Update Manager instance.	<code>run_migration_assistant</code>	Optional if the source vCenter Server appliance that you want to upgrade is connected to a VMware Update Manager instance that runs on a Windows virtual machine. Use this subsection if you want to automatically run Migration Assistant on the source VMware Update Manager instance. Contains the configuration parameters that describe the source VMware Update Manager instance, which will be migrated to the new upgraded vCenter Server appliance. See Table 74: Configuration Parameters in the source_vum Section, run_migration_assistant Subsection . Note: The Migration Assistant uses port 9123 by default. If port 9123 is used by another service on the Update Manager machine, Migration Assistant automatically finds another free port. You cannot set a custom port for Migration Assistant.
<code>ceip</code> - describes joining the VMware Customer Experience Improvement Program (CEIP).	<code>settings</code>	Contains only the <code>ceip_enabled</code> configuration parameter to join or not to join the VMware Customer Experience Improvement Program (CEIP). See Table 75: Configuration Parameters in the ceip Section, settings Subsection . Required only if you are upgrading a vCenter Server appliance with an embedded Platform Services Controller or a Platform Services Controller appliance. Note: If the <code>ceip_enabled</code> configuration parameter is set to <code>true</code> , you must run the CLI deployment command with the <code>--acknowledge-ceip</code> argument. For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i> .

Important: The string values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To set a value that contains a backslash (\) or quotation mark (") character, you must precede the character with the backslash (\) character. For example, `"password": "my\"password"` sets the password my"password, `"image": "G:\\vcsa\\VMware-vCenter-Server-Appliance-8.0.0.XXXX-YYYYYYY_OVF10.ova"` sets the path G:\vcsa\VMware-vCenter-Server-Appliance-8.0.0.XXXX-YYYYYYY_OVF10.ova.

The Boolean values must contain only lowercase characters. Can be either `true` or `false`. For example, `"ssh_enable": false`.

Configuration Parameters in the `new_vcsa` Section

Table 66: Configuration Parameters in the `new_vcsa` Section, `esxi` Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the target ESXi host on which you want to deploy the new appliance.
username	string	A user name with administrative privileges on the target ESXi host, for example, root.
password	string	The password of the user with administrative privileges on the target ESXi host.
deployment_network	string	The name of the network to which to connect the new appliance. The network must be part of the target ESXi host or vCenter Server appliance instance (as identified by the <code>managing_esxi_or_vc</code> configuration parameter) network configuration. Note: The network must be accessible from the source ESXi host or vCenter Server appliance instance (as identified by the <code>managing_esxi_or_vc</code> configuration parameter) on which resides the appliance that you want to upgrade. The network must be also accessible from the client machine from which you are performing the upgrade. Ignored if the target ESXi host has only one network.
datastore	string	The name of the datastore on which to store the virtual machine configuration files and virtual disks of the new appliance. The datastore must be available to the target ESXi host. Note: The datastore must have at least 25 GB of free space.
port	integer	The HTTPS reverse proxy port of the target ESXi host. The default port is 443. Use only if the target ESXi host uses a custom HTTPS reverse proxy port.
ssl_certificate_verification	string	The CLI verifies that a server's security certificate is signed by a Certificate Authority (CA), and establishes a secure connection. If the certificate is self-signed, the CLI stops the upgrade unless you specify one of the following SSL certificate configuration options: Specify the Secure Hash Algorithm 1 (SHA-1) certificate thumbprint. A certificate thumbprint is a hexadecimal string that uniquely identifies a certificate. The thumbprint is calculated from the content of the certificate using a thumbprint algorithm. <pre>"thumbprint": "certificate SHA-1 thumbprint"</pre> Set <code>verification_mode</code> to <code>NONE</code>. <pre>"verification_mode": "NONE"</pre> If you are connecting to a server with a self-signed certificate, and fail to either specify the SHA-1 certificate thumbprint or set the verification mode to <code>NONE</code>, the CLI displays the server's self-signed certificate thumbprint, and prompts you to accept or reject the certificate thumbprint. You may also specify that the CLI ignore the self-signed using the <code>vcsa-deploy upgrade</code> command parameter <code>--no-ssl-certificate-validation</code>. See Syntax of the CLI Upgrade Command.

Table 67: Configuration Parameters in the `new_vcsa` Section, `vc` Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the target vCenter Server instance on which you want to deploy the new appliance.
username	string	vCenter Single Sign-On administrator user name on the target vCenter Server instance, for example, administrator@vsphere.local.

Name	Type	Description
password	string	The password of the vCenter Single Sign-On administrator user on the target vCenter Server instance.
deployment_network	string	The name of the network to which to connect the new appliance. The network must part of the target ESXi host or DRS cluster network configuration. Note: The network must be accessible from the source ESXi host on which resides the appliance that you want to upgrade. The network must be also accessible from the client machine from which you are performing the upgrade. Ignored if the target ESXi host or DRS cluster has only one network.
datacenter	array	The vCenter Server data center that contains the target ESXi host or DRS cluster on which you want to deploy the new appliance. If the data center is located in a folder or a structure of folders, specify the as a comma-separated list of strings. For example, ["parent_folder", "child_folder", "datacenter_name"] If there is no folder path on the data center, then use only the data center name. For example, ["datacenter_name"] or "datacenter_name" Note: The value is case-sensitive.
datastore	string	The name of the datastore that you want to store all virtual machine configuration files and virtual disks of the new appliance. Note: The datastore must be available to the target ESXi host or DRS cluster. The datastore must have at least 25 GB of free space.
port	integer	The HTTPS reverse proxy port of the target vCenter Server instance. The default port is 443. Use only if the target vCenter Server instance uses a custom HTTPS reverse proxy port.
target	array	The target cluster, ESXi host, or resource pool on which you want to deploy the new appliance. This is the target you specified with the <code>datacenter</code> parameter. This path must end with a cluster name, ESXi hostname, or resource pool name. Important: You must provide the name that is displayed in the vCenter Server inventory. For example, if the name of the target ESXi host is an IP address in the vCenter Server inventory, you cannot provide an FQDN. Note: All values are case-sensitive. If you want the deployed appliance to be listed in a different location within the data center's hierarchy, use the <code>vm_folder</code> parameter described later in this section. If the target cluster, ESXi host, or resource pool is located in a folder or a structure of folders, specify the value as a comma-separated list of strings. For example, ["parent_folder", "child_folder", "esxi-host.domain.com"] If the target ESXi host is part of a cluster, specify the path as a comma-separated list of strings. For example, ["cluster_name", "esxi-host.domain.com"] If you are deploying to a resource pool, include the label <code>Resources</code> before the resource pool name. For example: ["cluster_name", "Resources", "resource_pool_name"] Note: Pre-checks verify only the memory of the resource pool.
vm_folder	string	Optional. The name of the virtual machine (VM) folder to which to add the new appliance.

Table 68: Configuration Parameters in the `new_vcsa` Section, `appliance` Subsection

Name	Type	Description
<code>thin_disk_mode</code>	Boolean	Set to <code>true</code> to deploy the new appliance with thin virtual disks.
<code>deployment_option</code>	string	The size for the new appliance. Note: You must consider the database size of the appliance that you want to upgrade. For an external database, see Determine the Oracle Database Size and the Storage Size for the New Appliance. - Set to <code>tiny</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the default storage size. Deploys an appliance with 2 CPUs, 10 GB of memory, and 300 GB of storage. - Set to <code>tiny-lstorage</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the large storage size. Deploys an appliance with 2 CPUs, 10 GB of memory, and 825 GB of storage. - Set to <code>tiny-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the x-large storage size. Deploys an appliance with 2 CPUs, 10 GB of memory, and 1700 GB of storage. - Set to <code>small</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the default storage size. Deploys an appliance with 4 CPUs, 16 GB of memory, and 340 GB of storage. - Set to <code>small-lstorage</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the large storage size. Deploys an appliance with 4 CPUs, 16 GB of memory, and 870 GB of storage. - Set to <code>small-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the x-large storage size. Deploys an appliance with 4 CPUs, 16 GB of memory, and 1750 GB of storage. - Set to <code>medium</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the default storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 525 GB of storage. - Set to <code>medium-lstorage</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the large storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 1025 GB of storage. - Set to <code>medium-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the x-large storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 1905 GB of storage. - Set to <code>large</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the default storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 740 GB of storage. - Set to <code>large-lstorage</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the large storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 1090 GB of storage. - Set to <code>large-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the x-large storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 1970 GB of storage. - Set to <code>xlarge</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the default storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 1180 GB of storage. - Set to <code>xlarge-lstorage</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the large storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 1230 GB of storage.

Name	Type	Description
		- Set to <code>xlarge-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the x-large storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 2110 GB of storage. - Set to <code>management-tiny</code> if you want to deploy a vCenter Server appliance with an external Platform Services Controller for up to 10 hosts and 100 virtual machines with the default storage size. Deploys an appliance with 2 CPUs, 10 GB of memory, and 300 GB of storage.
<code>image</code>	string	Optional. A local file path or URL to the vCenter Server appliance installation package. By default the installer uses the installation package that is included in the ISO file, in the <code>vcsa</code> folder.
<code>name</code>	string	The VM name for the new appliance. Must contain only ASCII characters except a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.
<code>ovftool_path</code>	string	Optional. A local file path to the OVF Tool executable file. By default the installer uses the OVF Tool instance that is included in the ISO file, in the <code>vcsa/ovftool</code> folder.
<code>license</code>	string	Optional. License key to be applied to the appliance. If the license key is not provided, evaluation license will be used.

Table 69: Configuration Parameters in the `new_vcsa` Section, `temporary_network` Subsection

Name	Type	Description
<code>ip_family</code>	string	IP version for the temporary network of the new appliance. Set to <code>ipv4</code> or <code>ipv6</code> .
<code>mode</code>	string	IP assignment for the temporary network of the new appliance. Set to <code>static</code> or <code>dhcp</code> .
<code>ip</code>	string	Temporary IP address for the new appliance. Required only if you use a static assignment, that is, if you set the <code>mode</code> parameter to <code>static</code> . You must set an IPv4 or IPv6 address that corresponds to the temporary network IP version, that is, to the value of the <code>ip_family</code> parameter. An IPv4 address must comply with the RFC 790 guidelines. An IPv6 address must comply with the RFC 2373 guidelines.
<code>dns_servers</code>	string or array	IP addresses of one or more DNS servers for the temporary network of the new appliance. To set more than one DNS server, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example, <code>["x.y.z.a", "x.y.z.b"]</code> or <code>"x.y.z.a, x.y.z.b"</code> Required only if you use static network mode for the temporary IP address allocation, that is, if you set the <code>mode</code> parameter to <code>static</code> .
<code>prefix</code>	string	Network prefix length for the temporary network of the new appliance. Use only if the <code>mode</code> parameter is set to <code>static</code> . Remove if the <code>mode</code> parameter is set to <code>dhcp</code> . The network prefix length is the number of bits that are set in the subnet mask. For example, if the subnet mask is 255.255.255.0, there are 24 bits in the binary version of the prefix length, so the network prefix length is 24. For IPv4 version, the value must be between 0 and 32.

Name	Type	Description
		For IPv6 version, the value must be between 0 and 128 .
gateway	string	IP address of the default gateway for the temporary network of the new appliance. For IPv6 version, the value can be default .

Table 70: Configuration Parameters in the `new_vcsa` Section, `user_options` Subsection

Name	Type	Description
vcdb_migrateSet	string	Select the types of data to migrate from the old appliance to the new appliance. Data is copied from the source vCenter Server to the target server. The original source of the data remains unchanged. - Set to <code>core</code> if you want to transfer only the configuration data. This option provides the fastest data migration, keeping system downtime to a minimum. - Set to <code>core_events_tasks</code> if you want to transfer the configuration and historical data (events and tasks) immediately. vCenter Server will not start until all data has been migrated from the source vCenter Server appliance. - Set to <code>all</code> if you want to transfer the configuration, historical, and performance metrics data immediately. vCenter Server will not start until all data has been migrated from the source vCenter Server for Windows. This option transfers the largest amount of data, and requires more downtime than other data migration options. Note: To minimize the upgrade time and the amount of storage required for the new vCenter Server appliance, use the <code>core</code> value. For more information on the types of data you can transfer from your existing vCenter Server to the new, upgraded vCenter Server, see Transferring Data from an Existing vCenter Server Appliance.

Table 71: Configuration Parameters in the `new_vcsa` Section, `sso` Subsection

Name	Type	Description
first_instance	string	Set <code>first_instance</code> to <code>true</code> if this is the first converge upgrade operation in your vCenter Server SSO domain, or if you do not want to set the replication partner for this node. If you want to set the replication partner for this node, set <code>first_instance</code> to <code>false</code> and provide the value of the <code>replication_partner_hostname</code> as the FQDN of a previously converged node which is currently in an embedded vCenter Server topology. For more information on specifying the <code>first_instance</code> and <code>replication_partner_hostname</code> parameters, see About Upgrading a vCenter Server with an External Platform Services Controller Using the CLI.
replication_partner_hostname	string	Host name of the replication partner. Remove if it is the <code>first_instance</code> .

Configuration Parameters in the `source_vc` Section**Table 72: Configuration Parameters in the `source_vc` Section, `managing_esxi_or_vc` Subsection**

Name	Type	Description
hostname	string	The IP address or FQDN of the source ESXi or vCenter Server host on which resides the appliance that you want to upgrade.

Name	Type	Description
username	string	A user name with administrative privileges on the source ESXi host, for example, root.
password	string	The password of the user with administrative privileges on the source ESXi host.
port	integer	The HTTPS reverse proxy port of the source ESXi host. The default port is 443. Use only if the source ESXi host uses a custom HTTPS reverse proxy port.

Table 73: Configuration Parameters in the `source_vc` Section, `vc_vcsa` Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the source appliance that you want to upgrade.
username	string	vCenter Single Sign-On administrator user on the source appliance, for example administrator@vsphere.local. Important: The user must be administrator@ <i>your_domain_name</i> .
password	string	The password of the vCenter Single Sign-On administrator user on the source appliance.
root_password	string	The password for the root user of the operating system of the source appliance.
export_dir	string	Directory to export source configuration and data.

Configuration Parameters in the `source_vum` Section**Table 74: Configuration Parameters in the `source_vum` Section, `run_migration_assistant` Subsection**

Name	Type	Description
esxi_hostname	string	The IP address or FQDN of the ESXi host on which resides the source VMware Update Manager instance. If an FQDN is provided, it must be resolvable from the client machine from which you run the upgrade.
esxi_username	string	A user name with administrative privileges on the ESXi host, for example, root.
esxi_password	string	The password of the user with administrative privileges on the ESXi host.
esxi_port	string	The HTTPS reverse proxy port of the ESXi host. The default port is 443. Use only if the ESXi host uses a custom HTTPS reverse proxy port.
vum_hostname	string	The IP address or FQDN of the Windows virtual machine on which the source VMware Update Manager instance runs. If an FQDN is provided, it must be resolvable from the client machine from which you run the upgrade.
vum_os_username	string	The administrator user name of the Windows virtual machine on which the source VMware Update Manager instance runs.
vum_os_password	string	The administrator password of the Windows virtual machine on which the source VMware Update Manager instance runs. If not provided, you are prompted to enter the password at the command console during the template verification.
export_dir	string	Directory to export source configuration and data.

Configuration Parameters in the `ceip` Section

Table 75: Configuration Parameters in the `ceip` Section, `settings` Subsection

Name	Type	Description
<code>ceip_enabled</code>	Boolean	Set to <code>true</code> to join the CEIP for the new upgraded appliance.

Upgrade a vCenter Server Appliance by Using the CLI

You can use the CLI installer to perform an unattended upgrade of a vCenter Server appliance or Platform Services Controller appliance. You must run the CLI upgrade from a Windows, Linux, or Mac machine that is in the same network as the appliance that you want to upgrade.

- See [Prerequisites for Upgrading the vCenter Server Appliance](#).
- [Prepare Your JSON Configuration File for CLI Upgrade](#).
- Review the arguments for running the CLI upgrade. See [Syntax of the CLI Upgrade Command](#).
- Verify that the user name with which you are logged in to your machine, the path to the vCenter Server appliance installer, the path to your JSON configuration file, and the string values in your JSON configuration file contain only ASCII characters. Extended ASCII and non-ASCII characters are not supported.

1. Navigate to the `vcasa-cli-installer` subdirectory for your operating system.
 - If you are running the upgrade on Windows, navigate to the `vcasa-cli-installer\win32` directory.
 - If you are running the upgrade on Linux, navigate to the `vcasa-cli-installer/lin64` directory.
 - If you are running the upgrade on Mac, navigate to the `vcasa-cli-installer/mac` directory.

2. Optional: Verify that you prepared the upgrade template correctly by running a basic template verification.

```
vcasa-deploy upgrade --verify-template-only path_to_the_json_file
```

3. Optional: Gather and validate the upgrade requirements by running a pre-upgrade check.

```
vcasa-deploy upgrade --precheck-only path_to_the_json_file
```

The pre-upgrade check installs the Upgrade Runner on the source appliance that you want to upgrade without upgrading the appliance.

The Upgrade Runner validates the configurations such as ESXi, network settings, and NTP servers. The Upgrade Runner also checks if you have selected a suitable deployment size and storage size for the new appliance against the compute resources required for the upgrade.

4. Perform the upgrade by running the following command.

```
vcasa-deploy upgrade --accept-eula optional_arguments path_to_the_json_file
```

Use `optional_arguments` to enter space-separated arguments to set additional execution parameters of the upgrade command.

For example, you can set the location of the log and other output files that the installer generates. This example also confirms participation in the VMware Customer Experience Improvement Program (CEIP). If the `ceip_enabled` parameter is set to `true` in the JSON deployment template, you must include the argument `--acknowledge-ceip`.

```
vcasa-deploy upgrade --accept-eula --acknowledge-ceip --log-dir=path_to_the_location path_to_the_json_file
```

Verify Your vCenter Server Upgrade or Migration Is Successful.

Related Links

[Syntax of the CLI Upgrade Command on page 1336](#)

You can use command arguments to set the execution parameters of the upgrade command.

Syntax of the CLI Upgrade Command

You can use command arguments to set the execution parameters of the upgrade command.

You can add a space-separated list of arguments to the CLI upgrade command.

```
vcsa-deploy upgrade path_to_the_json_filelist_of_arguments
```

Argument	Description
--accept-eula	Accepts the end-user license agreement. Required for executing the deployment command.
--acknowledge-ceip	Confirms your acknowledgement of your VMware Customer Experience Improvement Program (CEIP) participation. Required if the <code>ceip_enabled</code> parameter is set to <code>true</code> in the JSON deployment template.
-v, --verbose	Adds debug information to the console output.
-t, --terse	Hides the console output. Displays only warning and error messages.
--log-dir <i>LOG_DIR</i>	Specifies the location of the log and other output files that the installer generates.
--skip-ovftool-verification	Performs basic verification of the configuration parameters and deploys the vCenter Server appliance, but does not validate the OVF Tool parameters in the <code>ovftool_arguments</code> subsection of the JSON template. If you set arguments that the OVF Tool does not recognize, the deployment might fail.
--no-ssl-certificate-verification	Prohibits SSL verification of for all server connections. The CLI verifies that a server's security certificate is signed by a Certificate Authority (CA), and establishes a secure connection. If the certificate is self-signed, the CLI stops the upgrade unless you specify that the CLI ignore the self-signed certificate using the <code>--no-ssl-certificate-verification</code> command parameter. If you are connecting to a server with a self-signed certificate, and fail to specify that the CLI accept it, the CLI displays the server's self-signed certificate thumbprint, and prompts you to accept or reject it. You can also specify that the CLI ignore self-signed certificates using the <code>ssl_certificate_verification</code> configuration parameter in the JSON template. See Upgrade Configuration Parameters for CLI Upgrade of vCenter Server Appliance . Important: Avoid using this option as it may cause problems during or after upgrade due to an un-validated identity of the target host.
--operation-id	Lets you provide an identifier to track the concurrent installation, migration, or upgrade of multiple vCenter Server instances. If do not provide an operation ID, the CLI generates a universally unique identifier (UUID) which you can use to identify the different instances of vCenter Server and their installation or upgrade status.
--pause-on-warnings	Pauses and waits for acknowledgment of warnings.

Argument	Description
<code>--verify-template-only</code>	Performs basic template verification without installing Upgrade Runner, running prechecks, and upgrading or migrating the vCenter Server Appliance.
<code>--precheck-only</code>	Installs Upgrade Runner on the source appliance and runs a complete set of prechecks without performing the upgrade.
<code>-h, --help</code>	Displays the help message for the <code>vcsa-deploy upgrade</code> command.
<code>--template-help</code>	Displays the help message for the use of configuration parameters in the JSON upgrade file.

After the execution finishes, you can get the exit code of the command.

Exit Code	Description
0	Command ran successfully
1	Runtime error
2	Validation error
3	Template error

Related Links

[Upgrade a vCenter Server Appliance by Using the CLI on page 1335](#)

You can use the CLI installer to perform an unattended upgrade of a vCenter Server appliance or Platform Services Controller appliance. You must run the CLI upgrade from a Windows, Linux, or Mac machine that is in the same network as the appliance that you want to upgrade.

Reduced Downtime Upgrade

Reduced downtime upgrade is a migration-based approach to upgrade vCenter Server between major and minor versions that reduces downtime to less than 5 minutes under ideal network, CPU, memory, and storage conditions.

You can perform a reduced downtime upgrade for single self-managed vCenter Server instances, vCenter Server instances managed by another vCenter Server, vCenter Server instances enabled with vCenter HA, and vCenter Server instances connected in Enhanced Linked Mode configuration.

Supported Upgrade Paths

Source Version	Target Version
vSphere 8.0	vSphere 8.0 Update 2 or later
vSphere 8.0 Update 1	vSphere 8.0 Update 2 or later
vSphere 8.0 P02	vSphere 8.0 Update 2 or later
vSphere 8.0 Update 2	Releases later than vSphere 8.0 Update 2
vSphere 8.0 Update 2 with vSphere High Availability (vSphere HA) Note: Reduced downtime upgrade is not available for vCenter HA clusters of versions earlier than vSphere 8.0 Update 2 and for vCenter HA clusters that are manually configured.	vSphere 8.0 Update 3 and later releases.

About the Reduced Downtime Upgrade Process

The reduced downtime upgrade process uses a migration-based approach. In this approach, a new vCenter Server Appliance is deployed and the current vCenter data and configuration is copied to it. The reduced downtime upgrade does not replace the current vCenter Server upgrade process through the GUI installer.

During the reduced downtime upgrade process, the source vCenter Server Appliance and all resources remain online. The downtime occurs when the source vCenter Server Appliance is stopped, the configuration is switched over to the target vCenter, and the services are started. The downtime is expected to take less than 5 minutes under ideal network, CPU, memory, and storage provisioning.

The high-level steps for the reduced downtime upgrade are:

1. Configure the repository URL.
Note: Perform this step when you are upgrading from versions earlier than vSphere 8.0 Update 2 to vSphere 8.0 Update 2.
 Go to Step 2 if you are upgrading from vSphere 8.0 Update 2 to later versions.
2. Download and mount the ISO.
3. Verify that you have a backup of the source vCenter Server Appliance.
4. Upgrade the vCenter Server Life-cycle Manager service plug-in and run prechecks.
5. Configure the target vCenter Server Appliance.
6. Prepare the upgrade and switch over to the target vCenter Server Appliance.

Prechecks are run at every stage of the process. If you face issues, resolve the issue and retry the upgrade steps. If the upgrade is canceled or an error occurs at any stage, the configuration reverts to the source vCenter instance. The source vCenter instance does not shut down until the upgrade process is complete and the target vCenter is up and running.

Note: In the event of a failure, the upgrade process is cancelled automatically and the configuration reverts to the source vCenter instance. The vCenter Lifecycle plug-in does not revert.

Configure Repository URL

You can configure the vCenter Server Appliance to use the default or a custom repository URL as a source for ISO images. By default the current repository for URL-based patching is the default VMware repository URL.

Log in to the vCenter Server Appliance Appliance Management Interface as root.

If vCenter Server is not connected to the Internet or if your security policy does not allow you to connect to the Internet, you can build and configure a custom repository. The custom repository runs on a local Web server within your data center and replicates the data from the default repository.

Note: You must perform this procedure when you are upgrading from versions earlier than vSphere 8.0 Update 2 to vSphere 8.0 Update 2.

1. If you want to configure a custom repository URL, build the repository on your local web server.
2. Log in to Broadcom Support Portal.
3. Download the vCenter Server appliance ISO image. See the Knowledge Base article <https://knowledge.broadcom.com/external/article/366685/vmware-vsphere-downloads-vmware-converte.html>.
4. Confirm that the md5sum is correct by using an MD5 checksum tool.
5. On your Web server, create a repository directory under the root.
 For example, create the `vc_update_repo` directory.
6. Extract the ZIP file into the repository directory.

The extracted files are in the `manifest` and `package-pool` subdirectories.

7. In the vCenter Server Management Interface, click **Update**.

8. Click **Settings**.

9. Select the Repository settings.

Option	Description
Default	Uses the default VMware repository URL that is preset for the build profile of the appliance.
Specified	Uses a custom repository. You must enter the repository URL, for example, <i>https://web_server_name.your_company.com/vc_update_repo</i>. The repository URL must use a secure protocol such as HTTPS or FTPS.

10. Optional: If the specified repository requires authentication, enter the user name and password.

11. If you do not want to perform a security certificate check, deselect the **Check Certificate** check box.

If you trust the URL for the repository, you can choose to bypass the certificate check for the repository URL.

12. Click **SAVE**.

In the **Current version details** pane, you can view the vCenter Server version and build number.

Download and Mount the ISO Image

VMware releases the vCenter Server Appliance ISO image, which contains the installers for the vCenter Server Appliance.

- Create an account at <https://support.broadcom.com>.
- Verify that you have the following privileges:
 - **Virtual machine .Interaction.Configure CD media** on the virtual machine.
 - **Datastore.Browse datastore** on the datastore to which you upload the installation media ISO image.
 - **Datastore.Low level file operations** on the datastore to which you upload the installation media ISO image.
- Download the ISO image to the content library or datastore connected to the vCenter virtual machine or the managing vCenter.
- Mount the ISO image on the CD-ROM device of the vCenter Server virtual machine.

1. Log in to Broadcom Support Portal

2. Download the vCenter Server Appliance ISO image. See the Knowledge Base article <https://knowledge.broadcom.com/external/article/366685/vmware-vsphere-downloads-vmware-converte.html>.

3. Right-click a virtual machine in the vCenter Server inventory and select **Edit Settings**.

4. Expand **CD/DVD drive**, and select **Datastore ISO File** from the drop-down menu.

The **Select File** dialog box opens.

5. Browse to select the file that you downloaded and click **OK**.
6. From the **Virtual Device Node** drop-down menu, select the node that the drive uses in the virtual machine.
7. Select **Connect At Power On** to connect the device when the virtual machine powers on.
8. Click **OK**.
9. Expand the **VM Hardware** panel and click the **Connected** icon next to the datastore ISO file to connect the device.

Upgrade the vCenter Lifecycle Plug-in

After you download and mount the ISO image, upgrade the vCenter Lifecycle plug-in.

- Verify that you have downloaded and mounted the ISO image. See [Download and Mount the ISO Image](#).
 - Log in to the vCenter Server Management Interface as root.
 - Concurrent upgrades are not supported.
 - Perform steps 3-5 if you are upgrading from a version earlier than vSphere 8.0 Update 2.
 - If you are upgrading from version vSphere 8.0 Update 2 or later versions, you can mount the ISO file and upgrade the plug-in.
1. In the vSphere Client, select the vCenter Server that you want to upgrade.
 2. On the **Updates** tab, click **Update Planner**.
 3. Click **Target Version** and **SELECT VERSION**.
 4. Select the version and click **OK**.
 5. Click **NEXT**.
 6. In **Backup**, select the check box to indicate that you have backed up the vCenter Server.
The backups are listed.
 7. In **Upgrade Plug-in**, select **UPGRADE PLUG-IN** to update the vCenter Lifecycle service plug-in.
After the plug-in upgrades, the UI reloads with the version from the ISO file.
 8. Click **NEXT** to proceed with configuring the target vCenter Server Appliance.
Warning: Clicking **NEXT** before the upgrade completes might result in crash of the plug-in. Wait for the progress bar to reach 100% before you click **NEXT**.
 9. Click **RUN PRECHECKS** to run source prechecks and click **NEXT**.
 10. Optional: Click **DISCARD** to discard the upgrade workflow and return to the mount the ISO step. The vCenter Lifecycle plug-in does not revert.

Configure the Target vCenter Server Appliance

When the vCenter Lifecycle service plug-in update completes, configure the target vCenter Server Appliance.

After the upgrade, the target vCenter Server has the same IP address and identity as the source vCenter Server.

1. In **Upgrade**, go to the **Target Appliance** section.
2. Click **CONFIGURE TARGET APPLIANCE**.
The **Target VM deployment** wizard appears.

3. Read and accept the license agreement, and click **Next**.
4. Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.
For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.
5. Click **Next**.
6. Enter location details if the source vCenter Server is managed by another vCenter Server.
Enter the container location, HTTPS port, user name, and password of the source vCenter Server.
7. View and accept the certificate of the source vCenter Server and click **Next**.
8. Select where you want to deploy the target vCenter Server Appliance and click **Next**.
You can deploy the target vCenter Server Appliance on the same location as the source appliance or choose a different location. If you select a different location, the target location must be a different vCenter Server and you must enter the connection details such as, host name or IP address, HTTPS port, user name, and password of the vCenter Server.
The **Deployment Type** page is displayed if you choose to deploy the target on the same location as the source vCenter Server Appliance.
9. If you select a different location for the target vCenter Server, view and accept the certificate and click **Next**.
10. Select the deployment type of the target vCenter Server Appliance and click **Next**.
Select **Same configuration** to deploy the target vCenter Server Appliance with same configuration as the source.
Select **Detailed configuration** to enter all the configuration parameters for the target appliance.
 - a) Enter the data center or virtual machine folder where you want to deploy the target vCenter Server Appliance.
 - b) Select the compute resource required for the target deployment.
 - c) Click **Next**.
11. Configure the target VM appliance details and click **Next**.
Enter the VM name and set and confirm the temporary root password.
12. Select a deployment size and storage size for the target vCenter Server and click **Next**.
13. Select a datastore for the deployment and click **Next**.
Optionally, enable thin provisioning.
14. In the **Network Settings** page, enter a temporary IP address of the target vCenter Server for static deployments.
This IP address is used during the deployment of the target vCenter Server Appliance, during upgrade and for switchover.
If the temporary IP address mode is DHCP, you do not need to enter the IP address.
15. In the **Review** page, review the configuration and click **FINISH**.
The **Target Appliance** section under the **Upgrade** window appears.
16. To change the settings, click **Edit configuration** and make the changes required.
17. Click **Next** to proceed with the upgrade and switchover.
18. Click **DISCARD** to discard the upgrade workflow and return to the mount the ISO step. The vCenter Lifecycle plug-in does not revert.
You cannot discard the configuration when the upgrade is progress.

Prepare Upgrade and Switch Over to the Target vCenter Server Appliance

Once the target vCenter Server Appliance is configured, you can upgrade and switch over the services to it.

The upgrade and switchover happen in two stages. The first stage is the preparation stage where the upgrade is initialized and the target vCenter Server Appliance is created with the configuration that you entered. During this stage, data from the source machine is replicated on the target machine and the time taken for the process depends on the environment and vCenter usage. The second stage is the switchover stage where the target vCenter Server Appliance takes the identity of the source. During the switchover, there is a short downtime. You can select manual or automatic switchover.

1. In **Update Planner**, go to the **Upgrade** section.
2. Select the type of switchover you want to perform.
 - **Manual Switchover.** Select this option if want to manually switch over to the target vCenter Server Appliance. You must click **SWITCHOVER** to trigger the switchover after the target appliance is created with the selected configuration.
 - **Automated Switchover.** Select this option to continue with an automatic switchover after the target appliance is created with the selected configuration.
3. Click **START UPGRADE**

You can see the progress of the prepare stage. When the prepare stage completes and switchover starts and you can see the status of the switchover stage.

After the switchover completes, you can log in to the target vCenter Server Appliance.
4. Click **OPEN VSPHERE CLIENT** to log in.

If the upgrade fails, the configuration reverts to the source vCenter Server Appliance.

Migrating vCenter Server for Windows to vCenter Server Appliance

You can migrate a vCenter Server installation on Windows to a vCenter Server appliance installation while upgrading to version 8.0.


How to Migrate from vCenter Server on Windows to vCenter Server Appliance

VMware provides supported paths for migrating from vCenter Server version 6.7 installations on Windows to vCenter Server Appliance 8.0 installations.

You can migrate the following deployments:

Table 76: Supported vSphere Migration Paths

Source Configuration	Target Configuration
vCenter Server 6.7 with an embedded Platform Services Controller instance on Windows	vCenter Server Appliance 8.0
vCenter Server 6.7 instance on Windows	

Figure 48: High-level Tasks for vCenter Server on Windows Migration to vCenter Server Appliance8.0

You can use the GUI method or the CLI method to migrate your vCenter Server installation from Windows to an appliance.

- [GUI Migration of vCenter Server with an Embedded Platform Services Controller to an vCenter Server Appliance](#)

- [GUI Migration of vCenter Server with an External Platform Services Controller to an Appliance](#)
- [CLI Migration of a vCenter Server Installation from Windows to an Appliance](#)

Important: You cannot change your deployment type during migration.

Migration of Update Manager from Windows to a vCenter Server Appliance 8.0

For vSphere 6.7 and later releases, Update Manager is delivered as a 64-bit application and can be installed only on a 64-bit Windows operating system. In vSphere 6.7, Update Manager is provided as an optional service in vCenter Server Appliance 6.7. VMware provides supported paths for migrating Update Manager from a Windows operating system to a vCenter Server appliance version 8.0. In vSphere 8.0, Update Manager functionality is delivered as part of vSphere Lifecycle Manager.

You can migrate Update Manager in the following vCenter Server deployments:

Table 77: Supported Migration Paths for Update Manager That Runs on Windows to a vCenter Server Appliance

Source Configuration	Target Configuration
vCenter Server and Update Manager run on the same Windows machine	vCenter Server appliance 8.0 with embedded vSphere Lifecycle Manager
vCenter Server and Update Manager run on different Windows machines	vCenter Server appliance 8.0 with embedded vSphere Lifecycle Manager
Update Manager run on a Windows machine and is connected to a vCenter Server Appliance	vCenter Server appliance 8.0 with embedded vSphere Lifecycle Manager

You can use the GUI method or the CLI method to migrate your vCenter Server deployment that uses external Update Manager instance. If you use the GUI method, perform manual steps on the Update Manager Windows system. If you use the CLI method, add configuration parameters about Update Manager in your JSON template.

Important: Verify that the Update Manager source machine does not run additional extensions that are connected to other vCenter Server systems, which are not part of your migration.

Before the migration, Update Manager might use any of the supported Microsoft SQL Server, Oracle, or embedded database solutions. After the migration to the vCenter Server appliance, Update Manager is upgraded to embedded vSphere Lifecycle Manager, which uses the PostgreSQL database.

System Requirements for Migrating vCenter Server Deployments to vCenter Server Appliance Deployments

Your source and target systems must meet specific software and hardware requirements before you can migrate a vCenter Server, vCenter Single Sign-On, or Platform Services Controller deployment to a vCenter Server appliance.

Source System

- Synchronize the clocks on all machines running the source vCenter Server services. See [Synchronizing Clocks on the vSphere Network](#).
- Verify that your vCenter Server and Platform Services Controller certificates are valid for the vCenter Server or Platform Services Controller and have not expired.
- Verify that the system network name of the machines running the target vCenter Server services are valid, and are reachable from other machines in the network.
- Verify that the host name of the virtual machine or physical server from which you are migrating vCenter Server complies with RFC 1123 guidelines.

- If your vCenter Server service is running in a user account other than the Local System account, verify that the user account in which the vCenter Server service is running has the following permissions:
 - **Member of the Administrators group**
 - **Log on as a service**
 - **Act as part of the operating system (if the user is a domain user)**
 - **Replace a process level token**
 - Verify that the LOCAL SERVICE account has read permission on the folder in which vCenter Server is installed and on the HKLM registry.
 - Verify that the connection between the virtual machine or physical server and the domain controller is working.
 - Verify that the source vCenter Server instance or Platform Services Controller instance on Windows does not use a DHCP IP address as its system network name.
- Important:** Migration from a source Windows machine using a DHCP IP Address as its system network name to an appliance is not supported.

Target System

- Your target system must meet specific software and hardware requirements for vCenter Server appliance. See [System Requirements for the vCenter Server Appliance](#).
- When you use Fully Qualified Domain Names, make sure that the machine you use for deploying the vCenter Server appliance and the target ESXi host or vCenter Server instance are on the same DNS server.
- Synchronize the clocks of all target virtual machines on the vSphere network before beginning migration. Unsynchronized clocks might result in authentication problems and can cause the migration to fail or prevent the vCenter Server services from starting. See [Synchronizing Clocks on the vSphere Network](#).

Pre-migration Checks

When you migrate vCenter Server for Windows to vCenter Server Appliance, the upgrade installer performs a pre-check of the environment to ensure it meets requirements. For example, the pre-check verifies that enough space is available on the virtual machine or physical server where you are migrating, and that the external database, if any, can be successfully accessed.

Source Environment Checks

When you migrate vCenter Server for Windows versions 7.0, vCenter Single Sign-On is included as part of the Platform Services Controller. When you provide the information about the vCenter Single Sign-On service, the installer uses the administrator account to check the host name and password, to verify that the details of the vCenter Single Sign-On server you provided can be authenticated before proceeding with the migration process.

The pre-migration checker performs checks for the following aspects of the source environment:

- vCenter Server or Platform Services Controller to verify that migration is supported
- SSL certificates validity and compatibility with system names
- Network connections
- DNS resolution
- Internal and external ports used
- External database connectivity
- Administrator privileges on the Windows machine
- Required disk space for exporting configuration data
- NTP server validation
- Any credentials that you enter

Target Environment Checks

The pre-migration checker performs checks for the following aspects of the target environment:

- Minimum processor requirements
- Minimum memory requirements
- Minimum disk space requirements
- Administrator privileges on the target host
- Any credentials that you enter

Known Limitations

The current release has several known limitations.

The following list contains features or actions that are currently not supported:

- Local Windows OS users and groups are not migrated to the Photon OS of the vCenter Server 8.0. If you assigned vCenter Server permissions to any Local Windows OS users and groups, remove the permissions assignments before the migration. You can re-create Local OS users and groups on the Photon OS of the vCenter Server 8.0 after the migration.
- After the migration, the source vCenter Server is turned off and cannot be turned on to avoid network ID conflicts with the target vCenter Server appliance. After the source vCenter Server is turned off, all solutions that are installed on the source vCenter Server and that are not migrated become unavailable.
- Migration of deployments that use custom ports for services other than Auto Deploy, Update Manager, vSphere ESXi Dump Collector, and HTTP reverse proxy (RHTTP) are not supported.
- The migration process migrates only one network adapter setting to the target vCenter Server appliance. If the hostname of the source vCenter Server resolves to multiple IP addresses across multiple network adapters, you can select which IP address and network adapter settings to migrate. You cannot add the rest of the network adapters and settings to the target vCenter Server appliance.

Preparing for Migration

Before beginning to migrate any type of vCenter Server deployment to an appliance, you must complete the preparation tasks.

Preparation tasks:

- [Synchronize ESXi Clocks with a Network Time Server](#)
- [Preparing vCenter Server Databases for Migration](#)
- [Prepare Managed ESXi Hosts for Migration](#)
- [Download and Mount the vCenter Server Installer](#)
- [Download and Run VMware Migration Assistant on the Source Windows Machine](#)

Synchronize ESXi Clocks with a Network Time Server

Before you install vCenter Server, make sure all machines on your vSphere network have their clocks synchronized.

This task explains how to set up NTP from the VMware Host Client.

1. Start the VMware Host Client, and connect to the ESXi host.
2. Click **Manage**.
3. Under **System**, click **Time & date**, and click **Edit settings**.
4. Select **Use Network Time Protocol (enable NTP client)**.
5. In the NTP servers text box, enter the IP address or fully qualified domain name of one or more NTP servers to synchronize with.
6. From the **NTP Service Start-up Policy** drop-down menu, select **Start and stop with host**.
7. Click **Save**.

The host synchronizes with the NTP server.

Related Links

[Configuring Time Synchronization Settings in vCenter Server on page 1779](#)

You can change the time synchronization settings in vCenter Server after deployment.

Preparing vCenter Server Databases for Migration

The vCenter Server appliance instance requires a database to store and organize server data. Ensure your source vCenter Server database is prepared for migration to the target vCenter Server appliance.

Each vCenter Server appliance instance must have its own database. The bundled PostgreSQL database that is included in the vCenter Server appliance supports up to 2,500 hosts and 30,000 virtual machines.

To ensure your database is prepared for migration:

- Verify that passwords are current and not set to expire soon.
- Reduce the database size. For more information, see Knowledge Base article [KB 2110031](#).
- Verify that you have backed up your database. See your database documentation.
- Verify that vCenter Server can communicate with the local database.

During the migration of vCenter Server to vCenter Server appliance, the installer:

1. Exports the vCenter Server database.
2. Deploys the target vCenter Server appliance in an unconfigured state.
3. Copies exported data to the target vCenter Server appliance.
4. Starts the PostgreSQL service to import the source database data.
5. Upgrades the database schema to be compatible with the target vCenter Server appliance.
6. Starts the target vCenter Server appliance services.

When you configure the target vCenter Server appliance, you initialize and configure using the imported database with the old schema. You have a choice of migration options:

1. Inventory tables
2. Inventory tables with events and tasks
3. All database data

Prepare an Oracle Database for Migration

Ensure that you have the necessary credentials, and that you complete any necessary cleanup or other preparation before migrating your Oracle database from Windows to an embedded PostgreSQL database in the appliance.

Verify that you have confirmed basic interoperability before preparing your Oracle database for migration.

Verify that you have backed up your database. For information about backing up the vCenter Server database, see the Oracle documentation.

1. Verify that passwords are current and not set to expire soon.
2. Ensure that you have login credentials, the database name, and the database server name that the vCenter Server database is to use.
Look in the ODBC system for the connection name of the database source name for the vCenter Server database.
3. Use the Oracle SERVICE_NAME instead of SID to verify that your Oracle database instance is available.
 - Log in to the database server to read from the alert log: `$ORACLE_BASE/diag/rdbms/$instance_name/$INSTANCE_NAME/trace/alert_$ INSTANCE_NAME.log`.
 - Log in to the database server to read from the Oracle Listener status output.
 - If you have the SQL*Plus client installed, you can use `tnsping` for the vCenter Database instance. If the `tnsping` command does not work the first time, retry it after waiting a few minutes. If retrying does not work, restart the vCenter Database instance on the Oracle server and then retry `tnsping` to ensure it is available.
4. Verify that the JDBC driver file is included in the CLASSPATH variable.
5. Verify that permissions are set correctly.
6. Either assign the DBA role or grant the required permissions to the user.
7. Make a full backup of the vCenter Server database.

Your database is prepared for the vCenter Server migration to vCenter Server Appliance.

Prepare a Microsoft SQL Server Database for Migration

Ensure that you have the necessary credentials, and that you complete any necessary cleanup or other preparation before migrating your Microsoft SQL Server database on Windows to an embedded PostgreSQL database appliance.

Verify that you have backed up your database. For information about backing up the vCenter Server database, see the Microsoft SQL Server documentation.

Important: You cannot use Integrate Windows for your authentication method if the vCenter Server service is running under the Microsoft Windows built-in system account.

1. Verify that passwords are current and not set to expire soon.
2. Verify that JDK 1.6 or later is installed on the vCenter Server machine.
3. Verify that the `sqljdbc4.jar` file is added to the CLASSPATH variable on the machine where vCenter Server Appliance is to be migrated.
If the `sqljdbc4.jar` file is not installed on your system, the vCenter Server Appliance installer installs it.
4. Verify that your system database source name is using the Microsoft SQL Server Native Client 10 or 11 driver.
5. Make a full backup of the vCenter Server database.

Your database is prepared for the vCenter Server migration to vCenter Server Appliance.

Prepare PostgreSQL Database Before Migrating vCenter Server to an Appliance

Ensure that you have the necessary credentials, and that you complete any necessary cleanup or other preparation before migrating your PostgreSQL database installation on Windows to an appliance.

Verify that you have confirmed basic migration interoperability before preparing your PostgreSQL database for migrating vCenter Server.

For information about backing up the vCenter Server database, see the PostgreSQL documentation.

1. Verify that passwords are current and not set to expire soon.
2. For vCenter Server, locate the `cleanup_orphaned_data_PostgresSQL.sql` script in the ISO image and copy it to your PostgreSQL server.
3. Log in to vCenter Server Appliance as root user.
4. Run the cleanup script.
`/opt/vmware/vpostgres/9.4/bin/psql -U postgres -d VCDB -f path cleanup_orphaned_data_Postgres.sql`
 The cleanup script cleans and purges any unnecessary or orphaned data in your vCenter Server database that is not used by any vCenter Server component.
5. Make a full backup of the vCenter Server database.

Your database is prepared for the vCenter Server migration to vCenter Server Appliance.

Prepare Managed ESXi Hosts for Migration

You must prepare the ESXi hosts managed by your vCenter Server installation before migrating it from Windows to vCenter Server8.0.

To migrate vCenter Server or an external Platform Services Controller from Windows to vCenter Server8.0, your source and target ESXi hosts must meet the migration requirements.

- ESXi hosts must be at version 6.7 or greater. For information on ESXi compatibility, see the [VMware Compatibility Guide](#).
 - ESXi hosts must not be in lockdown or maintenance mode.
1. To keep your current SSL certificates, back up the SSL certificates that are on the vCenter Server system before you upgrade to vCenter Server8.0.
 The default location of the SSL certificates is `%allusersprofile%\VMware\VMware VirtualCenter`.
 2. If you have Custom or Thumbprint certificates, see [ESXi Host Upgrades and Certificates](#) to determine your preparatory steps.
 3. If you have vSphere HA clusters, SSL certificate checking must be enabled.
 If you have vSphere HA clusters, SSL certificate checking must be enabled.
 - a) Select the vCenter Server instance in the vSphere Client inventory tree.
 - b) Select the **Manage** tab and the **General** subtab.
 - c) Verify that **vCenter Server requires verified host SSL certificates** is selected.

Your ESXi hosts are ready for migration to the vCenter Server appliance.

Preparing vCenter Server Certificates for Migration

You must verify that your source vCenter Server certificates are prepared before you start the migration process.

In vSphere 6.0 and later certificates are stored in the VMware Endpoint Certificate Store. The migration process proceeds normally and preserves your certificates. For information about vCenter Server certificate locations, see Knowledge Base article [KB 2111411](#).

Certificate Files Location

The vCenter Server certificate files are located at `%ProgramData%\VMware\VMware VirtualCenter\SSL`

Supported Certificate Types

If your environment uses any of the supported certificate types, you can continue with the migration. The migration process proceeds normally and preserves your certificates.

- Your `ruicert.crt` file contains the entire chain including the leaf certificate. You can create this type of certificate by deploying and using the VMware SSL Certificate Automation Tool, see Knowledge Base article [KB 2057340](#).
- Your `ruicert.crt` file contains the leaf certificate and the corresponding `cacert.pem` is available in `%ProgramData%\VMware\VMware VirtualCenter\SSL` to validate the `ruicert.crt`.

Unsupported Certificate Types

If your environment uses any of the unsupported certificate types, you must prepare your certificates before you can proceed with the migration process.

- Your `ruicert.crt` contains only the leaf certificate, the `cacert.pem` is missing or invalid, and `cacert.pem` is not added to the Windows trust store.
Get the Certificate Authority certificate, including all intermediate certificates, and create a `cacert.pem` file, or replace the vCenter Server certificates with any of the supported formats.
- Your `ruicert.crt` contains only the leaf certificate and the `cacert.pem` is missing or invalid, but the `cacert.pem` is added to the Windows trust store.
Get the Certificate Authority certificate, including all intermediate certificates from the Windows trust store and create `cacert.pem`. Use OpenSSL to verify the certificate by running `verify -CAfile cacert.pem ruicert.crt` command.

For more information about vSphere security certificates, see the *vSphere Security* documentation.

System Requirements for the vCenter Server Installer

You can run the vCenter Server GUI or CLI installer from a network client machine that is running on a Windows, Linux, or Mac operating system of a supported version.

To ensure optimal performance of the GUI and CLI installers, use a client machine that meets the minimum hardware requirements.

Table 78: System Requirements for the GUI and CLI Installers

Operating System	Supported Versions	Minimum Hardware Configuration for Optimal Performance
Windows	• Windows 10, 11 • Windows 2016 x64 bit • Windows 2019 x64 bit • Windows 2022 x64 bit	4 GB RAM, 2 CPU having 4 cores with 2.3 GHz, 32 GB hard disk, 1 NIC
Linux	• SUSE 15 • Ubuntu 18.04, 20.04, 21.10	4 GB RAM, 1 CPU having 2 cores with 2.3 GHz, 16 GB hard disk, 1 NIC Note: The CLI installer requires 64-bit OS.
Mac	• macOS 10.15, 11, 12	8 GB RAM, 1 CPU having 4 cores with 2.4 GHz, 150 GB hard disk, 1 NIC

Operating System	Supported Versions	Minimum Hardware Configuration for Optimal Performance
	macOS Catalina, Big Sur, Monterey	

Note: For client machines that run on Mac 10.15 or later, concurrent GUI deployments of multiple appliances are unsupported. You must deploy the appliances in a sequence.

Note: Visual C++ redistributable libraries need to be installed to run the CLI installer on versions of Windows older than Windows 10. The Microsoft installers for these libraries are located in the `vcsa-cli-installer/win32/vcredist` directory.

Note: Deploying the vCenter Server appliance with the GUI requires a minimum resolution of 1024x768 to properly display. Lower resolutions can truncate the UI elements.

Determine the Oracle Database Size and the Storage Size for the New Appliance

Before upgrading a vCenter Server appliance or migrating a vCenter Server on Windows that uses an external Oracle database, you must determine the size of the existing database. Based on the size of the existing database, you can calculate the minimum storage size for the new vCenter Server appliance database using an embedded PostgreSQL database.

You must have the vCenter Server database login credentials.

You run scripts to determine the Oracle core table size, the events and tasks table size, and the statistics table size. The Oracle core table corresponds to the database (`/storage/db`) partition of the PostgreSQL database. The Oracle events and tasks and statistics tables correspond to the statistics, events, alarms, and tasks (`/storage/seat`) partition of the PostgreSQL database.

During the upgrade of the appliance, you must select a storage size for the new appliance that is at least twice the size of the Oracle tables size.

During the upgrade of the appliance, you can select the types of data to transfer to the new appliance. To minimize upgrade time and storage requirements for the new appliance, you can choose to transfer only the configuration data.

1. Log in to a SQL*Plus session with the vCenter Server database user.
2. Determine the core table size by running the following script.

```
SELECT ROUND(SUM(s.bytes)/(1024*1024)) SIZE_MB
FROM   user_segments s
WHERE  (s.segment_name,s.segment_type)
      IN (SELECT seg_name, seg_type FROM
          (SELECT t.table_name seg_name, t.table_name tname,
                 'TABLE' seg_type
           FROM   user_tables t
          UNION
           SELECT i.index_name, i.table_name,
                 'INDEX'
           FROM   user_indexes i
          ) ti
      WHERE (ti.tname LIKE 'VPX_%'
            OR ti.tname LIKE 'CL_%'
            OR ti.tname LIKE 'VDC_%')
            AND ti.tname NOT LIKE 'VPX_SAMPLE_TIME%'
            AND ti.tname NOT LIKE 'VPX_HIST_STAT%'
            AND ti.tname NOT LIKE 'VPX_TOPN%'
            AND ti.tname NOT LIKE 'VPX_SDRS_STATS_VM%')
```

```

AND ti.tname NOT LIKE 'VPX_SDRS_STATS_DATASTORE%'
AND ti.tname NOT LIKE 'VPX_TASK%'
AND ti.tname NOT LIKE 'VPX_EVENT%'
AND ti.tname NOT LIKE 'VPX_PROPERTY_BULLETIN%');

```

The script returns the database storage size in MB.

3. Determine the events and tasks table size by running the following script.

```

SELECT ROUND(SUM(s.bytes)/(1024*1024)) SIZE_MB
FROM   user_segments s
WHERE  (s.segment_name,s.segment_type)
        IN (SELECT seg_name, seg_type FROM
              (SELECT t.table_name seg_name, t.table_name tname,
                'TABLE' seg_type
               FROM   user_tables t
              UNION
               SELECT i.index_name, i.table_name,
                'INDEX'
               FROM   user_indexes i
              ) ti
        WHERE
              ti.tname LIKE 'VPX_TASK%'
            OR ti.tname LIKE 'VPX_EVENT%');

```

The script returns the events and tasks storage size in MB.

4. Determine the statistics table size by running the following script.

```

SELECT ROUND(SUM(s.bytes)/(1024*1024)) SIZE_MB
FROM   user_segments s
WHERE  (s.segment_name,s.segment_type)
        IN (SELECT seg_name, seg_type FROM
              (SELECT t.table_name seg_name, t.table_name tname,
                'TABLE' seg_type
               FROM   user_tables t
              UNION
               SELECT i.index_name, i.table_name,
                'INDEX'
               FROM   user_indexes i
              ) ti
        WHERE
              ti.tname LIKE 'VPX_SAMPLE_TIME%'
            OR ti.tname LIKE 'VPX_TOPN%'
            OR ti.tname LIKE 'VPX_TASK%'
            OR ti.tname LIKE 'VPX_EVENT%'
            OR ti.tname LIKE 'VPX_HIST_STAT%');

```

The script returns the statistics storage size in MB.

5. Calculate the minimum storage size for the new appliance that you are going to deploy during the upgrade.
 - a) The size of the database (`/storage/db`) partition of the embedded PostgreSQL database must be at least twice the size of the Oracle core table returned in [2](#).
 - b) The size of the statistics, events, alarms, and tasks (`/storage/seat`) partition of the embedded PostgreSQL database must be at least twice the sum of the sizes of the Oracle events and tasks and statistics tables returned in [3](#) and [4](#).

For example, if the Oracle core table is 100 MB, the events and tasks table is 1,000 MB, and the statistics table is 2,000 MB, then the Postgres `/storage/db` partition must be at least 200 MB and the `/storage/seat` partition must be at least 6,000 MB.

Related Links

[Storage Requirements for the vCenter Server Appliance on page 1215](#)

When you deploy the vCenter Server appliance, the ESXi host or DRS cluster on which you deploy the appliance must meet minimum storage requirements. The required storage depends not only on the size of the vSphere environment and the storage size, but also on the disk provisioning mode.

Determine the Microsoft SQL Server Database Size and the Storage Size for the New vCenter Server Appliance

Before upgrading a vCenter Server appliance, or migrating a vCenter Server on Windows that uses an external Microsoft SQL Server database, you must determine the size of the existing database. Based on the size of the existing database, you can calculate the minimum storage size for the new vCenter Server appliance. This storage size allows the embedded PostgreSQL database to assume the data from the old database with enough free disk space after the upgrade.

You must have the vCenter Server database login credentials.

You run scripts to determine the Microsoft SQL Server core table size, the events and tasks table size, and the statistics table size. The Microsoft SQL Server core table corresponds to the database (`/storage/db`) partition of the PostgreSQL database. The Microsoft SQL Server events and tasks and statistics tables correspond to the statistics, events, alarms, and tasks (`/storage/seat`) partition of the PostgreSQL database.

During the upgrade of the appliance, you must select a storage size for the new appliance that is at least twice the size of the Microsoft SQL Server tables size.

1. Log in to a SQL Management Studio session with the vCenter Server database user.
2. Determine the core table size by running the following script.

```
SELECT SUM(p.used_page_count * 8)/1024 AS disk_size
FROM sys.dm_db_partition_stats p
JOIN sys.objects o
  ON o.object_id = p.object_id
WHERE o.type_desc = 'USER_TABLE'
  AND o.is_ms_shipped = 0 AND UPPER(o.name) NOT LIKE 'VPX_HIST_STAT%'
  AND UPPER(o.name) NOT LIKE 'VPX_SAMPLE_TIME%'
  AND UPPER(o.name) NOT LIKE 'VPX_TOPN%'
  AND UPPER(o.name) NOT LIKE 'VPX_TASK%'
  AND UPPER(o.name) NOT LIKE 'VPX_EVENT%'
  AND UPPER(o.name) NOT LIKE 'VPX_SDRS_STATS_VM%'
  AND UPPER(o.name) NOT LIKE 'VPX_SDRS_STATS_DATASTORE%'
  AND UPPER(o.name) NOT LIKE 'VPX_PROPERTY_BULLETIN%';
```

The script returns the database storage size in MB.

3. Determine the events and tasks table size by running the following script.

```
SELECT SUM(p.used_page_count * 8)/1024 AS disk_size
```



```

FROM sys.dm_db_partition_stats p
JOIN sys.objects o
    ON o.object_id = p.object_id
WHERE o.type_desc = 'USER_TABLE'
    AND o.is_ms_shipped = 0 AND ( UPPER(o.name) LIKE 'VPX_TASK%'
    OR UPPER(o.name) LIKE 'VPX_EVENT%');

```

The script returns the events and tasks storage size in MB.

4. Determine the statistics table size by running the following script.

```

SELECT SUM(p.used_page_count * 8)/1024 AS disk_size
FROM sys.dm_db_partition_stats p
JOIN sys.objects o
    ON o.object_id = p.object_id
WHERE o.type_desc = 'USER_TABLE'
    AND o.is_ms_shipped = 0
    AND (    UPPER(o.name) LIKE 'VPX_HIST_STAT%'
        OR UPPER(o.name) LIKE 'VPX_SAMPLE_TIME%'
        OR UPPER(o.name) LIKE 'VPX_TOPN%');

```

The script returns the statistics storage size in MB.

5. Calculate the minimum storage size for the new appliance that you are going to deploy during the upgrade.

- a) The size of the database (/storage/db) partition of the embedded PostgreSQL database must be at least twice the size of the Microsoft SQL Server core table returned in [2](#).
- b) The size of the statistics, events, alarms, and tasks (/storage/seat) partition of the embedded PostgreSQL database must be at least twice the sum of the sizes of the Microsoft SQL Server events and tasks and statistics tables returned in [3](#) and [4](#).

For example, if the Microsoft SQL Server core table is 100 MB, the events and tasks table is 1,000 MB, and the statistics table is 2,000 MB, then the Postgres /storage/db partition must be at least 200 MB and the /storage/seat partition must be at least 6,000 MB.

Download and Run VMware Migration Assistant on the Source Windows Machine

You must download and run the VMware Migration Assistant on your source vCenter Server or Platform Services Controller to prepare it for migration from Windows to a vCenter Server appliance. If your deployment of vCenter Server has an external Update Manager that runs on Windows, download and run the VMware Migration Assistant on the source Windows machine. Update Manager prepares the Update Managerserver and database for migration from Windows to the vCenter Server appliance.

- [Download and Mount the vCenter Server Installer](#).
- Log in to the Windows machine as an administrator.

The VMware Migration Assistant performs the following tasks on the source Windows machine where you run it:

1. Discovers the source deployment type.
2. Runs pre-checks on the source.
3. Reports errors that must be addressed before starting the migration.
4. Provides information for the next steps in the migration process.

Ensure that the VMware Migration Assistant window remains open during the migration process. Closing the VMware Migration Assistant causes the migration process to stop.

1. In the vCenter Server appliance installer package, locate the directory that contains VMware Migration Assistant.
2. Copy the VMware Migration Assistant folder to the source Windows machine where either one of the following components runs:
 - Update Manager
 - Platform Services Controller
 - vCenter Server

CAUTION: If Update Manager runs on a different Windows machine than any other of the vCenter Server components that you are migrating, run VMware Migration Assistant on the Update Manager source machine first. If you do not run VMware Migration Assistant on the Update Manager source machine first, the VMware Migration Assistant on vCenter Server may fail.

3. Run the VMware Migration Assistant on the Windows machine.
 - For the GUI, double-click `VMware-Migration-Assistant.exe`
 - For the CLI, enter: `VMware-Migration-Assistant.exe -p <password of Administrator@vmdir.domain>`

To list all the available input parameters, enter: `VMware-Migration-Assistant.exe --help`.

Important: Leave the Migration Assistant window open until you finish the upgrade or the migration process of your vCenter Server deployment.

The VMware Migration Assistant runs pre-upgrade checks and prompts you to resolve any errors it finds before proceeding with the migration.

When the pre-checks are finished and any errors are addressed, your source system is ready for migration.

To begin the migration process, follow the VMware Migration Assistant instructions.

For detailed migration steps, see one of the following.

- [GUI Migration of vCenter Server with an Embedded Platform Services Controller to an vCenter Server Appliance](#)
- [GUI Migration of vCenter Server with an External Platform Services Controller to an Appliance](#)
- [CLI Migration of a vCenter Server Installation from Windows to an Appliance](#)

Prerequisites for Migrating vCenter Server

To ensure the successful migration of vCenter Server, you must perform a set of required tasks and pre-checks prior to running the migration.

General Prerequisites

- [Download and Mount the vCenter Server Installer](#).
- Verify that the clocks of all machines on the vSphere network are synchronized. See [Synchronizing Clocks on the vSphere Network](#).

Target System Prerequisites

- Verify that your system meets the minimum software and hardware requirements. See [System Requirements for the vCenter Server Appliance](#).

- If you plan to deploy the new appliance on an ESXi host, verify that the target ESXi host is not in lockdown or maintenance mode.
- If you plan to deploy the new appliance on an ESXi host, verify that the target ESXi host is not part of a fully automated DRS cluster.
- If you plan to deploy the new appliance on a DRS cluster of the inventory of a vCenter Server instance, verify that the cluster contains at least one ESXi host that is not in lockdown or maintenance mode.
- If you plan to deploy the new appliance on a DRS cluster of the inventory of a vCenter Server instance, verify that the cluster is not fully automated.

Source System Prerequisites

- Verify that the source machine that you want to migrate does not run on an ESXi host that is part of a fully automated DRS cluster.
- Verify that you have sufficient disk space on the source machine that you want to migrate so that you can accommodate the data for the migration.
- Create an image-based backup of the vCenter Server appliance you are migrating as a precaution in case there is a failure during the migration process. If you are migrating a vCenter Server appliance with an external Platform Services Controller, take a image-based backup of the Platform Services Controller appliance as well.

Important: To take a pre-migration image-based backup, power off all the vCenter Server and Platform Services Controller nodes in your environment, and take a backup of each node. After you have taken backups of all the nodes, you can restart them and proceed with the migration procedure.

If the migration fails, delete the newly deployed vCenter Server appliance, and restore the vCenter Server and Platform Services Controller nodes from their respective backups. You must restore all the nodes in the environment from their backups. Failing to do so will cause the replication partners to be out of synchronization with the restored node.

To learn about image-based back, see "Image-Based Backup and Restore of a vCenter Server Environment" in *vCenter Server Installation and Setup*.

- If you use an external database, back up the external database.

Network Prerequisites

- If you plan to assign a static IP address in the temporary network settings of the appliance, verify that you have configured the forward and reverse DNS records for the IP address.
- If you plan to assign a DHCP IP address in the temporary network settings of the new appliance, verify that the ESXi host on which you want to deploy the new appliance is in the same network as the ESXi host on which the existing vCenter Server appliance runs.
- If you plan to assign a DHCP IP address in the temporary network settings of the new appliance, verify that the ESXi host on which you want to deploy the new appliance is connected to at least one network that is associated with a port group which accepts MAC address changes. Note that the default security policy of a distributed virtual switch is to reject MAC address changes. For information about how to configure the security policy for a switch or port group, see *vSphere Networking*.
- Add the source vCenter Server IP address in the DNS records.

Required Information for Migrating vCenter Server from Windows to an Appliance

The vCenter Server migration wizard prompts you for the deployment and migration information when migrating a vCenter Server instance or a Platform Services Controller instance from Windows to an appliance. It is a best practice to keep a record of the values that you entered in case you must power off the appliance and restore the source installation.

You can use this worksheet to record the information that you need for migrating a vCenter Server instance with a Platform Services Controller from Windows to an appliance.

Important: The user name that you use to log in to the machine from which you want run the GUI installer, the path to the vCenter Server Appliance installer, and your values including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

Local OS users existing on source windows machine are not migrated to the target vCenter Server Appliance and must be recreated after migration is complete. If any local OS user names are used to log in to the vCenter Single Sign-On, you must recreate them and reassign permissions in the Platform Services Controller appliance.

If the source vCenter Server machine is joined to an Active Directory domain, the account you use must have permissions to rejoin the machine to the domain. For more information, see <http://kb.vmware.com/kb/2146454>.

Table 79: Information Required for Migrating vCenter Server from Windows to vCenter Server Appliance

Required Information		Default Value	Your Entry
Required source vCenter Server migration data	vCenter Server IP address or FQDN		
	vCenter Single Sign-On administrator user name	administrator@vsphere.local Important: The user must be administrator@your_domain_name.	
	Password of the vCenter Single Sign-On administrator		
	Migration Assistant port number		
	vCenter Server version		
	Temporary upgrade files path	%LOCALAPPDATA%\VMware\Migration-Assistant\export	
	IP address or FQDN of the source ESXi host on which the source vCenter Server resides		
	Source ESXi host user name with administrative rights on the source ESXi host		
	Source ESXi host password		
	Migrate performance & other historical data	Deactivated by default	
	Migration Assistant thumbprint		
	Active Directory administrator credentials		
	Service account credentials, if vCenter Server is running under a customer user account		
Required target vCenter Server Appliance data	IP address or FQDN of the target ESXi host or vCenter Server instance where you deploy the new vCenter Server Appliance to which you migrate the source vCenter Server		

Required Information		Default Value	Your Entry
	User name with administrative privileges for the target ESXi host or vCenter Server instance, data center or data center folder, and resource pool of an ESXi host or DRS cluster to which to migrate the source installation		
	Password for the target ESXi host. or vCenter Server instance, data center or data center folder, and resource pool of an ESXi host or DRS cluster		
	vCenter Single Sign-On username		
	vCenter Single Sign-On password		
	Target vCenter Server Appliance name		
	Password of the root user		
	vCenter Server Appliance size. The options vary depending on the size of your vSphere environment. - Tiny (up to 10 hosts, 100 virtual machines) - Small (up to 100 hosts, 1,000 virtual machines) - Medium (up to 400 hosts, 4,000 virtual machines) - Large (up to 1,000 hosts, 10,000 virtual machines) - X-Large (up to 2,000 hosts, 35,000 virtual machines)	Tiny (up to 10 hosts, 100 virtual machines)	
	Storage size of the vCenter Server Appliance for your vSphere environment. Increase the default storage size if you want larger volume for SEAT data (stats, events, alarms, and tasks). Default For tiny deployment size, deploys the appliance with 250 GB of storage. For small deployment size, deploys the appliance with 290 GB of storage. For medium deployment size, deploys the appliance with 425 GB of storage. For large deployment size, deploys the appliance with 640 GB of storage. For x-large deployment size, deploys the appliance with 980 GB of storage.	Default	

Required Information		Default Value	Your Entry
	- Large For tiny deployment size, deploys the appliance with 775 GB of storage. For small deployment size, deploys the appliance with 820 GB of storage. For medium deployment size, deploys the appliance with 925 GB of storage. For large deployment size, deploys the appliance with 990 GB of storage. For x-large deployment size, deploys the appliance with 1030 GB of storage. - X-Large For tiny deployment size, deploys the appliance with 1650 GB of storage. For small deployment size, deploys the appliance with 1700 GB of storage. For medium deployment size, deploys the appliance with 1805 GB of storage. For large deployment size, deploys the appliance with 1870 GB of storage. For x-large deployment size, deploys the appliance with 1910 GB of storage.		
	Name of the datastore on which the new version of the vCenter Server Appliance is deployed		
	Activate or deactivate thin disk mode.	Deactivated by default	
	Join or do not participate in the VMware Customer Experience Improvement Program (CEIP). For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i> .	Join the CEIP	
Temporary network for communication between the source vCenter Server and the target vCenter Server Appliance	IP address version	IPv4	
	IP address allocation method	DHCP	
Static assignment settings	Network address		
	Subnet mask		
	Network gateway		
	Network DNS servers, separated with commas		
Activate or deactivate SSH		Deactivated by default	

GUI Migration of vCenter Server with an Embedded Platform Services Controller to a vCenter Server Appliance


You can use the GUI method to migrate a vCenter Server instance with an embedded Platform Services Controller to a vCenter Server appliance.

When you migrate from vCenter Server with an embedded Platform Services Controller on Windows to vCenter Server appliance, the entire deployment is migrated in a single step.

If you use Update Manager in the deployment on Windows that you migrate, and Update Manager runs on a separate machine from any other of the components, take an extra step to migrate Update Manager to an appliance.

1. If your deployment on Windows uses an external Update Manager, run Migration Assistant on the Update Manager machine to start the migration of the Update Manager server and database to the Update Manager appliance.
2. Migrate the vCenter Server instance with embedded Platform Services Controller from Windows to a vCenter Server appliance.

Figure 49: vCenter Server 6.7 with Embedded Platform Services Controller Deployment Before and After Migration



GUI tasks for migrating vCenter Server with an embedded Platform Services Controller from Windows to a vCenter Server appliance.

1. [Download and Mount the vCenter Server Installer ISO file](#) on a network virtual machine or physical server from which you want to perform the migration.
2. [Download and Run VMware Migration Assistant on the Source Windows Machine.](#)
Note: If you are migrating a system that uses an external instance of Update Manager that runs on a separate Windows machine, first run Migration Assistant on the Update Manager machine.
3. Assemble the [Required Information for Migrating vCenter Server from Windows to an Appliance.](#)
4. [Deploy the OVA File for Migrating to the Target vCenter Server Appliance.](#)
5. [Set Up the Target vCenter Server Appliance](#)

Important: The user name that you use to log in to the machine from which you want run the GUI installer, the path to the vCenter Server appliance installer, and your values including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

The installer:


- Deploys a new target appliance.
- Exports the required files from the source vCenter Server.
- Copies the required files to the new vCenter Server appliance.
- Runs the migration process on the new vCenter Server appliance as specified in the Summary.
- Imports and updates the files and settings of the source vCenter Server installation to the new vCenter Server appliance.

Deploy the OVA File for Migrating to the Target vCenter Server Appliance

To start the migration process, you use the GUI installer to deploy the OVA file that is included in the installer ISO file as the target vCenter Server appliance.

- See [Prerequisites for Migrating vCenter Server](#)

Figure 50: vCenter Server 6.7 with Embedded Platform Services Controller Deployment Before and After Migration



1. In the vCenter Server installer, navigate to the `vcasa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.

2. On the Home page, click **Migrate**.
3. Review the Introduction page to understand the migration process and click **Next**.
4. Read and accept the license agreement, and click **Next**.
5. On the Connect to source page, enter the details for the source vCenter Server instance, and click **Next**.
 - a) Enter the IP address or FQDN.
 - b) Enter the user name and password of a user who has administrative privileges on the vCenter Server instance, for example, the administrator@your_domain_name user.
 - c) Enter the Migration Assistant Port you received in the Migration Assistant instructions.
6. Optional: Review the warning message and try to resolve the warnings, if any, and click **Yes**.
7. Connect to the target server to which you want to migrate the source vCenter Server.

Option	Steps
You can connect to an ESXi host on which to deploy the target appliance.	1. Enter the FQDN or IP address of the ESXi host. 2. Enter the HTTPS port of the ESXi host. 3. Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the target appliance.	1. Enter the FQDN or IP address of the vCenter Server instance. 2. Enter the HTTPS port of the vCenter Server instance. 3. Enter the user name and password of a vCenter Single Sign-On user with administrative privileges on the vCenter Server instance, for example, the administrator@your_domain_name user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes. 6. Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next. Note: You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode. 7. Select the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next.

8. On the Set up target appliance VM page, enter a name for the target vCenter Server appliance, set the password for the root user, and click **Next**.

The password must contain at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

Important: The local operating system password is not migrated to the target appliance.

9. Select the deployment size for the new vCenter Server appliance for your vSphere inventory.

Note: You cannot select a deployment size that is smaller than the source deployment.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 vCPUs and 14 GB of memory.

Deployment Size Option	Description
	Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 vCPUs and 21 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 vCPUs and 30 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 vCPUs and 39 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 vCPUs and 58 GB of memory. Suitable for environments with up to 2,500 hosts or 45,000 virtual machines

Note: At the bottom of the deployment size table, a row shows the size information of the source machine. This size information is reported by the migration assistant and might help understand why you cannot select some deployment sizes.

10. Select the storage size for the new vCenter Server appliance, and click **Next**.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 579 GB of storage.	Deploys an appliance with 694 GB of storage.	Deploys an appliance with 908 GB of storage.	Deploys an appliance with 1358 GB of storage.	Deploys an appliance with 2283 GB of storage.
Large	Deploys an appliance with 2019 GB of storage.	Deploys an appliance with 2044 GB of storage.	Deploys an appliance with 2208 GB of storage.	Deploys an appliance with 2258 GB of storage.	Deploys an appliance with 2383 GB of storage.
X-Large	Deploys an appliance with 4279 GB of storage.	Deploys an appliance with 4304 GB of storage.	Deploys an appliance with 4468 GB of storage.	Deploys an appliance with 4518 GB of storage.	Deploys an appliance with 4643 GB of storage.

11. From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**. NFS datastores are thin provisioned by default.

12. Configure the temporary network for communication between the source vCenter Server and the target vCenter Server appliance, and click **Next**.

Option	Action
Choose a network	Select the network to which to connect the new appliance temporarily.

Option	Action
	The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are unsupported and are not displayed in the drop-down menu. Important: If you want to assign a temporary IPv4 address with DHCP allocation, you must select a network that is associated with a port group which accepts MAC address changes.
IP Address family	Select the version for the temporary IP address of the new appliance. Can be either IPv4 or IPv6.
Network type	Select the allocation method for the temporary IP address of the appliance. • Static The wizard prompts you to enter the temporary IP address and network settings. • DHCP A DHCP server is used to allocate the temporary IP address. Select this option only if a DHCP server is available in your environment.

13. On the Ready to complete stage 1 page, review the deployment settings for the target vCenter Server appliance and click **Finish** to start the OVA deployment process.

14. Wait for the OVA deployment process to finish and click **Continue** to proceed with stage 2 of the migration process to transfer the data from the source vCenter Server and start the services of the target appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface of the newly deployed target vCenter Server appliance to transfer the data from the source vCenter Server and set up the services.

The newly deployed target vCenter Server appliance version 8.0 is running on the target server but is not yet configured.

Important: The data from the source vCenter Server is not yet transferred and the services of the target appliance are not started.

Set Up the Target vCenter Server Appliance

When the OVA deployment completes, you are redirected to stage 2 of the migration process to transfer the data from the source vCenter Server and start the services of the newly deployed vCenter Server appliance version 8.0.

Your window of downtime does not begin until you begin to set up the target appliance. You cannot cancel or interrupt the process until it completes with the shut down of the source deployment. Your window of downtime ends when the target appliance starts.

1. Review the introduction to stage 2 of the migration process and click **Next**.
2. On the Select source vCenter Server page, enter the vCenter Single Sign-On administrator password and the root password of the source vCenter Server, enter the password of the user with administrative privileges on the vCenter Server instance, and click **Next**.
3. Optional: Accept the warning message, if any, by clicking **Yes**.
4. If your source Windows machine is connected to an Active Directory domain, enter the credentials for an administrator domain user with permission to add the target machine to the Active Directory domain, and click **Next**.

Note: The installer verifies the entered credentials, but does not check the required privileges to add the target machine to the Active Directory domain. Verify that the user credentials have all the required permissions to add a machine to the Active Directory domain.

5. Review the VMware Customer Experience Improvement Program (CEIP) page and choose if you want to join the program.
For information about the CEIP, see the Configuring Customer Experience Improvement Program section in *vCenter Server and Host Management*.
6. On the ready to complete page, review the migration settings, accept the backup acknowledgment, and click **Finish**.
7. Click **OK** to confirm the shutdown of the source vCenter Server.
8. Wait for the data transfer and setup process to finish and click **OK** to go to the vCenter Server Getting Started page.

The source vCenter Server instance is migrated from Windows to a vCenter Server appliance. The source vCenter Server instance is powered off and the new target appliance starts.

Verify that your migration to an appliance was successful. For verification steps, see [Verify Your vCenter Server Upgrade or Migration Is Successful](#). For post-migration steps, see [After Upgrading or Migrating vCenter Server](#).


GUI Migration of vCenter Server with an External Platform Services Controller to an Appliance

You can use the GUI to migrate vCenter Server with an external Platform Services Controller to a vCenter Server appliance.

Before migrating a vCenter Server deployment with an external Platform Services Controller, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the migration. See Knowledge Base article [KB 68137](#).

If you use Update Manager in the deployment on Windows that you migrate, and Update Manager runs on a separate machine from any other of the components, take an extra step to migrate Update Manager to an appliance.

1. If your deployment on Windows uses an external Update Manager, run Migration Assistant on the Update Manager machine to start the migration of the Update Manager server and database to the Update Manager appliance.
2. Migrate the vCenter Server instance from Windows to a vCenter Server appliance.
3. Verify the newly migrated vCenter Server appliance.
4. Decommission the Platform Services Controller.

Figure 51: vCenter Server 6.7 with External Platform Services Controller Before and After Migration

GUI tasks for migrating an external Platform Services Controller instance from Windows to an appliance.

1. [Download and Mount the vCenter Server Installer ISO file](#) on a network virtual machine or physical server from which you want to perform the migration.
2. [Download and Run VMware Migration Assistant on the Source Windows Machine.](#)
Note: If you are migrating a system that uses an external instance of Update Manager that runs on a separate Windows machine, first run Migration Assistant on the Update Manager machine.
3. Assemble the [Required Information for Migrating vCenter Server from Windows to an Appliance](#) for each Platform Services Controller or vCenter Server instance.
4. [Deploy the OVA File for the Target vCenter Server Appliance](#)
5. [Set Up the Target vCenter Server Appliance](#)

Important: The user name you use to log in to the physical machine from which you want to run the GUI installer, the path to the vCenter Server appliance installer, and values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

For each node to be migrated, the installer:

- Deploys a new target appliance.
- Exports the required files from the source Platform Services Controller or vCenter Server instance.
- Copies the required files to the target appliance for migration.
- Runs the migration process on the target appliance as indicated in the Summary.
- Imports and updates the files and settings of the source Platform Services Controller or vCenter Server instance to the new vCenter Server appliance.

Deploy the OVA File for the Target vCenter Server Appliance

To start the migration process, you use the GUI installer to deploy the OVA file that is included in the installer ISO file as the target vCenter Server appliance.

- See [Prerequisites for Migrating vCenter Server](#)

1. In the vCenter Server installer, navigate to the `vcsa-ui-installer` directory, go to the subdirectory for your operating system, and run the installer executable file.
 - For Windows OS, go to the `win32` subdirectory, and run the `installer.exe` file.
 - For Linux OS, go to the `lin64` subdirectory, and run the `installer` file.
 - For Mac OS, go to the `mac` subdirectory, and run the `Installer.app` file.
2. On the Home page, click **Migrate**.
3. Review the Introduction page to understand the migration process and click **Next**.
4. Read and accept the license agreement, and click **Next**.
5. Connect to the target server to which you want to migrate the source vCenter Server.

Option	Steps
You can connect to an ESXi host on which to deploy the target appliance.	1. Enter the FQDN or IP address of the ESXi host. 2. Enter the HTTPS port of the ESXi host. 3. Enter the user name and password of a user with administrative privileges on the ESXi host, for example, the root user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes.
You can connect to a vCenter Server instance and browse the inventory to select an ESXi host or DRS cluster on which to deploy the target appliance.	1. Enter the FQDN or IP address of the vCenter Server instance. 2. Enter the HTTPS port of the vCenter Server instance. 3. Enter the user name and password of a vCenter Single Sign-On user with administrative privileges on the vCenter Server instance, for example, the <code>administrator@your_domain_name</code> user. 4. Click Next. 5. Accept the certificate warning, if any, by clicking Yes. 6. Select the data center or data center folder that contains the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next. Note: You must select a data center or data center folder that contains at least one ESXi host that is not in lockdown or maintenance mode. 7. Select the ESXi host or DRS cluster on which you want to deploy the new appliance, and click Next.

6. Optional: Review the warning message and try to resolve the warnings, if any, and click **Yes**.
7. On the Set up target appliance VM page, enter a name for the target vCenter Server Appliance, set the password for the root user, and click **Next**.

The password must contain at least eight characters, a number, uppercase and lowercase letters, and a special character, for example, an exclamation mark (!), hash key (#), at sign (@), or brackets (()).

Important: The local operating system password is not migrated to the target appliance.

8. On the Connect to source page, enter the details for the source vCenter Server instance, and click **Next**.
 - a) Enter the IP address or FQDN.
 - b) Enter the user name and password of a user who has administrative privileges on the vCenter Server instance, for example, the administrator@*your_domain_name* user.
 - c) Enter the Migration Assistant Port you received in the Migration Assistant instructions.

9. On the Connect to source instance page, enter the details for the source Windows installation that you want to migrate.

Option	Action
vCenter Server IP Address/FQDN	Enter the IP address or FQDN of the vCenter Server Appliance that you want to upgrade.
vCenter Single Sign-On administrator user name	Enter the vCenter Single Sign-On administrator user name.
vCenter Single Sign-On administrator password	Enter the password of the vCenter Single Sign-On administrator.
vCenter Server HTTPS Port	Optionally, change the default vCenter Server HTTPS port number. The default value is 443.

10. Optional: Accept the warning message, if any, by clicking **Yes**.

11. Select the deployment size for the new vCenter Server Appliance for your vSphere inventory.

Deployment Size Option	Description
Tiny	Deploys an appliance with 2 vCPUs and 14 GB of memory. Suitable for environments with up to 10 hosts or 100 virtual machines
Small	Deploys an appliance with 4 vCPUs and 21 GB of memory. Suitable for environments with up to 100 hosts or 1,000 virtual machines
Medium	Deploys an appliance with 8 vCPUs and 30 GB of memory. Suitable for environments with up to 400 hosts or 4,000 virtual machines
Large	Deploys an appliance with 16 vCPUs and 39 GB of memory. Suitable for environments with up to 1,000 hosts or 10,000 virtual machines
X-Large	Deploys an appliance with 24 vCPUs and 58 GB of memory. Suitable for environments with up to 2,500 hosts or 45,000 virtual machines

Note: At the bottom of the deployment size table, a row shows the size information of the source machine. This size information is reported by the migration assistant and might help understand why you cannot select some deployment sizes.

12. Select the storage size for the new vCenter Server Appliance, and click **Next**.

Storage Size Option	Description for Tiny Deployment Size	Description for Small Deployment Size	Description for Medium Deployment Size	Description for Large Deployment Size	Description for X-Large Deployment Size
Default	Deploys an appliance with 579 GB of storage.	Deploys an appliance with 694 GB of storage.	Deploys an appliance with 908 GB of storage.	Deploys an appliance with 1358 GB of storage.	Deploys an appliance with 2283 GB of storage.
Large	Deploys an appliance with 2019 GB of storage.	Deploys an appliance with 2044 GB of storage.	Deploys an appliance with 2208 GB of storage.	Deploys an appliance with 2258 GB of storage.	Deploys an appliance with 2383 GB of storage.
X-Large	Deploys an appliance with 4279 GB of storage.	Deploys an appliance with 4304 GB of storage.	Deploys an appliance with 4468 GB of storage.	Deploys an appliance with 4518 GB of storage.	Deploys an appliance with 4643 GB of storage.

13. From the list of available datastores, select the location where all the virtual machine configuration files and virtual disks will be stored and, optionally, enable thin provisioning by selecting **Enable Thin Disk Mode**. NFS datastores are thin provisioned by default.

14. Configure the temporary network for communication between the source vCenter Server and the target vCenter Server appliance, and click **Next**.

Option	Action
Choose a network	Select the network to which to connect the new appliance temporarily. The networks displayed in the drop-down menu depend on the network settings of the target server. If you are deploying the appliance directly on an ESXi host, non-ephemeral distributed virtual port groups are unsupported and are not displayed in the drop-down menu. Important: If you want to assign a temporary IPv4 address with DHCP allocation, you must select a network that is associated with a port group which accepts MAC address changes.
IP Address family	Select the version for the temporary IP address of the new appliance. Can be either IPv4 or IPv6.
Network type	Select the allocation method for the temporary IP address of the appliance. • Static The wizard prompts you to enter the temporary IP address and network settings. • DHCP A DHCP server is used to allocate the temporary IP address. Select this option only if a DHCP server is available in your environment.

15. On the Ready to complete stage 1 page, review the deployment settings for the target vCenter Server appliance and click **Finish** to start the OVA deployment process.
16. Wait for the OVA deployment to finish, and click **Continue** to proceed with stage 2 of the deployment process to set up and start the services of the newly deployed appliance.

Note: If you exit the wizard by clicking **Close**, you must log in to the vCenter Server Management Interface to set up and start the services.

The newly deployed target vCenter Server appliance version 8.0 is running on the target server but is not yet configured.

Important: The data from the source vCenter Server is not yet transferred and the services of the target appliance are not started.

Set Up the Target vCenter Server Appliance

When the OVA deployment finishes, you are redirected to stage 2 of the migration process to transfer the data from the source vCenter Server and start the services of the newly deployed target vCenter Server appliance.

The period of time during which your system is unavailable does not begin until you start to set up the target appliance. You cannot cancel or interrupt the process until it finishes with the shutdown of the source deployment. The time your system is unavailable ends when the target appliance starts.

1. Review the introduction to stage 2 of the migration process and click **Next**.
2. On the Select source vCenter Server page, enter the vCenter Single Sign-On administrator password and the root password of the source vCenter Server, enter the password of the user with administrative privileges on the vCenter Server instance, and click **Next**.
3. Optional: Accept the warning message, if any, by clicking **Yes**.
4. If your source Windows machine is connected to an Active Directory domain, enter the credentials for an administrator domain user with permission to add the target machine to the Active Directory domain, and click **Next**.

Note: The installer verifies the entered credentials, but does not check the required privileges to add the target machine to the Active Directory domain. Verify that the user credentials have all the required permissions to add a machine to the Active Directory domain.

5. Specify the replication topology for the vCenter Server. When converging vCenter Server instance with an external Platform Services Controller you must specify the replication topology.

Convergence is the process of converting a vCenter Server instance with an external Platform Services Controller to a vCenter Server instance with those services embedded in the appliance.

- This is the first vCenter Server in the topology that I want to converge.
- This is a subsequent vCenter Server.

If this is a subsequent vCenter Server, provide the IP address of its partner vCenter Server and its HTTPS port.

6. On the Select migration data page, choose the types of data that you want to transfer from the source vCenter Server to the target appliance.

The large amount of data requires more time to be transferred to the new appliance.

7. On the ready to complete page, review the migration settings, accept the backup acknowledgment, and click **Finish**.
8. Click **OK** to confirm the shutdown of the source vCenter Server.
9. Wait for the data transfer and setup process to finish. Click **OK** to go to the vCenter Server Getting Started page.

vCenter Server is migrated from Windows to a newly deployed target vCenter Server appliance. The source vCenter Server is powered off and the target appliance starts.

- [Verify Your vCenter Server Upgrade or Migration Is Successful.](#)
- Upgrade all vCenter Server instances in the vCenter Single Sign-On domain.
- After converging a vCenter Server with an external Platform Services Controller node to a vCenter Server appliance, you must decommission the original external Platform Services Controller. Decommissioning a Platform Services Controller shuts it down and removes it from the single sign-on domain. See [Decommission the Platform Services Controller.](#)
- For post-migration steps, see [After Upgrading or Migrating vCenter Server.](#)
- You can configure high availability for the vCenter Server appliance. For information about providing vCenter Server appliance high availability, see *vSphere Availability*.

CLI Migration of a vCenter Server Installation from Windows to an Appliance

You can use the CLI installer to automatically migrate a vCenter Server or Platform Services Controller from Windows to an appliance.

The installer ISO file contains example templates of JSON files that contain the minimum configuration parameters required for migrating a vCenter Server or Platform Services Controller instance from Windows to an appliance. The example templates are located in the `vcasa-cli-installer/templates/migrate` directory.

Note: Before migrating a vCenter Server deployment with an external Platform Services Controller, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the migration. See Knowledge Base article [KB 68137](#).

CLI tasks for migrating your vCenter Server installation from Windows to an appliance:

1. [Download and Mount the vCenter Server Installer.](#)
2. [Download and Run VMware Migration Assistant on the Source Windows Machine.](#)
3. [Prepare JSON Configuration Files for CLI Migration.](#)
4. [Run the precheck before performing the migration.](#)
5. [Perform a CLI Migration of vCenter Server from Windows to an Appliance.](#)

You can run the CLI installer multiple times with different JSON files to perform multiple CLI migrations, or you can run the CLI migrations concurrently in a batch mode.

Important: The user name that you use to log in to the machine from which you want to run the CLI installer, the path to the vCenter Server appliance installer, the path to your JSON configuration file, and the string values in your JSON configuration file, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

Prepare JSON Configuration Files for CLI Migration

When you use the CLI to migrate a source vCenter Server installation to a target vCenter Server Appliance, you must prepare a JSON template with the configuration values for the new appliance.

Verify your environment meets the requirements for migration. See [System Requirements for Migrating vCenter Server Deployments to vCenter Server Appliance Deployments](#),

Prepare your environment for migration. See [Preparing for Migration](#).

You can migrate vCenter Server or Platform Services Controller instances from Windows to a vCenter Server Appliance by setting values to the configuration parameters in the templates that are available in the installer ISO file. The configuration parameters that are not included in the templates are set to their default values. You can add configuration parameters in the templates to set their values for your migration specification.

The `vcasa-cli-installer/templates/migrate` directory contains example migration templates for CLI migration of vCenter Server 6.7 to an appliance.

For a complete list of the configuration parameters and their descriptions, navigate to the installer subdirectory for your operating system and run the `vcsa-deploy migrate --template-help` command.

Important: The user name that you use to log in to the machine from which you want to run the CLI installer, the path to the vCenter Server Appliance installer, the path to your JSON configuration file, and the string values in your JSON configuration file, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To perform the batch migration, place the JSON templates defining your deployment in a single directory. When invoked, the CLI installer migrates your existing deployment to the new vCenter Server version using the topology defined in the JSON templates.

1. Open the `migrate` subfolder in the `vcsa-cli-installer/templates` directory.
2. Copy the migration template from the `migrate` subfolder to your workspace. For vCenter Server 6.7, use the `migrate/winvc6.7/` folder.
3. Open the template file for your use case in a text editor.
To ensure the correct syntax of your JSON configuration file, use a JSON editor.
4. Fill in values for the required configuration parameters and, optionally, enter additional parameters and their values.

Important: To set a value that contains the backslash (\) or quotation mark (") character, you must precede the character with the backslash (\) character. For example, `"password": "my\"password"` sets the password `my"password`, `"image": "C:\\vmware\\vcsa"` sets the path `C:\vmware\vcsa`.

The boolean values must contain only lowercase characters. Can be either `true` or `false`. For example, `"ssh_enable": false`

5. Save in UTF-8 format and close the file.

Your file is ready to use for migration.

You can create and save as many templates as are needed for your specific environment. When your template is ready, run the pre-check before using it to run the migration. See [Run a Pre-Check Before a CLI Migration to vCenter Server Appliance](#).

JSON Templates for CLI Migration of vCenter Server for Windows

The vCenter Server installer contains JSON templates that are located in the `vcsa-cli-installer/templates` directory. In the `migrate` subfolder, you can find the JSON templates with the minimum configuration parameters for all migration types.

Table 80: Migration JSON Templates Included in the vCenter Server Installer

Location	Template	Description
vcsa-cli-installer\templates\ migrate\vcsa\winvc6.7	embedded_win_vc_to_embedded_v CSA_on_ESXi.json	Contains the minimum configuration parameters required to migrate vCenter Server 6.7 on Windows with an embedded Platform Services Controller to vCenter Server appliance 8.0 on an ESXi host.
	embedded_win_vc_to_embedded_v CSA_on_VC.json	Contains the minimum configuration parameters required to migrate vCenter Server 6.7 on Windows with an embedded Platform Services Controller to vCenter Server Appliance 8.0 on a vCenter Server instance.

Location	Template	Description
	<code>win_vc_to_vCSA_on_ESXi.json</code>	Contains the minimum configuration parameters required to migrate vCenter Server 6.7 on Windows with an external Platform Services Controller to vCenter Server appliance 8.0 on an ESXi host.
	<code>win_vc_to_vCSA_on_VC.json.js on</code>	Contains the minimum configuration parameters required to migrate vCenter Server 6.7 on Windows with an external Platform Services Controller to vCenter Server appliance 8.0 on a vCenter Server instance.

About Migrating a vCenter Server for Windows with an External Platform Services Controller Using the CLI

When migrating a vCenter Server for Windows that uses an external Platform Services Controller, it is converted to a vCenter Server instance with those services embedded in the appliance. For this reason, you must specify replication partner parameters in the JSON migration templates.

When migrating a vCenter Server for Windows that uses an external Platform Services Controller, you must specify replication partner parameters in the JSON migrate template. These parameters indicate if the migration is for one of the following vCenter Server and Platform Services Controller migration scenarios.

- A single vCenter Server and a single Platform Services Controller instance within an SSO domain.
- The first vCenter Server and Platform Services Controller instance within an SSO domain containing more than one vCenter Server.
- Replicating vCenter Server and Platform Services Controller instances within an SSO domain, which you must point to an existing vCenter Server 7.0 instance as the replication partner.

Important: Before migrating a vCenter Server deployment using an external Platform Services Controller, you must first converge the external Platform Services Controller to an embedded Platform Services Controller and then perform the migration. For more information, see [Removal of Platform Services Controller](#).

Deployments using an external Platform Services Controller are often called an $M \times N$ deployment, where M represents the vCenter Server, and N represents the external Platform Services Controller. When representing multiple instances of vCenter Server and Platform Services Controller in a domain, this is expressed as $Mn \times Nn$, where n is the number of instances. For example, $M2 \times N2$ indicates two vCenter Server instances and two Platform Services Controller instances within the same domain. When migrating such a deployment, you first migrate one instance to version 8.0, and then migrate the second instance to 8.0, specifying the first of the two instances as the replicating partner.

The replication parameters in the JSON migration template are located in the `ssso` section of the template, and are `first_instance` and `replication_partner_hostname`.

Parameter Name	Description
<code>first_instance</code>	The <code>first_instance</code> parameter identifies the vCenter Server as either the first instance in your vCenter Server SSO domain, or as a replicating partner. If it is the first vCenter Server instance, set <code>first_instance</code> to <code>true</code>. For additional instances of vCenter Server in the same SSO domain, set <code>first_instance</code> to <code>false</code>.

Parameter Name	Description
<code>replication_partner_hostname *</code>	If you set <code>first_instance</code> to <code>false</code> , you must specify the FQDN or IP address of a vCenter Server that has already been migrated to version 8.0, and whose external PSC has been converged into the vCenter Server appliance.

When you perform the CLI migration of a single vCenter Server for Windows instance that uses an external Platform Services Controller, the process is as follows:

1. Migrate the vCenter Server in your topology using the `win_vc_to_vCSA_on_ESXi.json` template when migrating on ESXi, or the `win_vc_to_vCSA_on_VC.json` template when migrating on a vCenter Server deployment.
2. For a deployment topology consisting of a single vCenter Server and a Platform Services Controller (an M1xN1 deployment), edit the `sso` section of the template such that the `first_instance` parameter is set to `true`, and remove the `replication_partner_hostname` parameter.

```
"sso": {
  "__comments": [ This is the first instance (M1) vCenter Server management node, with "first_instance" set
    to "true" and "replication_partner_hostname" removed.
  ],
  "first_instance": true
}
```

3. Specify the remaining configuration parameters in the JSON template, and migrate vCenter Server to version 8.0 using the CLI. See [Perform a CLI Migration of vCenter Server from Windows to an Appliance](#).
4. After successfully migrating vCenter Server, decommission the Platform Services Controller. See [Decommission the Platform Services Controller](#).

Important: Decommission the Platform Services Controller instances only after migrating and converging all vCenter Server instances in your topology.

For a deployment consisting of two or more vCenter Server and Platform Services Controller instances (an M2xN2 deployment), you first migrate the M1 vCenter Server instance to version 8.0, and then migrate the second, replicating vCenter Server instance.

1. Migrate the vCenter Server in your topology using the `win_vc_to_vCSA_on_ESXi.json` template when migrating on ESXi, or the `win_vc_to_vCSA_on_VC.json` template when migrating on a vCenter Server deployment.
2. For the first vCenter Server instance (M1), edit the `sso` section of the template such that the `first_instance` parameter is set to `true`, and remove the `replication_partner_hostname` parameter.

```
"sso": {
  "__comments": [ This is the first instance (M1) vCenter Server management node, with "first_instance" set
    to "true" and "replication_partner_hostname" removed.
  ],
  "first_instance": true
}
```

3. For the remaining vCenter Server instance (M2), edit the `sso` section of the template such that the `first_instance` parameter is set to `false`, and provide the FQDN or IP address of the vCenter Server instance to which you want to replicate the second node. The replication partner is the first vCenter Server instance (M1) you migrated to version 8.0, which has been converged with its Platform Services Controller instance.

```
"sso": {
  "__comments": [ This is the second instance (M2) vCenter Server management node, with "first_instance" set
    to "false" and "replication_partner_hostname" set to the hostname of the replication partner.],
  "first_instance": false,
  "replication_partner_hostname": "FQDN_or_IP_address"
```

4. Specify the remaining configuration parameters in the JSON templates, and migrate vCenter Server to version 8.0 using the CLI. See [Perform a CLI Migration of vCenter Server from Windows to an Appliance](#).

- After successfully migrating all vCenter Server instances in the topology, decommission the Platform Services Controller instances. See [Decommission the Platform Services Controller](#).

Important: Decommission the Platform Services Controller instances only after migrating and converging all vCenter Server instances in your topology.

Migration Configuration Parameters

When using the CLI installer to migrate your vCenter Server installation to an appliance, you must provide the parameters with values for your migration specification.

The table lists the configuration parameters that you use to provide input data for the source vCenter Server.

Important: The path to the vCenter Server appliance installer, the path to your JSON configuration file, and the string values in your JSON configuration file, including passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are not supported.

To set a value that contains the backslash (\) or quotation mark (") character, you must precede the character with the backslash (\) character. For example, "password": "my\"password" sets the password my"password, "image": "C:\\vmware\\vcsa" sets the path C:\vmware\vcsa.

The boolean values must contain only lowercase characters. Can be either `true` or `false`. For example, "ssh_enable":false.

Sections and Subsections of Configuration Parameters in the JSON Migration Templates

The configuration parameters in the JSON migration templates are organized in sections and subsections.

Table 81: Sections and Subsections of Configuration Parameters in the JSON Migration Templates

Section	Subsection	Description
new_vcsa - describes the target appliance that you want to migrate to	esxi	Use only if you want to deploy the appliance directly on an ESXi host. Contains the configuration parameters that describe the target ESXi host. Note: You must fill in either the <code>esxi</code> or the <code>vc</code> subsection.
	vc	Use only if you want to deploy the appliance on the inventory of a vCenter Server instance. Contains the configuration parameters that describe the target ESXi host or DRS cluster from the vCenter Server inventory. Note: You must fill in either the <code>vc</code> or the <code>esxi</code> subsection.
	appliance	Contains the configuration parameters that describe the appliance.
	os	Contains the configuration parameters that describe the operating system settings for the appliance.
	ovftool_arguments	Optional subsection for adding arbitrary arguments and their values to the OVF Tool command that the installer generates. Important: The vCenter Server appliance installer does not validate the configuration parameters in the <code>ovftool_arguments</code> subsection. If you set arguments that the OVF Tool does not recognize, the deployment might fail.
	temporary_network	Contains the configuration parameters that describe the temporary network for migrating the data from the source to the new target appliance.

Section	Subsection	Description
	user-options	Use only when the source is a vCenter Server instance. Contains the configuration parameters that let you control aspects of the migration process for particular components.
source_vc - describes the source vCenter Server, vCenter Single Sign-On, or Platform Services Controller	vc_win	Contains the configuration parameters that describe the source Windows installation of vCenter Server or Platform Services Controller
	run_migration_assistant	Use only if the source Windows installation is running as a virtual machine and you want to automate the invocation of Migration Assistant. For a source Windows installation running on a physical machine, or if you are running Migration Assistant manually on the source Windows machine, copy and paste the thumbprint value from the Migration Assistant console output on the source machine to the <code>migration_ssl_thumbprint</code> key in the <code>vc_win</code> subsection, and remove the <code>run_migration_assistant</code> section.
ceip - describes joining the VMware Customer Experience Improvement Program (CEIP)	settings	Contains only the <code>ceip_enabled</code> configuration parameter to join or not to join the VMware Customer Experience Improvement Program (CEIP). Required only if you are deploying a vCenter Server appliance. Note: If set to <code>true</code> , you must run the CLI deployment command with the <code>--acknowledge-ceip</code> argument. For information about the CEIP, see the Configuring Customer Experience Improvement Program section in <i>vCenter Server and Host Management</i> .

Configuration Parameters in the `new_vcsa` Section

Table 82: Configuration Parameters in the `new_vcsa` Section, `esxi` Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the target ESXi host on which you want to deploy the appliance.
username	string	A user name with administrative privileges on the target ESXi host, for example, root.
password	string	The password of the user with administrative privileges on the target ESXi host.
deployment_network	string	The name of the network to which to connect the appliance. Note: The network must be accessible from the target ESXi host. Ignored if the target ESXi host has only one network.
datastore	string	The name of the datastore that you want to store all virtual machine configuration files and virtual disks of the appliance. Note: The datastore must be accessible from the ESXi host. The datastore must have enough free space.
port	integer	The port number of the ESXi host. The default port is 443.
ssl_certificate_verification	string	The CLI verifies that a server's security certificate is signed by a Certificate Authority (CA), and establishes a secure connection. If the certificate is self-signed, the CLI stops the upgrade unless you specify one of the following SSL certificate configuration options. Specify the Secure Hash Algorithm 1 (SHA-1) certificate thumbprint. A certificate thumbprint is a hexadecimal string that uniquely identifies a certificate. The thumbprint is calculated from the content of the certificate using a thumbprint algorithm. <code>"thumbprint": "certificate SHA-1 thumbprint"</code> Set <code>verification_mode</code> to <code>NONE</code> . <code>"verification_mode": "NONE"</code>

Name	Type	Description
		If you are connecting to a server with a self-signed certificate, and fail to either specify the SHA-1 certificate thumbprint, or set the verification mode to <code>NONE</code>, the CLI will display the server's self-signed certificate's thumbprint, and prompt you to accept or reject the certificate thumbprint. You can also specify that the CLI ignore the self-signed using the <code>vcasa-deploy upgrade</code> command parameter <code>--no-ssl-certificate-validation</code>. See Syntax of the CLI Migrate Command.

Table 83: Configuration Parameters in the `new_vcasa` Section, `vc` Subsection

Name	Type	Description
hostname	string	The IP address or FQDN of the target vCenter Server instance on which you want to deploy the appliance.
username	string	vCenter Single Sign-On administrator user name on the target vCenter Server instance, for example, <code>administrator@vsphere.local</code> .
password	string	The password of the vCenter Single Sign-On administrator user on the target vCenter Server instance.
deployment_network	string	The name of the network to which to connect the appliance. Note: The network must be accessible from the target ESXi host or DRS cluster on which you want to deploy the appliance. Ignored if the target ESXi host or DRS cluster has only one network.
datacenter	string or array	The vCenter Server data center that contains the target ESXi host or DRS cluster on which you want to deploy the appliance. If the datacenter is located in a folder or a structure of folders, the value must be either a comma-separated list of strings or a comma-separated list as a single string. For example, <pre>["parent_folder", "child_folder", "datacenter_name"]</pre> or <pre>"parent_folder, child_folder, datacenter_name"</pre> If there is no folder path on the data center, then use only the datacenter name. For example, <pre>["datacenter_name"]</pre> or <pre>"datacenter_name"</pre> Note: The value is case-sensitive.
datastore	string	The name of the datastore that you want to store all virtual machine configuration files and virtual disks of the appliance. Note: The datastore must be accessible from the target ESXi host or DRS cluster. The datastore must have at least 25 GB of free space.
port	integer	The port number of the vCenter Server. The default port is 443.
target	string or array	The target cluster, ESXi host, or resource pool on which you want to deploy the new appliance. This is the target you specified with the <code>datacenter</code> parameter. This path must end with a cluster name, ESXi hostname, or resource pool name. Important: You must provide the name that is displayed in the vCenter Server inventory. For example, if the name of the target ESXi host is an IP address in the vCenter Server inventory, you cannot provide an FQDN. Note: All values are case-sensitive.

Name	Type	Description
		If you want the deployed appliance to be listed in a different location within the data center's hierarchy, use the <code>vm_folder</code> parameter described later in this section. If the target cluster, ESXi host, or resource pool is located in a folder or a structure of folders, the value must be a comma-separated list of strings or a comma-separated list as a single string. For example, <pre>["parent_folder", "child_folder", "esxi-host.domain.com"]</pre> or <pre>"parent_folder, child_folder, esxi-host.domain.com"</pre> If the target ESXi host is part of a cluster, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example, <pre>["cluster_name", "esxi-host.domain.com"]</pre> or <pre>"cluster_name, esxi-host.domain.com"</pre> If you are deploying to a resource pool, include the label <code>Resources</code> before the resource pool name. For example: <pre>["cluster_name", "Resources", "resource_pool_name"]</pre> Note: Pre-checks verify only the memory of the resource pool.
<code>vm_folder</code>	string	Optional. The name of the virtual machine (VM) folder to which to add the appliance.

Table 84: Configuration Parameters in the `new_vcsa` Section, `appliance` Subsection

Name	Type	Description
<code>thin_disk_mode</code>	Boolean	Set to <code>true</code> to deploy the appliance with thin virtual disks.
<code>deployment_option</code>	string	The size of the appliance. - Set to <code>tiny</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the default storage size. Deploys an appliance with 2 CPUs, 10 GB of memory, and 300 GB of storage. - Set to <code>tiny-lstorage</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the large storage size. Deploys an appliance with 2 CPUs, 10 GB of memory, and 825 GB of storage. - Set to <code>tiny-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 10 hosts and 100 virtual machines with the x-large storage size. Deploys an appliance with 2 CPUs, 10 GB of memory, and 1700 GB of storage. - Set to <code>small</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the default storage size. Deploys an appliance with 4 CPUs, 16 GB of memory, and 340 GB of storage. - Set to <code>small-lstorage</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the large storage size. Deploys an appliance with 4 CPUs, 16 GB of memory, and 870 GB of storage. - Set to <code>small-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 100 hosts and 1,000 virtual machines with the x-large storage size. Deploys an appliance with 4 CPUs, 16 GB of memory, and 1750 GB of storage. - Set to <code>medium</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the default storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 525 GB of storage. - Set to <code>medium-lstorage</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the large storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 1025 GB of storage.

Name	Type	Description
		- Set to <code>medium-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 400 hosts and 4,000 virtual machines with the x-large storage size. Deploys an appliance with 8 CPUs, 24 GB of memory, and 1905 GB of storage. - Set to <code>large</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the default storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 740 GB of storage. - Set to <code>large-lstorage</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the large storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 1090 GB of storage. - Set to <code>large-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 1,000 hosts and 10,000 virtual machines with the x-large storage size. Deploys an appliance with 16 CPUs, 32 GB of memory, and 1970 GB of storage. - Set to <code>xlarge</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the default storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 1180 GB of storage. - Set to <code>xlarge-lstorage</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the large storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 1230 GB of storage. - Set to <code>xlarge-xlstorage</code> if you want to deploy a vCenter Server appliance for up to 2,000 hosts and 35,000 virtual machines with the x-large storage size. Deploys an appliance with 24 CPUs, 48 GB of memory, and 2110 GB of storage.
<code>image</code>	string	Optional. A local file path or URL to the vCenter Server appliance installation package. By default the installer uses the installation package that is included in the ISO file, in the <code>vcsa</code> folder.
<code>name</code>	string	The VM name for the appliance. Must contain only ASCII characters except a percent sign (%), backslash (\), or forward slash (/) and must be no more than 80 characters in length.
<code>ovftool_path</code>	string	Optional. A local file path to the OVF Tool executable file. By default the installer uses the OVF Tool instance that is included in the ISO file, in the <code>vcsa/ovftool</code> folder.

Table 85: Configuration Parameters in the `new_vcsa` Section, `os` Subsection

Name	Type	Description
<code>password</code>	string	The password for the root user of the appliance operating system. The password must contain between 8 and 20 characters, at least one uppercase letter, at least one lowercase letter, at least one number, and at least one special character, for example, a dollar sign (\$), hash key (#), at sign (@), period (.), or exclamation mark (!). All characters must be lower ASCII characters without spaces.
<code>ssh_enable</code>	Boolean	Set to <code>true</code> to enable the SSH administrator login to the appliance.

Table 86: Configuration Parameters in the `new_vcsa` Section, `temporary_network` Subsection

Name	Type	Description
<code>ip_family</code>	string	IP version for the network of the appliance. Set to <code>ipv4</code> or <code>ipv6</code> .
<code>mode</code>	string	IP assignment for the network of the appliance. Set to <code>static</code> or <code>dhcp</code> .

Name	Type	Description
ip	string	IP address for the appliance. Required only if you use a static assignment, that is, if you set the <code>mode</code> parameter to <code>static</code> . You must set an IPv4 or IPv6 address that corresponds to the network IP version, that is, to the value of the <code>ip_family</code> parameter. An IPv4 address must comply with the RFC 790 guidelines. An IPv6 address must comply with the RFC 2373 guidelines.
dns_servers	string or array	IP addresses of one or more DNS servers. To set more than one DNS server, use a comma-separated list of strings or a comma-separated list as a single string to provide the path. For example, [<code>"x.y.z.a"</code> , <code>"x.y.z.b"</code>] or <code>"x.y.z.a, x.y.z.b"</code> Required only if you use a static assignment, that is, if you set the <code>mode</code> parameter to <code>static</code> .
prefix	string	Network prefix length. Required only if you use assignment, that is, if you set the <code>mode</code> parameter to <code>static</code> . For IPv4 version, the value must be between 0 and 32. For IPv6 version, the value must be between 0 and 128.
gateway	string	IP address of the default gateway. For IPv6 version, the value can be <code>default</code> .

Table 87: Configuration Parameters in the `new_vcsa` Section, `user_options` Subsection

Name	Type	Description
vcdb_migrateSet	string	Select the types of data to migrate from the old appliance to the new appliance. Data is copied from the source vCenter Server to the target server. The original source of the data remains unchanged. - Set to <code>core</code> if you want to transfer only the configuration data. This provides the fastest data migration, keeping system downtime to a minimum. - Set to <code>core_events_tasks</code> if you want to transfer the configuration and historical data (events and tasks) immediately. vCenter Server will not start until all data has been migrated from the source vCenter Server for Windows. - Set to <code>all</code> if you want to transfer the configuration, historical, and performance metrics data immediately. vCenter Server will not start until all data has been migrated from the source vCenter Server for Windows. This option transfers the largest amount of data, and requires more downtime than other data migration options. - Set to <code>transfer_events_tasks_after_upgrade</code> if you want to transfer the historical data (events and tasks) in the background after the upgrade finishes. During this time, vCenter Server performance might not be optimal. - Set to <code>transfer_stats_events_tasks_after_upgrade</code> if you want to transfer historical data and performance metrics data in the background after the upgrade finishes. During this time, vCenter Server performance might not be optimal. Note: To minimize the migration time and the amount of storage required for the new vCenter Server appliance, use the <code>core</code> value. For more information on the types of data you can transfer from your existing vCenter Server to the new, upgraded vCenter Server, see Transferring Data from an Existing vCenter Server Appliance.

Table 88: Configuration Parameters in the `new_vcsa` Section, `sso` Subsection

Name	Type	Description
<code>first_instance</code>	string	Set <code>first_instance</code> to <code>true</code> if this is the first converge migration operation in your vCenter Server SSO domain, or if you do not want to set the replication partner for this node. If you want to set the replication partner for this node, set <code>first_instance</code> to <code>false</code> and provide the value of the <code>replication_partner_hostname</code> as the FQDN of a previously converged node which is currently in an embedded vCenter Server topology. For more information on specifying the <code>first_instance</code> and <code>replication_partner_hostname</code> parameters, see About Migrating a vCenter Server for Windows with an External Platform Services Controller Using the CLI .
<code>replication_partner_hostname</code>	string	Host name of the replication partner. Remove if it is the <code>first_instance</code>

Requirements for the Automatic Invocation of Migration Assistant

You use the `run_migration_assistant` subsection to automate the invocation of Migration Assistant. Automatic invocation works only if the source Windows installation is running as a virtual machine.

The user account that you specify in the `os_username` or `vum_os_username` parameters needs a privilege elevation to Administrator. For example:

- The built-in Windows Administrator account.
- A user account with a user name other than Administrator that is a member of the local Windows Administrators group.
- The Domain Administrator account with the user name Administrator that is a member of the local Windows Administrators group.
- The user name ID must be in the format `your_domain_name\user_ID` or `user_ID@your_domain_name`.

Restriction: Automatic invocation of Migration Assistant does not work with a Windows account that requires a privilege elevation to Administrator. Instead, run Migration Assistant manually on the source Windows machine. Copy and paste the thumbprint value from the Migration Assistant console output on the source machine to the `migration_ssl_thumbprint` key in the `vc_win` subsection, and remove the `run_migration_assistant` section.

Configuration Parameters in the `source_vc` Section**Table 89: Configuration Parameters in the `source_vc` Section, `vc_win` Subsection**

Name	Type	Description
<code>hostname</code>	string	The host name or IP address of the source Windows installation of vCenter Server or Platform Services Controller that you want to migrate.
<code>username</code>	string	A vCenter Single Sign-On user name with administrative privileges for the vCenter Server, vCenter Single Sign-On, or Platform Services Controller instance that you want to migrate.
<code>password</code>	string	The password of the vCenter Server or Platform Services Controller instance that you want to migrate.
<code>migration_port</code>	string	Migration Assistant port number shown in the Migration Assistant console. The default port is 9123.
<code>active_directory_domain</code>	string	The name of the Active Directory domain to which the source vCenter Server instance is joined.

Name	Type	Description
active_directory_username	string	Administrator user name of the Active Directory domain to which the source vCenter Server instance is joined.
active_directory_password	string	Administrator password of the Active Directory domain to which the source vCenter Server instance is joined. Note: The installer verifies the entered credentials, but does not check the required privileges to add the target machine to the Active Directory domain. Verify that the user credentials have all the required permissions to add a machine to the Active Directory domain.
migration_ssl_thumbprint	string	The SSL thumbprint of Migration Assistant.

Table 90: Configuration Parameters in the `source_vc` Section, `run_migration_assistant` Subsection

Name	Type	Description
esxi_hostname	string	FQDN or IP address of ESXi on which the source vCenter Server, vCenter Single Sign-On, or Platform Services Controller instance resides.
esxi_username	string	User name of a user with administrative privileges on the ESXi host.
esxi_password	string	The password of the ESXi host user. If omitted, you are prompted to enter the password at the command console during the template verification.
esxi_port	string	The port number of the ESXi host. The default port is 443.
os_username	string	Administrator user name for the source Windows machine.
os_password	string	Administrator user password for the source Windows machine. If omitted, you are prompted to enter it at the command console during the template verification.
migration_ip	string	The IP address of the network adapter to be migrated.
migration_port	string	Migration Assistant port number shown in the Migration Assistant console. The default port is 9123.
export_dir	string	Directory to export source configuration and data.
sa_password	string	The IP address of the network vCenter Server service account user password. This option is only required if the vCenter Server service is running under a non-LocalSystem account. If omitted, you are prompted to enter it at the command console during the template verification.

Table 91: Configuration Parameters in the `source_vum` Section, `run_migration_assistant` Subsection

Name	Type	Description
esxi_hostname	string	FQDN or IP address of ESXi on which the source vCenter Server, vCenter Single Sign-On, or Platform Services Controller instance resides.
esxi_username	string	User name of a user with administrative privileges on the ESXi host.
esxi_password	string	The password of the ESXi host user. If omitted, you are prompted to enter the password at the command console during the template verification.
esxi_port	string	The port number of the ESXi host. The default port is 443.
vum_hostname	string	FQDN or IP address of ESXi on which the source Update Manager instance resides.
vum_os_username	string	Administrator user name for the source Windows machine.

Name	Type	Description
vum_os_password	string	Administrator user password for the source Update Manager Windows machine. If omitted, you are prompted to enter it at the command console during the template verification.
migration_port	string	Migration Assistant port number shown in the Migration Assistant console. The default port is 9123.
export_dir	string	Directory to export source configuration and data.

Configuration Parameters in the `ceip` Section

Table 92: Configuration Parameters in the `ceip` Section, `settings` Subsection

Name	Type	Description
ceip_enabled	Boolean	Set to <code>true</code> to join the CEIP for this appliance.

Run a Pre-Check Before a CLI Migration to vCenter Server Appliance

You can run a pre-check to verify that the migration requirements are met and resolve any problems before migration of your vCenter Server deployment.

[Prepare JSON Configuration Files for CLI Migration](#) using the example templates and [configuration parameters](#).

Before migrating your vCenter Server deployment to an appliance, you can run a pre-check to find out the disk space requirement, the estimated upgrade time, and the extensions registered with the vCenter Server appliance. Running pre-upgrade is an optional but highly recommended step when planning your upgrade.

1. Verify your template without deploying the appliance by entering the command: `vcsa-deploy migrate --verify-template-only path_to_json_file`.
2. Run CLI Migrate using the `--precheck-only` option.
You can adjust your migration plans using the disk space requirements and estimated migration time. If you receive an error, you can troubleshoot and resolve the problem before the actual migration.
3. After resolving any errors, run the CLI Migrate command using the `--verify-template-only` option again until all errors are resolved.

You are now prepared for an error-free CLI migration process.

[Perform a CLI Migration of vCenter Server from Windows to an Appliance.](#)

Perform a CLI Migration of vCenter Server from Windows to an Appliance

You can migrate vCenter Server to an appliance from a machine that is in your vSphere network.

- See [Prerequisites for Migrating vCenter Server](#)
- Create a snapshot of the deployment that you want to migrate as a precaution in case of failure during the migration process.
- Download the installer ISO file from the VMware website to a machine that is in your vSphere network. The installer ISO filename is `VMware-VCSA-all-8.0.0-yyyyyy.iso`, where `yyyyyy` is the build number. See [Download and Mount the vCenter Server Installer](#).
- [Prepare JSON Configuration Files for CLI Migration](#).
- [Run a Pre-Check Before a CLI Migration to vCenter Server Appliance](#) to identify problems and refine your migration plan.

- Review the optional arguments for running the migration. See [Syntax of the CLI Migrate Command](#).
1. Navigate to the software CLI installer directory for your operating system.
 - If you are deploying the appliance from a machine with Windows OS, navigate to the `vcsa-cli-installer\win32` directory.
 - If you are deploying the appliance from a machine with Linux OS, navigate to the `vcsa-cli-installer/lin64` directory.
 - If you are deploying the appliance from a machine with Mac OS, navigate to the `vcsa-cli-installer/mac` directory.
 2. Select the CLI installer: `vcsa-deploy.exe`.
 3. Run the migration command.

```
vcsa-deploy migrate --accept-eula optional_arguments path_to_the_json_file
```

The `optional_arguments` variable is a space-separated list of optional arguments to set additional configurations.

For example, you can set the location of the log and other output files that the installer generates.

```
vcsa-deploy migrate --accept-eula --log-dir=path_to_the_location path_to_the_json_file
```

The migration template is deployed. You can [Verify Your vCenter Server Upgrade or Migration Is Successful](#).

Syntax of the CLI Migrate Command

You can use one or more command arguments to set the execution parameters of the migrate command.

You can add a space-separated list of arguments to the CLI upgrade command.

```
vcsa-deploy migrate list_of_arguments path_to_the_json_file
```

The required `template` argument provides the path of a JSON file that describes the vCenter Server appliance deployment procedure. You can place multiple JSON files in a directory, and the CLI will migrate all of the deployments in batch mode. For more information on performing concurrent migrations, see [Prepare JSON Configuration Files for CLI Migration](#).

Important: The string values, including the passwords, must contain only ASCII characters. Extended ASCII and non-ASCII characters are unsupported.

To set a value that contains the backslash (\) or quotation mark (") character, you must precede the character with the backslash (\) character. For example, `"password": "my\"password"` sets the password `my"password`, `"image": "C:\\vmware\\vcsa"` sets the path `C:\vmware\vcsa`.

The boolean values must contain only lowercase characters. Can be either `true` or `false`. For example, `"ssh.enable": false`.

Optional Argument	Description
<code>--accept-eula</code>	Accepts the end-user license agreement. Required for executing the deployment command.
<code>-h, --help</code>	Displays the help message for the command.
<code>--template-help</code>	Displays the help message for the configuration parameters in the JSON deployment file. You can use <code>vcsa-deploy [subcommand] --help</code> for a list of subcommand-specific arguments.
<code>-v, --verbose</code>	Adds debug information to the console output.
<code>-t, --terse</code>	Hides the console output. Displays only warning and error messages.

Optional Argument	Description
<code>--log-dir, LOG_DIR</code>	Specifies the location of the log and other output files that the installer generates.
<code>--skip-ovftool-verification</code>	Performs basic verification of the configuration parameters and deploys the vCenter Server appliance, but does not validate the OVF Tool parameters in the <code>ovftool_arguments</code> subsection of the JSON template. If you set arguments that the OVF Tool does not recognize, the deployment might fail.
<code>--no-ssl-certificate-verification</code>	Prohibits SSL verification of ESXi connections. The CLI verifies that a server's security certificate is signed by a Certificate Authority (CA), and establishes a secure connection. If the certificate is self-signed, the CLI will stop the upgrade unless you specify that the CLI ignore the self-signed certificate using the <code>--no-ssl-certificate-validation</code> command parameter. If you are connecting to a server with a self-signed certificate, and fail to specify that the CLI accept it, the CLI displays the server's self-signed certificate thumbprint, and prompts you to accept or reject the it. You can also specify that the CLI ignore self-signed certificates using the <code>ssl_certificate_verification</code> configuration parameter in the JSON template. See Upgrade Configuration Parameters for CLI Upgrade of vCenter Server Appliance. Important: Avoid using this option as it may cause problems during or after upgrade due to an un-validated identity of the target host.
<code>--operation-id</code>	Lets you provide an identifier to track the concurrent installation, migration, or upgrade of multiple vCenter Server instances. If do not provide an operation ID, the CLI generates a universally unique identifier (UUID) which you can use to identify the different instances of vCenter Server and their installation or upgrade status.
<code>--verify-template-only</code>	Performs basic template verification without installing Upgrade Runner, running prechecks, and upgrading or migrating the vCenter Server appliance.
<code>--precheck-only</code>	Installs Migration Assistant on the source vCenter Server virtual machine, and runs a complete set of prechecks without performing the migration.
<code>--acknowledge-ceip</code>	Confirms acknowledgement of your VMware Customer Experience Improvement Program (CEIP) participation. This argument is required if <code>ceip.enabled</code> is set to true in the upgrade template.

Exit Code	Description
0	Command ran successfully
1	Runtime error
2	Validation error
3	Template error

After Upgrading or Migrating vCenter Server

After you upgrade to vCenter Server, consider the post-upgrade options and requirements.

- Complete any component reconfigurations that might be required for changes during upgrade.
- Verify that you understand the authentication process and identify your identity sources.
- If you migrated vCenter Server on Windows to a target vCenter Server appliance and you use any local OS user names to log in to the vCenter Single Sign-On, you must recreate them and reassign permissions.
- If you performed an upgrade, upgrade any additional modules that are linked to this instance of vCenter Server, such as Update Manager. If you performed a migration from vCenter Server for Windows to a vCenter Server appliance, the Update Manager module is migrated as well into vSphere Lifecycle Manager.
- Optionally, upgrade or migrate the ESXi hosts in the vCenter Server inventory to the same version as the vCenter Server instance.
- If you use Update Manager in your vCenter Server deployment, and Update Manager and vCenter Server were running on separate machines prior the migration, consider to shut down or delete the Update Manager host machine after the migration is complete. Before disposing of the Update Manager host machine, take into account the following:
 - You might need the host machine for rolling back purposes of your upgraded or migrated environment.
 - You might have other software that runs on that machine.
- If you are using smart card authentication, ensure to keep the smart card port open in the client environment. By default, the smart card port is open in vCenter Server. For details about the smart card port, see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com>
- If you plan to install Windows 11 as a guest operating system on a virtual machine, you must configure a key provider. Installing Windows 11 requires a Trusted Platform Module (TPM) 2.0. When installing Windows 11 as a guest operating system on a virtual machine, instead of using a physical TPM, you can use a virtual Trusted Platform Module (vTPM). A vTPM is a software-based representation of a physical TPM 2.0 chip. A vTPM depends on virtual machine encryption to secure vital TPM data, and so requires that you configure a key provider. For information about the key providers supported by vSphere, see the *Virtual Machine Encryption* chapter in the *vSphere Security* documentation. The easiest way is to configure a VMware vSphere® Native Key Provider™. vSphere Native Key Provider is included in all vSphere editions and does not require an external key server. For information about configuring a vSphere Native Key Provider, see *Configuring and Managing vSphere Native Key Provider* chapter in the *vSphere Security* documentation. As with all security solutions, consider the system design, implementation considerations, and tradeoffs of using vSphere Native Key Provider.
- If you had modified the DRS cluster automation level before upgrade, you can either continue to use the modified settings or revert the automation level to fully automated.
- If you upgrade vCenter Server to a version that starts with the same number, you can keep the same license. If you upgrade vCenter Server to a major version that starts with a different number, the evaluation period restarts and you must assign a new license. For more information about licensing, see vSphere License Management in *vCenter Server and Host Management*.

Verify Your vCenter Server Upgrade or Migration Is Successful

You can verify the success of your vCenter Server upgrade or migration.

You must be logged into the upgraded or migrated vCenter Server instance. If you created a reference of required information based on a CLI template, you can use it to validate the upgrade or migration success.

1. Verify that the IP address is correct.
2. Verify that the Active Directory registration has not changed.
3. Verify the Network registration is correct.
4. Verify the Domain is correct.
5. Verify the certificates are valid.
6. Verify the inventory data is correctly migrated.
 - a) Review the events history.
 - b) Review the performance charts.
 - c) Review the users, permissions, and roles.

If the postupgrade or postmigration configuration conforms to your required information or CLI template reference and expectations, the vCenter Server upgrade or migration is complete.

You can troubleshoot unexpected behavior by reviewing logs. You can also perform a rollback to the source configuration. See [Roll Back a vCenter Server Upgrade or vCenter Server on Windows Migration](#)

How to Log in to Your vCenter Server by Using the vSphere Client

Learn how to use the vSphere Client to log in to vCenter Server and manage your vSphere inventory.

The vSphere Client is automatically installed as part of the vCenter Server appliance deployment. This way, the vSphere Client always points to the same vCenter Single Sign-On instance.

1. Open a Web browser and enter the URL for your vCenter Server instance:

`https://vcenter_server_ip_address_or_fqdn`

As an alternative, you can open a Web browser and enter the URL for the vSphere Client:

`https://vcenter_server_ip_address_or_fqdn/ui.`

2. If a warning message about a potential security risk appears, select to continue to the website.

Browser	Action
Microsoft Edge	1. Click Details. 2. Under the additional message that appears, click Go on to the webpage.
Mozilla Firefox	1. Click Advanced. 2. Under the additional message that appears, click Accept the Risk and Continue.
Google Chrome	1. Click Advanced. 2. Under the additional message that appears, click Proceed to vcenter_server_ip_address_or_fqdn.

3. On the vSphere Welcome page, select **Launch vSphere Client (HTML5)**.
4. If the warning message about a potential security risk appears again, repeat Step 2.
5. Enter the credentials of a user who has permissions on vCenter Server and click **Login**.

The vSphere Client connects to all the vCenter Server systems on which the specified user has permissions, and you can view and manage the vSphere inventory.

Decommission the Platform Services Controller

After you upgrade to vCenter Server appliance, can decommission external Platform Services Controller instances.

Make sure no vCenter Server instances are pointing to the Platform Services Controller before decommissioning.

Reconfigure any products deployed into the environment that use the Platform Services Controller to use the newly deployed vCenter Server appliance.

Make a backup of the Platform Services Controller before decommissioning to ensure no loss of data.

After converging an external Platform Services Controller node to a vCenter Server appliance, decommission the original external Platform Services Controller. Decommissioning a Platform Services Controller shuts it down and removes it from the single sign-on domain.

Decommission an External Platform Services Controller Appliance

Decommission and remove an external Platform Services Controller appliance from the vSphere inventory.

Make sure no vCenter Server instances are pointing to the Platform Services Controller before decommissioning.

Reconfigure any products deployed into the environment that use the Platform Services Controller to use the newly deployed vCenter Server appliance.

Make a backup of the Platform Services Controller before decommissioning to ensure no loss of data.

After converging an external Platform Services Controller node to a vCenter Server appliance, decommission the original external Platform Services Controller. Decommissioning a Platform Services Controller shuts it down and removes it from the single sign-on domain.

1. Stop the Platform Services Controller that you no longer need.
2. Log in as root to the appliance shell of one of the Platform Services Controller appliances within the domain.
3. Enable access to the Bash shell using the `shell.set --enabled true` command.
4. Start the Bash shell using the `shell` command, and log in.
5. Unregister the stopped Platform Services Controller appliance using the `cmsso-util unregister` command.

Where, *Platform_Services_Controller_System* is the FQDN or IP address of the

Platform Services Controller that you want to decommission. Run this command only on one of the Platform Services Controller replication partners, as the synchronization removes the entries from all other Platform Services Controller replication partners. Enter the `vCenter_Single_Sign_On_password` in quotes.

Note: The `cmsso-util unregister` command restarts the services on the Platform Services Controller on which you are running the command. It can take several minutes for the command to finish.

```
cmsso-util unregister --node-pnid Platform_Services_Controller_System --username administrator@your_domain_name --passwd 'vCenter_Single_Sign_On_password'
```

6. Remove the Platform Services Controller appliance that you no longer need from the vSphere inventory.

You can delete the virtual machine for the decommissioned Platform Services Controller.

Ensure that any external solutions or products are registered with the new vCenter Server appliance.

Decommission an External Platform Server Controller for Windows

Decommission and remove an external Platform Services Controller for Windows from the vSphere inventory.

Make sure no vCenter Server instances are pointing to the Platform Services Controller before decommissioning.

Reconfigure any products deployed into the environment that use the Platform Services Controller to use the newly deployed vCenter Server appliance.

Make a backup of the Platform Services Controller before decommissioning to ensure no loss of data.

After converging an external Platform Services Controller node to a vCenter Server appliance, decommission the original external Platform Services Controller. Decommissioning a Platform Services Controller shuts it down and removes it from the single sign-on domain.

1. On the Windows computer running Platform Services Controller, click **Start > Run**, type `cmd`, and click **OK**.
The Command Prompt opens.
2. Navigate to the directory `C:\Program Files\VMware\vCenter Server\bin\`
3. Unregister the stopped Platform Services Controller appliance using the `cmsso-util unregister` command.

Where, *Platform_Services_Controller_System* is the FQDN or IP address of the Platform Services Controller that you want to decommission. Run this command only on one of the Platform Services Controller replication partners, as the synchronization removes the entries from all other Platform Services Controller replication partners. Enter the *vCenter_Single_Sign_On_password* in quotes.

Note: The `cmsso-util unregister` command restarts the services on the Platform Services Controller on which you are running the command. It can take several minutes for the command to finish.

```
cmsso-util unregister --node-pnid Platform_Services_Controller_System --username administrator@your_domain_name --passwd 'vCenter_Single_Sign_On_password'
```

4. Remove the Platform Services Controller appliance that you no longer need from the vSphere inventory.

You can delete the virtual machine for the decommissioned Platform Services Controller.

Ensure that any external solutions or products are registered with the new vCenter Server appliance.

Identity Sources for vCenter Server with vCenter Single Sign-On

You can use identity sources to attach one or more domains to vCenter Single Sign-On. A domain is a repository for users and groups that the vCenter Single Sign-On server can use for user authentication.

Note: At any time, only one default domain exists. If a user from a non-default domain logs in, that user must add the domain name to authenticate successfully. The domain name is in the form:

`DOMAIN\user`

The following identity sources are available.

- Active Directory over LDAP. vCenter Single Sign-On supports multiple Active Directory over LDAP identity sources.
- OpenLDAP versions 2.4 and later. vCenter Single Sign-On supports multiple OpenLDAP identity sources.

Note: A future update to Microsoft Windows will change the default behavior of Active Directory to require strong authentication and encryption. This change will impact how vCenter Server authenticates to Active Directory. If you use Active Directory as your identity source for vCenter Server, you must plan to enable LDAPS. For more information about this Microsoft security update, see <https://portal.msrc.microsoft.com/en-US/security-guidance/advisory/ADV190023> and <https://blogs.vmware.com/vsphere/2020/01/microsoft-ldap-vsphere-channel-binding-signing-adv190023.html>.

Reregister Plug-In Solution in vCenter Server After Upgrade or Migration

You reregister a previously registered plug-in solution and any third-party client plug-in packages with vCenter Server after the SSL certificate has been refreshed following an upgrade or migration.

Consult the vendor documentation for any solution-based vCenter Server extensions and client plug-ins for instructions to re-register after a vCenter Server upgrade or migration.

If the procedure provided by your plug-in solution vendor fails to reregister the plug-in, you can use the following procedure to remove the plug-in registration, and then register it again with vCenter Server. For information on registering plug-ins, see the *vCenter Server and Host Management* documentation. For information on removing or disabling unwanted plug-ins from vCenter Server, see Knowledge Base article [KB 102536](#).

1. In a Web browser, navigate to the Managed Object Browser of your vCenter Server.
`https://vcenter_server_ip_address_or_fqdn/mob/?moid=ExtensionManager`
2. Log in with your vCenter Server credentials.
3. On the `ManagedObjectReference:ExtensionManager` page, under Methods, click **UnregisterExtension**.
4. On the `void UnregisterExtension` page, in the text box inside the Value column, enter the value for the `key` property of the `Extension` data object of your vSphere Client extension.
5. To unregister the extension click **Invoke Method**.

Go to the solution registration page and register the plug-in.

Verify that your extension is registered successfully with vCenter Server by using one of the following approaches.

- In the vSphere Client, go to Administration and under Solutions, select **Client Plug-Ins** and click **Check for New Plug-Ins**.
- Log out and log in again to the vSphere Client. The vSphere Client checks for new plug-ins for each new user session

Roll Back a vCenter Server Upgrade or vCenter Server on Windows Migration

You can roll back a vCenter Server upgrade or migration by reverting to the source appliance or vCenter Server on Windows.

You must have access to the source vCenter Server appliance or vCenter Server on Windows.

The rollback procedure applies in the following upgrade and migration contexts:

- vCenter Server appliance with an embedded Platform Services Controller .
- vCenter Server appliance with an external Platform Services Controller .

To revert a failed migration of vCenter Server, see Knowledge Base article [KB 2146453](#).

Monitor and Manage Historical Data Migration

You can monitor and manage the background migration of historical data using the vCenter Server Management Interface.

- Verify that vCenter Server is successfully deployed and running.
- You must have selected to import historical data from the external database in use by the previous version of vCenter Server, to the embedded PostgreSQL database in vCenter Server 8.0. See [Transferring Data from an Existing vCenter Server Appliance](#).

Using the vCenter Server Appliance Management Interface, you can perform the following data management tasks:

- Monitor the progress of the data migration.

- Pause the data migration.
 - Cancel the data migration.
1. In a Web browser, go to the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
 2. Log in as root.
The default root password is the password you set when deploying vCenter Server .
 3. A status bar at the top of the vCenter Server Management Interface displays the percentage of data that has been copied from the source vCenter Server to the PostgreSQL database embedded in the target vCenter Server appliance.
 4. To pause or cancel the data migration, click **Manage**.

Option	Description
Pause	During the data migration, vCenter Server performance might not be optimal. Pausing the migration lets vCenter Server run with better performance until such time that you can resume the process when it does not interfere with the needs of your business.
Resume	You can resume the data migration at a time that does not interfere with the business needs of your vCenter Server deployment.
Cancel	If you cancel the data import, the historical data will fail to import to the embedded PostgreSQL database. You cannot recover the data once you cancel the operation. If you cancel the import process, and want to import the historical data later, you must restart the upgrade or migration process from Stage 1 of the GUI installer.

When the data has been migrated, a success message displays in the vCenter Server Management Interface status bar.

Managing Software Updates, Upgrades, and Product Compatibility

You can view available vCenter Server updates and upgrades and produce interoperability reports about VMware products associated with vCenter Server. The reports help you plan vCenter Server updates and upgrades in your environment.

You can perform an interoperability check of VMware products within your environment against both source and target vCenter Server versions. You can also generate a pre-update report to ensure your environment meets the minimum software and hardware requirements for a successful upgrade of vCenter Server. You can plan vCenter Server updates and upgrades based on the information in the reports you generate.

Interoperability report for the current vCenter Server

Lists VMware products in your environment and their compatibility with the current, or source, version of vCenter Server.

Interoperability reports for target versions of vCenter Server

Lists the VMware products in your environment and their compatibility against the target version of vCenter Server to which you want to upgrade.

Pre-update reports

With the pre-update report, you can ensure that your system meets the minimum software and hardware requirements for a successful upgrade of vCenter Server. The report provides information about problems that might prevent the completion of a software upgrade, and actions you can take to remedy those problems.

Monitor Interoperability for the Current vCenter Server Version

You can view an interoperability report that lists products in your environment associated with the current vCenter Server version and their compatibility.

To use Update Planner, you must join the VMware Customer Experience Improvement Program (CEIP). For information about the CEIP, see the "Configuring Customer Experience Improvement Program" section in *vCenter Server and Host Management*.

The product interoperability report lists available products in your environment and their compatibility with the selected vCenter Server version.

1. In the vSphere Client navigate to a vCenter Server for which you want to view an interoperability report.
 2. On the **Monitor** tab, click **Interoperability**.
The Product Interoperability report lists all available products in your environment associated with the selected vCenter Server.
 3. (Optional) If there are VMware products in your environment that go undetected, you can manually add them to the list of products to check and regenerate the interoperability report.
 - a) Click **Add Product** in the Product Interoperability pane.
 - b) Select the VMware product and version you want to include in the interoperability report.
Repeat this step for each VMware product you want to add to the report.
 - c) To add the selected products to the product inventory click **Done**.
 - d) Regenerate the report and verify that the products you added are included in the report.
 4. (Optional) Click **Export** in the Product Interoperability pane to export and save a copy of the report as a comma-separated values (CSV) file, which allows data to be saved in a tabular format.
- Using the information in the report, you can apply patches and updates to your vCenter Server environment. See [Patching and Updating vCenter Server 8.0 Deployments](#).
 - You can generate a pre-upgrade report that lists actions that you must address to ensure a successful upgrade of vCenter Server in your environment. See [Generate an Interoperability Report for Target vCenter Server Versions](#).

Generate an Interoperability Report for Target vCenter Server Versions

You can perform pre-checks and generate reports that provide product interoperability information for recommended update and upgrade versions of vCenter Server.

To use this feature, you must join the VMware Customer Experience Improvement Program (CEIP). For information about the CEIP, see the "Configuring Customer Experience Improvement Program" section in *vCenter Server and Host Management*.

You can create a vCenter Server pre-update interoperability report to check the compatibility of VMware products in your environment against a pending version of vCenter Server. You can view the list of products in your environment that are

associated with the current vCenter Server version, and whether or not those products are compatible with a suggested update version.

1. In the vSphere Client, select a vCenter Server for which you want to view an interoperability report.
 2. On the **Updates** tab, click **Update**.
 3. Click **Target Version**.
 4. Select the target version from the list of available updates and click **OK**. The list contains minor updates and major upgrade versions of vCenter Server.
 5. Click **Product Interoperability** to view the report.
 6. (Optional) If there are VMware products in your environment that go undetected, you can manually add them to the list of products to check and regenerate the interoperability report.
 - a) Click **Modify Product List**.
 - b) Select the VMware product and version you want to include in the interoperability report.
Repeat this step for each VMware product you want to add to the report.
 - c) To add the selected products to the Update product inventory click **Done**.
 - d) Regenerate the report and verify that the products you added are included in the report.
 7. (Optional) Click **Export** to save the report as a comma-separated values (CSV) file, which allows data to be saved in a tabular format.
- Using the information in the pre-upgrade report, you can apply patches and updates to your vCenter Server environment. See [Patching and Updating vCenter Server 8.0 Deployments](#).
 - You can generate a pre-upgrade report that lists actions that you must address to ensure a successful upgrade of vCenter Server in your environment. See [Generate an Interoperability Report for Target vCenter Server Versions](#).

Generate Pre-Update Reports

You can run pre-checks and generate reports that provide pre-update information. You use this information to make sure that your system meets the minimum software and hardware requirements to ensure a successful upgrade of vCenter Server.

When you generate a pre-update report, the report lists the actions you must address to ensure a successful upgrade of vCenter Server in your environment. The report also provides information about problems that might prevent the completion of a software upgrade or update.

1. In the vSphere Client, select a vCenter Server for which you want to view a pre-update report.
 2. On the **Updates** tab, click **Update**.
 3. Click **Target Version**.
 4. Select the target version from the list of available updates and click **OK**. The list contains minor updates and major upgrade versions of vCenter Server.
 5. Select **Source Pre-Checks** to view pre-update report.
 6. (Optional) Click **Export** to save the report as a comma-separated values (CSV) file, which allows data to be saved in a tabular format.
- Using the information in the pre-upgrade report, you can apply patches and updates to your vCenter Server environment. See [Patching and Updating vCenter Server 8.0 Deployments](#).
 - You can generate an interoperability report that lists available VMware products in your environment, and the versions compatible with the vCenter Server version for which you are generating the report. See [Generate an Interoperability Report for Target vCenter Server Versions](#).

Patching and Updating vCenter Server8.0 Deployments

You can update the vCenter Server with patches by using the `software-packages` utility available in the vCenter Server appliance shell.

Patching vCenter Server

VMware regularly releases patches for the vCenter Server. You can use the Appliance Management Interface or the appliance shell to apply patches to a vCenter Server.

VMware makes patches available on a monthly basis. These patches can only be applied in between major releases of vCenter Server. For example, patches released for the initial release of vCenter Server8.0, are not applicable to vCenter Server8.0 Update 1, as any patches previously made available will be included with the Update 1 release.

These patches can be for core product functionality, other packages in the vCenter Server such as Photon, or both.

Note: You must use only the patches provided by VMware to update the packages in your vCenter Server. Updating these packages through any other means may impact the product functionality.

VMware distributes the available patches in two forms, one for ISO-based and one for URL-based models of patching.

- You can download the patch ISO images from <https://www.broadcom.com/support>. VMware publishes a single type of ISO image that contains patches.

Download Filename	Description
<code>VMware-vCenter-Server-Appliance-product_version-build_number-patch-FP.iso</code>	Full product patch for the appliance, which contains the VMware software patches and the fixes related to security and third-party products (e.g. JRE and Photon OS components).

- You can configure the vCenter Server appliance to use a repository URL as a source of available patches.

Note:

The default online depot URLs are now unique for each customer. To update the default URLs of the online depot, see the following KB articles: [KB 390098](#) and [KB 390121](#).

Patching vCenter Server Using the vCenter Server Management Interface

You can use the vCenter Server Management Interface to view the installed patches, check for new patches and install them, and configure automatic checks for available patches.

To perform ISO-based patching, you download an ISO image, attach the ISO image to the CD/DVD drive of the appliance, check for available patches in the ISO image, and install the patches.

To perform URL-based patching, you check for available patches in a repository URL and install the patches. The vCenter Server is preset with a default VMware repository URL for the build profile of the appliance. You can configure the appliance to use the default VMware repository URL or a custom repository URL, for example, a repository URL that you previously built on a local Web server running within your data center.

Automatic Creation of LVM Snapshot

The vCenter Server patching orchestrator automatically creates a Logical Volume Manager (LVM) snapshot of the vCenter Server before performing a patch installation. If the patching fails, the orchestrator provides options to either continue patching or roll back to the automatically created LVM snapshot.

The roll back option is only available when you are patching vCenter Server 8.0 Update 2 and higher versions.


When you initiate the installation of a patch, prechecks are performed to verify that a snapshot can be taken and there is sufficient disk space for the LVM snapshot. If the precheck returns an error, the patching workflow exits.

After prechecks run successfully, the VMware Directory Services (`vmmdir`) is stopped and the orchestrator takes an image-based backup (LVM snapshot). The snapshot is saved as snapshot disk on the vCenter Server Appliance. If there is a failure while taking the snapshot, the `vmmdir` is started in normal replication mode, an error is thrown and the patching workflow exits and the snapshot is cleaned up.

If the snapshot is taken successfully, the patching resumes with `vmmdir` in standalone mode.

If an error occurs during patching, the vCenter Server Management Interface displays options to resume or rollback the process. You can select **Resume the vCenter installation** and correct the problem preventing the patching to finish staging the patches or select **Revert to the last vCenter backup version** to roll back.

If you choose to revert, a patching rollback API is called. The API checks if an LVM snapshot is available and reverts to the snapshot. After reverting, the snapshot is removed to reclaim the disk space and `vmmdir` is set to normal replication mode. The API also performs a cleanup of any changes made by the patching install workflow before creating the backup. If a valid snapshot is not available, the installation exits with an error.

Figure 52: Workflow for automatic creation of LVM snapshot and roll back

Log In to the vCenter Server Management Interface

Log in to the vCenter Server Management Interface to access the vCenter Server configuration settings.

Verify that the vCenter Server is successfully deployed and running.

Note: The login session expires if you leave the vCenter Server Management Interface idle for 10 minutes.

1. In a Web browser, go to the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
2. Log in as root.
The default root password is the password that you set while deploying vCenter Server.

Check for and Stage Patches to the vCenter Server Appliance

Before you install available patches, you can stage the patches to the appliance. You can use the vCenter Server Management Interface to stage patches either from a local repository by attaching an ISO image to the appliance, or from a remote repository directly by using a repository URL.

- If you are staging patches from an ISO image that you previously downloaded from <https://my.vmware.com/group/vmware/patch>, you must attach the ISO image to the CD/DVD drive of the vCenter Server. You can configure the ISO image as a datastore ISO file for the CD/DVD drive of the appliance using the vSphere Client. See *vSphere Virtual Machine Administration*.
- If you are staging patches from a remote repository, verify that you have configured the repository settings and that the current repository URL is accessible. See [Configure URL-Based Patching](#).

During the process of staging, the vCenter Server Management Interface validates that a patch is a VMware patch, that the staging area has enough free space, and that the patches are not altered. Only new patches or patches for existing packages that can be upgraded are staged.

If you are patching vCenter Server 8.0 Update 2 and higher versions, the patching orchestrator automatically creates a backup of the vCenter Server before performing a patch installation. For more information, see [Automatic Creation of LVM Snapshot](#).

If a problem that prevents the successful staging of the patches occurs, vCenter Server suspends the staging process. Review the error messages, correct the problem, and, in many cases, you can resume staging the patches from the point at which vCenter Server encountered the problem.

1. Log into the vCenter Server Management Interface as root.
The default root password is the password you set when deploying vCenter Server.
2. Click **Update**.
3. Click **Check Updates** and select a source.

Option	Description
Check URL	Scans the configured repository URL for available patches.
Check CDROM	Scans the ISO image that you attached to the CD/DVD drive of the appliance for available patches.

In the Available Updates pane, you can view the details about the available patches in the source that you selected.

Important: Some updates might require a reboot of the system. You can see information about these updates in the Available Updates pane.

4. You can run a pre-check of an update to verify that it is compatible with your current deployment.
5. Click the staging option you would like to use.

Option	Description
Stage	Stages the selected patches to the vCenter Server appliance for installation at a later time.
Stage and Install	Stages and installs the selected patches to the vCenter Server appliance. For more information on installing patches, see Install vCenter Server Patches .
Unstage	Unstages the selected patches.
Resume the vCenter installation	If vCenter Server encounters a problem staging the patches, the vCenter Server Management Interface displays the Resume button. Correct the problem preventing the patches from staging, and click Resume to finish staging the patches. The staging process resumes from the point at which vCenter Server encountered the problem.
Revert to last vCenter backup option	Select this option to roll back to the automatically created snapshot. If you choose to revert, a patching rollback API is called. The API checks if a snapshot is available and reverts to the snapshot. Note: The automatic creation of a backup is available only when you are patching vCenter Server 8.0 Update 2 and higher versions.

If you decided to stage the available patches for installation at a later time, you can now install. See [Install vCenter Server Patches](#).

Configure the Repository for URL-Based Patching

For URL-based patching, by default the vCenter Server appliance is configured to use the default VMware repository URL that is preset for the build profile of the appliance. You can configure a custom repository URL as the current source of patches for your environment's requirements.

Log in to the vCenter Server Management Interface as root.

By default the current repository for URL-based patching is the default VMware repository URL.

If vCenter Server is not connected to the Internet or if your security policy requires it, you can build and configure a custom repository. The custom patching repository runs on a local Web server within your data center and replicates the data from the default repository. Optionally, you can set up an authentication policy for accessing the Web server that hosts the custom patching repository.

1. If you want to configure a custom repository URL, build the repository on your local web server.
2. Log in to Broadcom Support.
3. Download the vCenter Server appliance ISO image. See the Knowledge Base article <https://knowledge.broadcom.com/external/article/366685/vmware-vsphere-downloads-vmware-converte.html>.
4. Confirm that the md5sum is correct by using an MD5 checksum tool.
5. On your Web server, create a repository directory under the root.
For example, create the `vc_update_repo` directory.
6. Extract the ZIP file into the repository directory.
The extracted files are in the manifest and package-pool subdirectories.

7. In the vCenter Server Management Interface, click **Update**.

8. Click **Settings**.

9. Select the Repository settings.

Option	Description
Default repository	Uses the default VMware repository URL that is preset for the build profile of the appliance.
Specified repository	Uses a custom repository. You must enter the repository URL, for example, <code>https://web_server_name.your_company.com/vc_update_repo</code>. The repository URL must use a secure protocol such as HTTPS or FTPS.

10. If the specified repository requires authentication, enter the user name and password.

11. (Optional) If you do not want to perform a security certificate check, deselect the **Check Certificate** check box.

If you trust the URL for the repository, you can choose to bypass the certificate check for the repository URL.

12. Click **OK**.

Install vCenter Server Patches

Install vCenter Server Patches

You can check for and install patches either from an ISO image or directly from a repository URL.

- Log in to the vCenter Server Management Interface as root.
- Before you can install available patches, you check for new patches and stage the patches to the vCenter Server appliance. See [Check for and Stage Patches to the vCenter Server Appliance](#).
- If you are patching the appliance from an ISO image that you previously downloaded from <https://my.vmware.com/group/vmware/patch>, you must attach the ISO image to the CD/DVD drive of the vCenter Server appliance. You can configure the ISO image as a datastore ISO file for the CD/DVD drive of the appliance by using the vSphere Client. See *vSphere Virtual Machine Administration*.
- If you are patching the appliance from a repository URL, verify that you have configured the repository settings and that the current repository URL is accessible. See [Configure the Repository for URL-Based Patching](#).
- Create an image-based backup and take a powered-off snapshot of the vCenter Server Appliance you are patching as a precaution in case there is a failure during the patching process.

Important: The services running in the vCenter Server appliance become unavailable during the installation of the patches. You must perform this procedure during a maintenance period. As a precaution if there is a failure, you can back up the vCenter Server. For information on backing up and restoring vCenter Server, see *vCenter Server Installation and Setup*.

1. In the vCenter Server Management Interface, click **Update**.

In the **Current version details** pane, you can view the vCenter Server version and build number.

In the **Available updates** pane, you can view the available updates with update priority and severity.

The update priority indicates how soon you must install the update. The values include:

- **HIGH** - Install as soon as possible.
- **MEDIUM** - Install at the earliest convenience.
- **LOW** - Install at your discretion.

The **Update severity** defines the severity of the issues fixed in the update. The values include the following:

- **CRITICAL** - Vulnerabilities that can be exploited by an unauthenticated attacker from the Internet or those that break the guest/host Operating System isolation. The exploitation results in the complete compromise of confidentiality, integrity, and availability of user data and/or processing resources without user interaction. Exploitation could be leveraged to propagate an Internet worm or execute arbitrary code between Virtual Machines and/or the Host Operating System.
 - **IMPORTANT** - Vulnerabilities that are not rated critical but whose exploitation results in the complete compromise of confidentiality and/or integrity of user data and/or processing resources through user assistance or by authenticated attackers. This rating also applies to those vulnerabilities which could lead to the complete compromise of availability when exploitation is by a remote unauthenticated attacker from the Internet or through a breach of virtual machine isolation.
 - **MODERATE** - Vulnerabilities where the ability to exploit is mitigated to a significant degree by configuration or difficulty of exploitation, but in certain deployment scenarios could still lead to the compromise of confidentiality, integrity, or availability of user data and/or processing resources.
 - **LOW** - All other issues that have a security impact. Vulnerabilities where exploitation is believed to be extremely difficult, or where successful exploitation would have minimal impact.
2. Select the range of staged patches to apply and click **Install**.
Important: Some updates might require a reboot of the system. You can see information about these updates in the Available Updates pane.
 3. Read and accept the End User License Agreement.
 4. A system pre-check verifies that the patches can be successfully installed with the provided information.
If the pre-check discovers missing or incorrect information, or other problems preventing a successful installation, you are prompted to correct the problem and resume the installation.
 5. After the installation finishes, click **OK**.
 6. If the patch installation requires the appliance to reboot, click **Summary**, and click **Reboot** to reset the appliance.

In the Available Updates pane, you can see the changed update status of the vCenter Server appliance.

Enable Automatic Checks for vCenter Server Patches

You can configure vCenter Server to perform automatic checks for available patches in the configured repository URL at a regular interval.

- Log in to the vCenter Server Management Interface as root.
 - Verify that you have configured the repository settings and that the current repository URL is accessible. See [Configure the Repository for URL-Based Patching](#).
1. In the vCenter Server Management Interface, click **Update**.
 2. Click **Settings**.
 3. Select **Check for updates automatically**, and select the day and time in UTC to perform automatic checks for available patches.
 4. Click **OK**.

vCenter Server appliance performs regular checks for available patches in the configured repository URL. In the Available Updates pane, you can view information about the available patches. You can also view the vCenter Server health status for notifications about available patches. See *vCenter Server Configuration*.

Patching the vCenter Server Appliance by Using the Appliance Shell

You can use the `software-packages` utility in the appliance shell of a vCenter Server appliance to see the installed patches, stage new patches, and install new patches.

To perform ISO-based patching, you download an ISO image, mount the ISO image to the CD/DVD drive of the appliance, optionally stage the available patches from the ISO image to the appliance, and install the patches. For steps to mount the ISO image to the CD/DVD drive, see [Configure a Datastore ISO File for the CD/DVD Drive](#) section in the vSphere Virtual Machine Administration.

To perform URL-based patching, you optionally stage the available patches from a repository URL to the appliance and install the patches. The vCenter Server appliance is preset with a default VMware repository URL for the build profile of the appliance. You can use the `update.set` command to configure the appliance to use the default VMware repository URL or a custom repository URL, for example, a repository URL that you previously built on a local Web server running within your data center. You can also use the `proxy.set` command to configure a proxy server for the connection between the vCenter Server appliance and the repository URL.

Related Links

[Patching a vCenter High Availability Environment on page 1405](#)

To patch a vCenter Server Appliance configured in a vCenter High Availability (HA) cluster, you must remove the vCenter Server High Availability configuration and delete the Passive and Witness nodes. After patching the vCenter Server Appliance, you must re-create your vCenter Server High Availability clusters.

View a List of All Installed Patches in the vCenter Server Appliance

You can use the `software-packages` utility to see a list of the patches currently applied to the vCenter Server appliance. You can also view the list of the installed patches in chronological order and details about a specific patch.

1. Access the appliance shell and log in as a user who has a super administrator role.
The default user with a super administrator role is `root`.
2. To view the full list of patches and software packages installed in the vCenter Server appliance, run the following command:

```
software-packages list
```

3. To view all patches applied to the vCenter Server appliance in chronological order, run the following command:

```
software-packages list --history
```

You see the list in chronological order. A single patch in this list can be an update of multiple different packages.

4. To view details about a specific patch, run the following command:

```
software-packages list --patch patch_name
```

For example, if you want to view the details about the `VMware-vCenter-Server-Appliance-Patch1` patch, run the following command:

```
software-packages list --patch VMware-vCenter-Server-Appliance-Patch1
```

You can see the complete list of details about the patch, such as vendor, description, and installation date.

Configure URL-Based Patching

For URL-based patching, the vCenter Server appliance is preset with a default VMware repository URL for the build profile of the appliance. You can use the `update.set` command to configure the appliance to use the default or a custom repository URL as the current source of patches and enable automatic checks for patches.

By default the current repository for URL-based patching is the default VMware repository URL.

Note: You can use the `proxy.set` command to configure a proxy server for the connection between vCenter Server and the repository URL. For more information about the API commands in the appliance shell, see *vCenter Server Configuration*.

If vCenter Server is not connected to the Internet or if your security policy requires it, you can build and configure a custom repository. The custom patching repository runs on a local Web server within your data center and replicates the data from the default repository. Optionally, you can set up an authentication policy for accessing the Web server that hosts the custom patching repository.

1. If you want to configure a custom repository URL, build the repository on your local web server.
2. Access the appliance shell and log in as a user who has a super administrator role.
The default user with a super administrator role is root.
3. To see information about the current URL-based patching settings, run the `update.get` command.
You can see information about the current repository URL, the default repository URL, the time at which the appliance last checked for patches, the time at which the appliance last installed patches, and the current configuration of automatic checks for patches.

4. Configure the current repository for URL-based patching.

- To configure the appliance to use the default VMware repository URL, run the following command:

```
update.set --currentURL default
```

- To configure the appliance to use a custom repository URL, run the following command:

```
update.set --currentURL https://web_server_name.your_company.com/vc_update_repo [--username username]
[--password password]
```

The square brackets `[]` enclose the command options.

If the custom repository requires authentication, use the `--username username` and `--password password` options.

5. To activate automatic checks for vCenter Server appliance patches in the current repository URL at regular intervals, run the following command:

```
update.set --CheckUpdates enabled [--day day] [--time HH:MM:SS]
```

The square brackets `[]` enclose the command options.

Use the `--day day` option to set the day for performing the regular checks for patches. You can set a particular day of the week, for example, `Monday`, or `Everyday`. The default value is `Everyday`.

Use the `--time HH:MM:SS` option to set the time in UTC for performing the regular checks for patches. The default value is `00:00:00`.

The appliance performs regular checks for available patches in the current repository URL.

6. To deactivate automatic checks for vCenter Server patches, run the following command:

```
update.set --CheckUpdates disabled
```

If you configured the appliance to perform automatic checks for available patches, you can regularly view the vCenter Server appliance health status for notifications about available patches. See *vCenter Server Configuration*.

Stage Patches to the vCenter Server Appliance

Before you install available patches, you can stage the patches to the appliance. You can use the `software-packages` utility to stage patches either from a local repository by attaching an ISO image to the appliance, or from a remote repository directly by using a repository URL.

- If you are staging patches from an ISO image that you previously downloaded from <https://my.vmware.com/group/vmware/patch>, you must attach the ISO image to the CD/DVD drive of the vCenter Server appliance. You can

configure the ISO image as a datastore ISO file for the CD/DVD drive of the appliance by using the vSphere Client. See *vSphere Virtual Machine Administration*.

- If you are staging patches from a remote repository, verify that you have configured the repository settings and that the current repository URL is accessible. See [Configure URL-Based Patching](#).

1. Access the appliance shell and log in as a user who has a super administrator role.

The default user with a super administrator role is root.

2. Stage the patches.

- To stage the patches included in the attached ISO image, run the following command:

```
software-packages stage --iso
```

- To stage the patches included in the current repository URL, run the following command:

```
software-packages stage --url
```

By default the current repository URL is the default VMware repository URL.

If you want to stage only the third-party patches, use the `--thirdParty` option.

- To stage the patches included in a repository URL that is not currently configured in the appliance, run the following command:

```
software-packages stage --url URL_of_the_repository
```

If you want to stage only the third-party patches, use the `--thirdParty` option.

If you want to directly accept the End User License Agreement, use the `--acceptEulas` option.

For example, to stage only the third-party patches from the current repository URL with directly accepting the End User License Agreement, run the following command:

```
software-packages stage --url --thirdParty --acceptEulas
```

In the process of staging, the command validates that a patch is a VMware patch, that the staging area has enough free space, and that the patches are not altered. Only completely new patches or patches for existing packages that can be upgraded are staged.

3. Optional: To see information about the staged patches, run the following command:

```
software-packages list --staged
```

Each patch includes a metadata file that contains information such as patch version, product name, whether a restart of the system is required, and so on.

4. Optional: To view a list of the staged patches, run the following command:

```
software-packages list --staged --verbose
```

5. Optional: To unstage the staged patches, run the following command:

```
software-packages unstage
```

All directories and files generated by the staging process are removed.

Install the staged patches. See [Install vCenter Server Patches](#).

Important: If you staged the patches from an ISO image, keep the ISO image attached to the CD/DVD drive of the appliance. The ISO image must be attached to the CD/DVD drive of the appliance throughout the staging and installation processes.

Install vCenter Server Patches

You can use the `software-packages` utility to install the staged patches. You can also use the `software-packages` utility to install patches directly from an attached ISO image or repository URL without staging the patch payload.

- If you are installing staged patches, verify that you staged the correct patch payload. See [Stage Patches to the vCenter Server Appliance](#).
- If you are installing patches that you previously staged from an ISO image, verify that the ISO image is attached to the CD/DVD drive of the vCenter Server appliance. See [Stage Patches to the vCenter Server Appliance](#).
- If you are installing patches directly from an ISO image that you previously downloaded from <https://my.vmware.com/group/vmware/patch>, you must attach the ISO image to the CD/DVD drive of the vCenter Server Appliance. You can configure the ISO image as a datastore ISO file for the CD/DVD drive of the appliance by using the vSphere Client. See *vSphere Virtual Machine Administration*.
- If you are installing patches directly from a repository, verify that you have configured the repository settings and that the current repository URL is accessible. See [Configure URL-Based Patching](#).
- Create an image-based backup and take a powered-off snapshot of the vCenter Server Appliance you are patching as a precaution in case there is a failure during the patching process.

Important: The services running in the appliance become unavailable during the installation of the patches. You must perform this procedure during a maintenance period. As a precaution in case of failure, you can back up the vCenter Server appliance. For information about backing up and restoring vCenter Server, see *vCenter Server Installation and Setup*.

1. Access the appliance shell and log in as a user who has a super administrator role.

The default user with a super administrator role is root.

2. Install the patches.

- To install staged patches, run the following command:

```
software-packages install --staged
```

- To install patches directly from an attached ISO image, run the following command:

```
software-packages install --iso
```

- To install patches directly from the current repository URL, run the following command:

```
software-packages install --url
```

By default the current repository URL is the default VMware repository URL.

- To install patches directly from a repository URL that is not currently configured, run the following command:

```
software-packages install --url URL_of_the_repository
```

If you want to directly accept the End User License Agreement, use the `--acceptEulas` option.

For example, to install patches from the current repository URL without staging the patches with directly accepting the End User License Agreement, run the following command:

```
software-packages install --url --acceptEulas
```

3. If the patch installation requires a reboot of the appliance, run the following command to reset the appliance.

```
shutdown now -r "patch reboot"
```

Patching a vCenter High Availability Environment

To patch a vCenter Server Appliance configured in a vCenter High Availability (HA) cluster, you must remove the vCenter Server High Availability configuration and delete the Passive and Witness nodes. After patching the vCenter Server Appliance, you must re-create your vCenter Server High Availability clusters.

This procedure describes how to remove the vCenter HA configuration.

1. In the vSphere Client, log in to the vCenter Server from which you want to remove the vCenter HA configuration.
2. Click the **Configure** tab, and select **vCenter HA**.
3. Select **Remove vCenter HA** .
 - The vCenter HA cluster's configuration is removed from the Active, Passive, and Witness nodes.
 - The Active node continues to run as a standalone vCenter Server Appliance.
4. Delete the Passive and Witness nodes.

Patch the vCenter Server Appliance as described in either [Patching vCenter Server Using the vCenter Server Management Interface](#) or [Patching the vCenter Server Appliance by Using the Appliance Shell](#).

After patching the vCenter Server Appliance, you can configure vCenter HA. For information about configuring vCenter HA, see *vSphere Availability*.

Related Links

[Patching the vCenter Server Appliance by Using the Appliance Shell on page 1401](#)

You can use the `software-packages` utility in the appliance shell of a vCenter Server appliance to see the installed patches, stage new patches, and install new patches.

Troubleshooting a vSphere Upgrade

The installation and upgrade software enables you to identify problems on the host machine that can cause an installation, upgrade, or migration to fail.

For interactive installations, upgrades, and migrations, the errors or warnings are displayed on the final panel of the installer, where you are asked to confirm or cancel the installation or upgrade. For scripted installations, upgrades, or migrations, the errors or warnings are written to the installation log file. You can also consult the product release notes for known problems.

vSphere Update Manager provides custom messages for these errors or warnings. To see the original errors and warnings returned by the precheck script during an Update Manager host upgrade scan, review the Update Manager log file `vmware-vum-server-log4cpp.log`.

The *vSphere Upgrade* guide describes how to use VMware products and their features. If you encounter problems or error situations that are not described in this guide, you may find a solution in VMware Knowledge Base. You can also use VMware Community Forums to find others with same problem or ask for help, or you can open Support Request to get help from VMware service professional.

Collect Installation Logs for vCenter Server Appliance

You can collect installation log files and check these files to identify the source of a failure if vCenter Server appliance stops responding during initial startup.

1. Access the vCenter Server appliance shell.

Option	Description
If you have direct access to the appliance	Press Alt+F1.
To connect remotely	Use SSH or another remote console connection to start a session to the appliance.

2. Enter a user name and password that the appliance recognizes.
3. In the appliance shell, run the `pi shell` command to access the Bash shell.
4. In the Bash shell, run the `vc-support.sh` script to generate a support bundle.

This command generates a `.tgz` file in `/storage/log`.

5. Export the generated support bundle to the `user@x.x.x.x:/tmp` folder.

```
scp /var/tmp/vc-etco-vm-vlan11-dhcp-63-151.eng.vmware.com-2014-02-28--21.11.tgz user@x.x.x.x:/tmp
```

6. Determine which `firstboot` script failed.

```
cat /var/log/firstboot/firstbootStatus.json
```

To identify potential causes of the failure, examine the log file of the `firstboot` script that failed.

Errors and Warnings Returned by the Installation and Upgrade Precheck Script

The installation and upgrade precheck script runs tests to identify problems on the host machine that can cause an installation, upgrade, or migration to fail.

For interactive installations, upgrades, and migrations, the errors or warnings are displayed on the final screen of the GUI installer, where you are asked to confirm or cancel the installation or upgrade. For scripted installations, upgrades, or migrations, the errors and warnings are written to the installation log file.

vSphere Update Manager provides custom messages for these errors or warnings. To see the original errors and warnings returned by the precheck script during an Update Manager host upgrade scan, review the Update Manager log file `vmware-vum-server-log4cpp.log`.

Table 93: Error and Warning Codes That Are Returned by the Installation and Upgrade Precheck Script

Error or Warning	Description
64BIT_LONGMODESTATUS	The host processor must be 64-bit.
COS_NETWORKING	Warning. An IPv4 address was found on an enabled service console virtual NIC that has no corresponding address in the same subnet in the VMkernel. A separate warning appears for each such occurrence.
CPU_CORES	The host must have at least two cores.
DISTRIBUTED_VIRTUAL_SWITCH	If the Cisco Virtual Ethernet Module (VEM) software is found on the host, the test checks that the upgrade also contains the VEM software. The test also determines whether the upgrade supports the same version of the Cisco Virtual Supervisor Module (VSM) as the existing version on the host. If the software is missing or is compatible with a different version of the VSM, the test returns a warning. The result indicates which version of the VEM software was expected on the upgrade ISO and which versions, if any, were found. You can use ESXi Image Builder CLI to create a custom installation ISO that includes the appropriate version of the VEM software.
HARDWARE_VIRTUALIZATION	Warning. If the host processor does not have hardware virtualization, or if hardware virtualization is not turned on in the host BIOS, host performance reduces. You can enable hardware virtualization in the host machine boot options panel. See your hardware vendor's documentation.

Error or Warning	Description
MD5_ROOT_PASSWORD	This test checks that the root password is encoded in MD5 format. If a password is not encoded in MD5 format, it might be significant only to eight characters. In this case, any characters after the first eight are no longer authenticated after the upgrade, which can create a security problem. To work around this problem, see Knowledge Base article KB 1024500 .
MEMORY_SIZE	The host requires the specified amount of memory to upgrade.
PACKAGE_COMPLIANCE	vSphere Update Manager upgrade to vSphere Lifecycle Manager only. This test checks the existing software on the host against the software contained on the upgrade ISO to determine whether the host has been successfully upgraded. If any of the packages are missing or are an older version than the package on the upgrade ISO, the test returns an error. The test results indicate which software was found on the host and which software was found on the upgrade ISO.
PARTITION_LAYOUT	You can upgrade or migrate software only if at most one VMFS partition on the disk is being upgraded. The VMFS partition must begin after sector 1843200.
POWERPATH	This test checks for the installation of EMC PowerPath software, consisting of a CIM module and a kernel module. If either of these components is found on the host, the test checks that matching components, such as CIM or the VMkernel and module, also exist in the upgrade. If they do not exist, the test returns a warning that indicates which PowerPath components were expected on the upgrade ISO and which, if any, were found.
PRECHECK_INITIALIZE	This test checks that the precheck script can be run.
SANE_ESX_CONF	The <code>/etc/vmware/esx.conf</code> file must exist on the host.
SPACE_AVAIL_ISO	vSphere Update Manager only. The host disk must have enough free space to store the contents of the installer CD or DVD.
SPACE_AVAIL_CONFIG	vSphere Update Manager upgrade to vSphere Lifecycle Manager only. The host disk must have enough free space to store the legacy configuration between reboots.
SUPPORTED_ESX_VERSION	You can upgrade or migrate to ESXi8.0 only from version 6.0 ESXi hosts.
TBOOT_REQUIRED	This message applies only to vSphere Update Manager upgrade to vSphere Lifecycle Manager only. The upgrade fails with this error when the host system is running in trusted boot mode (tboot), but the ESXi upgrade ISO does not contain any tboot VIBs. This test prevents an upgrade that can make the host less secure.
UNSUPPORTED_DEVICES	Warning. This test checks for unsupported devices. Some PCI devices are not supported in ESXi8.0.

Error or Warning	Description
UPDATE_PENDING	This test checks the host for VIB installations that require a reboot. This test fails if one or more such VIBs are installed, but the host has not yet been rebooted. In these conditions, the precheck script is unable to determine reliably which packages are currently installed on the host. If this test fails, it may not be safe to rely on the rest of the precheck tests to determine whether an upgrade is safe. If you encounter this error, restart the host and retry the upgrade.

Upgrade Issues with vCenter Server Containing Host Profiles

Most common issues that might occur during a vCenter Server upgrade to version 8.0, that contains host profiles.

- For issues occurring during a vCenter Server upgrade or ESXi upgrade, see *Troubleshooting a vSphere Upgrade*.
- If upgrading vCenter Server 7.0 or 7.0, containing host profiles with version earlier than 7.0, results with a failure, see [KB 52932](#).
- For error `There is no suitable host in the inventory as reference host for the profile Host Profile`. The profile does not have any associated reference host, see [KB 2150534](#).
- If an error occurs when you import a host profile to an empty vCenter Server inventory, see *vSphere Host Profiles* for Reference Host is Unavailable.
- If a host profile compliance check fails for NFS datastore, see *vSphere Host Profiles* for Host Profile without NFS Datastore.
- If compliance check fails with an error for the `UserVars.ESXiVPsDisabledProtocols` option, when an ESXi host upgraded to version 8.0 is attached to a host profile with version 7.0, see VMware vSphere 8.0 Release Notes.

Roll Back a vCenter Server Instance on Windows When vCenter Server Upgrade Fails

You can roll back or restore a vCenter Server instance on Windows when an upgrade of vCenter Server with an external Platform Services Controller fails after the export stage and the legacy environment has been uninstalled.

The roll back or restore of vCenter Server applies when all of the following conditions apply:

- You must have access to the vCenter Server for Windows machine.
- The vCenter Server instance is attached to an external Platform Services Controller.
- The upgrade of the vCenter Server instance attached to the Platform Services Controller instance is in a failed state after the export stage and uninstallation of the legacy vCenter Server.
- Ensure that vCenter Server rollback happened properly in case of upgrade failure and that no stale failed upgrade log entries remain.

For Rollback Method 1:

- To clean up vCenter Server appliance 7.0 data from the legacy Platform Services Controller, see [KB 2106736](#).
- Use a vCenter Server database snapshot taken before the start of the vCenter Server upgrade.

For Rollback Method 2:

- Use a powered off snapshot of the vCenter Server before the vCenter Server upgrade.
- Use a Platform Services Controller snapshot taken after the Platform Services Controller node upgrade and before the start of the vCenter Server upgrade.
- Use a vCenter Server snapshot taken after the Platform Services Controller upgrade and before the start of the vCenter Server upgrade.

You can restore the legacy vCenter Server using Rollback Method 1 or Rollback Method 2.

- Use Rollback Method 1.
 - a. Manually clean up vCenter Server appliance data from Platform Services Controller.
 - b. Restore the legacy vCenter Server database from a backup which was taken before the upgrade.
 - c. Point the legacy vCenter Server instance to the Platform Services Controller, and also point to the database with the restored data.
 - d. Ensure that the vCenter Server services are up and running.
- Use Rollback Method 2.
 - a. Restore the Platform Services Controller instance from a snapshot to the point where you were about to start vCenter Server upgrade. You can use a backup for a Windows configuration or use another backup and restore approach to revert the snapshot.
 - b. Restore the vCenter Server instance from a snapshot.
 - c. Restore the vCenter Server database from a snapshot.
 - d. Ensure that the vCenter Server services are up and running.

For Rollback Method 2, you will lose all data written to Platform Services Controller after the vCenter Server upgrade has been started when you restore from the Platform Services Controller snapshot taken before that point in time.

Collect Logs to Troubleshoot ESXi Hosts

You can collect installation or upgrade log files for ESXi that help you identify the source of the failure if an installation or upgrade fails.

1. Enter the `vm-support` command in the ESXi Shell or through SSH.
2. Navigate to the `/var/tmp/` directory.
3. Retrieve the log files from the `.tgz` file.

vSphere Authentication

The *vSphere Authentication* documentation provides information to help you perform common tasks such as certificate management and vCenter Single Sign-On configuration.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

vSphere Authentication explains how you can manage certificates for vCenter Server and related services, and set up authentication with vCenter Single Sign-On.

Table 94: vSphere Authentication Highlights

Topics	Content Highlights
Getting Started with Authentication	- Managing authentication services. - Managing vCenter Server using the vCenter Server Management Interface.
vSphere Security Certificates	- Certificate model, and options for replacing certificates. - Replace certificates from the UI (simple cases). - Replace certificates using the Certificate Manager utility. - Replace certificates using the CLI (complex situations). - Certificate management CLI reference.
vSphere Authentication with vCenter Single Sign-On	- Architecture of the authentication process. - How to add identity sources so users in your domain can authenticate. - Two-factor authentication. - Managing users, groups, and policies. - vCenter Server Identity Provider Federation

What Happened to the Platform Services Controller

Beginning in vSphere 7.0, deploying a new vCenter Server or upgrading to vCenter Server 7.0 requires the use of the vCenter Server appliance, a preconfigured virtual machine optimized for running vCenter Server. The new vCenter Server contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, tags, and licensing. It is no longer necessary nor possible to deploy and use an external Platform Services Controller. All Platform Services Controller services are consolidated into vCenter Server, and deployment and administration are simplified.

As these services are now part of vCenter Server, they are no longer described as a part of Platform Services Controller. In vSphere 7.0, the *vSphere Authentication* publication replaces the *Platform Services Controller Administration* publication. The new publication contains complete information about authentication and certificate management. For information about upgrading or migrating from vSphere 6.5 and 6.7 deployments using an existing external Platform Services Controller to vSphere 7.0 using vCenter Server appliance, see the *vSphere Upgrade* documentation.

Related Documentation

A companion document, *vSphere Security*, describes available security features and the measures that you can take to safeguard your environment from attack. That document also explains how you can set up permissions, and includes a reference to privileges.

In addition to these documents, VMware publishes the *vSphere Security Configuration Guide* (formerly known as the *Hardening Guide*) for each release of vSphere, accessible at <https://core.vmware.com/security>. The *vSphere Security Configuration Guide* contains guidelines on security settings that can or should be set by the customer, and security settings delivered by VMware that should be audited by the customer to ensure that they are still set to default.

Intended Audience

This information is intended for administrators who want to configure vCenter Server authentication and manage certificates. The information is written for experienced Linux system administrators who are familiar with virtual machine technology and data center operations.

Getting Started with vSphere Certificate Management and Authentication

vSphere provides common infrastructure services to manage certificates for both vCenter Server and ESXi components, and to manage authentication with vCenter Single Sign-On.

How Do You Manage vSphere Certificates

By default, vSphere enables you to provision vCenter Server components and ESXi hosts with VMware Certificate Authority (VMCA) certificates. You can also use custom certificates, which are stored in the VMware Endpoint Certificate Store (VECS). For more information, see [What Options Do You Have to Manage vSphere Certificates](#).

What Is vCenter Single Sign-On

vCenter Single Sign-On allows vSphere components to communicate with each other through a secure token mechanism. vCenter Single Sign-On uses specific terms and definitions that are important to understand.

Table 95: vCenter Single Sign-On Glossary

Term	Definition
Principal	An entity that can be authenticated, such as a user.
Identity Provider	A service that manages identity sources and authenticates principals. Examples: Microsoft Active Directory Federation Services (AD FS) and vCenter Single Sign-On.
Identity Source (Directory Service)	Stores and manages principals. Principals consist of a collection of attributes about a user or a service account such as name, address, email, and group membership. Examples: Microsoft Active Directory and VMware Directory Service (vmdir).
Authentication	The means of determining whether someone or something is, in fact, who or what it declares itself to be. For example, users are authenticated when they provide their credentials, such as smart cards, user name and correct password, and so on.
Authorization	The process of verifying what objects principals have access to.
Token	A signed collection of data comprising the identity information for a given principal. A token might include not only basic information about the principal such as email address and full name, but also, depending on the token type, the principal's groups and roles.

Term	Definition
vmdir	VMware Directory Service. The internal (local) LDAP repository in vCenter Server that contains user identities, groups, and configuration data.
OAuth 2.0	An open authorization standard that enables the exchange of information among principals and web services without exposing principals' credentials.
OpenID Connect (OIDC)	Authentication protocol based on OAuth 2.0 that augments OAuth with user-identifying information. It is represented by the ID token that the authorization server returns together with the access token during OAuth authentication. vCenter Server uses OIDC capabilities when interacting with Active Directory Federation Services (AD FS), Okta, Microsoft Entra ID, and PingFederate.
System for Cross-domain Identity Management (SCIM)	The standard for automating the exchange of user identity information between identity domains or IT systems.
VMware Identity Services	Starting in version 8.0 Update 1, VMware Identity Services is a built-in container within vCenter Server that you can use for identity federation to external identity providers. It serves as an independent identity broker within vCenter Server and comes with its own set of APIs. Currently, VMware Identity Services support Okta, Microsoft Entra ID, and PingFederate as external identity providers.
Tenant	A VMware Identity Services concept. A tenant provides a logical separation of data from other tenants' data in one and the same virtual environment.
JSON Web Token (JWT)	A token format defined by the OAuth 2.0 specification. A JWT token carries authentication and authorization information about a principal.
Relying party	A relying party "relies" on the authorization server, VMware Identity Services or AD FS, for identity management. For example, through federation, vCenter Server establishes relying party trust to VMware Identity Services or AD FS.
Security Assertion Markup Language (SAML)	An XML-based open standard for exchanging authentication and authorization data between parties that is used by vCenter Server. Principals obtain a SAML token from vCenter Single Sign-On and then send it to the vSphere Automation API endpoint for a session identifier.

What Are the vCenter Single Sign-On Authentication Types

vCenter Single Sign-On uses different types of authentication, depending on whether the built-in vCenter Server identity provider or an external identity provider is involved.

Table 96: vCenter Single Sign-On Authentication Types

Authentication Type	What Acts as the Identity Provider?	Does vCenter Server Handle the Password?	Description
Token-Based Authentication	External identity provider. For example, AD FS.	No	vCenter Server contacts the external identity provider through a particular protocol and obtains a token, which represents a particular user identity.
Simple Authentication	vCenter Server	Yes	The user name and password are passed directly to vCenter Server, which validates the credentials through its identity sources.

Managing vCenter Server Certificates

You manage vCenter Server certificates from the vSphere Client, or by using an API, scripts, or CLIs.

The following table describes the interfaces you can use to manage vCenter Server certificates.

Table 97: Interfaces for Managing vSphere Certificates

Interface	Description
vSphere Client	Web interface (HTML5-based client). See Managing Certificates Using the vSphere Client .
vSphere Automation API	See <i>VMware vSphere Automation SDKs Programming Guide</i> .
Certificate Management utility	Command-line tool that supports Certificate Signing Request (CSR) generation and certificate replacement. See Managing Certificates Using the vSphere Certificate Manager Utility .
CLIs for managing certificate and directory services	Set of commands for managing certificates, the VMware Endpoint Certificate Store (VECS), and VMware Directory Service (vmdir). See vSphere Certificates and Services CLI Command Reference .

Manage vCenter Server Certificates Using the vSphere Client

You can manage vCenter Server certificates from the vSphere Client.

1. Log in to a vCenter Server as a user with administrator privileges in the local vCenter Single Sign-On domain.
The default domain is vsphere.local.
2. Select **Administration**.
3. Under **Certificates**, click **Certificate Management**.
Certificate tabs for the different types of certificates appear.
4. Perform certificate tasks, such as viewing certificate details, renewing or refreshing a certificate, and adding a Trusted Root certificate.
For more information, see [Managing Certificates Using the vSphere Client](#).

Manage vCenter Server Certificates Using CLIs

vCenter Server includes CLIs for generating Certificate Signing Requests (CSRs), managing certificates, and managing services.

Enable SSH login to vCenter Server. You can use the **Access** tab in the vCenter Server Management Interface (https://vcenter_server_ip:5480) for SSH login activation and deactivation.

For example, you can use the `certool` command to generate CSRs and to replace certificates.

Use the CLIs for management tasks that the vSphere Client does not support, or to create custom scripts for your environment.

Table 98: CLIs for Managing vCenter Server Certificates and Associated Services

CLI	Description	Links
<code>certool</code>	Generate and manage certificates and keys. Part of VMware Certificate Authority (VMCA).	certool Initialization Commands Reference
<code>vecs-cli</code>	Manage the contents of VMware Certificate Store instances. Part of VMware Authentication Framework Daemon (VMAFD).	vecs-cli Command Reference
<code>dir-cli</code>	Create and update certificates in VMware Directory Service. Part of VMAFD.	dir-cli Command Reference
<code>sso-config</code>	Update Security Token Service (STS) certificates.	Replace a vCenter Server STS Certificate Using the Command Line
<code>service-control</code>	Command for starting, stopping, and listing services.	Run this command to stop services before running other CLI commands.

1. Log in to the vCenter Server shell.

Usually, you have to be the root or Administrator user. See [vSphere Certificates and Services CLI Command Reference](#) for details.

2. Access a CLI at one of the following default locations.

The required privileges depend on the task that you want to perform. Sometimes, you are prompted for the password twice to safeguard sensitive information.

```
/usr/lib/vmware-vmafd/bin/vecs-cli
/usr/lib/vmware-vmafd/bin/dir-cli
/usr/lib/vmware-vmca/bin/certool
/opt/vmware/bin/sso-config.sh
```

The `service-control` command does not require that you enter the path.

For more information, see [Manual vSphere Certificate Replacement](#).

Managing vCenter Server Authentication Services

You manage authentication services from the vSphere Client, or by using the CLI. You can also manage the vCenter Server Identity Provider Federation configuration process by using an API.

You can manage vCenter Server authentication using different interfaces.

Table 99: Interfaces for Managing vCenter Server Authentication Services

Interface	Description
vSphere Client	Web interface (HTML5-based client).
API	Manage the vCenter Server Identity Provider Federation configuration process.
sso-config	Command-line utility for configuring the vCenter Server built-in identity provider.

Manage vCenter Server Authentication Services Using the vSphere Client

You can manage vCenter Server authentication services from the vSphere Client.

1. Log in to a vCenter Server as a user with administrator privileges in the local vCenter Single Sign-On domain.
The default domain is vsphere.local.
2. Select **Administration**.
3. Under **Single Sign On**, click **Configuration** to manage identity providers and configure password and lockout policies.
For more information, see [vSphere Authentication with vCenter Single Sign-On](#).

Manage vCenter Server Authentication Services Using Scripts

vCenter Server includes a utility, `sso-config`, for managing authentication services.

Enable SSH login to vCenter Server. You can use the **Access Settings** tab in the vCenter Server Management Interface (https://vcenter_server_ip:5480) for SSH login activation and deactivation.

Use the `sso-config` utility for management tasks that the vSphere Client does not support, or to create custom scripts for your environment.

Table 100: CLIs for Managing Authentication and Associated Services

CLI	Description	Links
<code>sso-config</code>	Command-line utility for configuring the vCenter Server built-in identity provider.	Refer to the <code>sso-config</code> help by running <code>sso-config.sh -help</code> , or see the VMware knowledge base article at https://kb.vmware.com/s/article/67304 for usage examples.
<code>service-control</code>	Command for starting, stopping, and listing services.	Run this command to stop services before running other CLI commands. The <code>service-control</code> command does not require that you specify the path.

1. Log in to the vCenter Server shell.
Usually, you have to be the root or Administrator user. See [vSphere Certificates and Services CLI Command Reference](#) for details.
2. Access the `sso-config` utility at following default location.
`/opt/vmware/bin/sso-config.sh`
The required privileges depend on the task that you want to perform. Sometimes, you are prompted for the password twice to safeguard sensitive information.

Managing the vCenter Server

You can manage the vCenter Server from the vCenter Server Management Interface or from the vCenter Server shell. For more information about managing vCenter Server, see *vCenter Server Configuration*.

Table 101: Interfaces for Managing the vCenter Server

Interface	Description
vCenter Server Management Interface	Use this interface to reconfigure the system settings. See Manage vCenter Server Using the Management Interface .
vCenter Server shell	Use this command-line interface to perform service management operations on VMCA, VECS, and VMDIR. See Managing Certificates Using the vSphere Certificate Manager Utility and vSphere Certificates and Services CLI Command Reference .

Manage vCenter Server Using the Management Interface

You can use the vCenter Server Management Interface to configure the system settings.

Settings in the vCenter Server Management Interface include time synchronization, network settings, and SSH login settings. You can also change the root password, join the appliance to an Active Directory domain, and leave an Active Directory domain.

Note: Under the **Networking** pane, the virtual NIC 0 is reserved for management traffic. You cannot reassign the traffic from NIC 0 to another NIC. If you are using VCHA, this traffic uses NIC 1. You can add NICs to the vCenter Server Appliance. See the VMware knowledge base article at <https://kb.vmware.com/article/2147155>.

1. In a browser, go to the Web interface at `https://vcenter_server_ip:5480`.
2. If a warning message about an untrusted SSL certificate appears, resolve the issue based on your company security policy and the browser that you are using.
3. Log in as root.
The default root password is the root password that you set when deploying the vCenter Server.

You see the Summary page of the vCenter Server Management Interface.

Manage vCenter Server Using the vCenter Server Shell

You can use service management utilities and CLIs from the vCenter Server shell. You can use TTY1 to log in to the console, or can use SSH to connect to the shell.

1. Activate SSH login if necessary.
 - a) Log in to the vCenter Server Management Interface at `https://vcenter_server_ip:5480`.
 - b) In the Navigator, select **Access** and click **Edit**.
 - c) Toggle on **Activate SSH Login** and click **OK**.

You can follow the same steps to activate the Bash shell for the vCenter Server.
2. Access the shell.
 - If you have direct access to the vCenter Server console, select **Log in**, and press Enter.
 - To connect remotely, use SSH or another remote console connection to start a session to the vCenter Server.

3. Log in as root with the password that you set when you initially deployed the vCenter Server.
If you changed the root password, use the new password.

Add a vCenter Server to an Active Directory Domain

If you want to add an Active Directory identity source to vCenter Server, you must join the vCenter Server to an Active Directory domain.

If you are unable to use vCenter Server Identity Provider Federation, or Active Directory over LDAPS, vCenter Server supports Integrated Windows Authentication (IWA). To use IWA, you must join the vCenter Server to your Active Directory domain.

1. Using the vSphere Client, log in to vCenter Server as a user with administrator privileges in the local vCenter Single Sign-On domain (vsphere.local by default).
2. Select **Administration**.
3. Expand **Single Sign On** and click **Configuration**.
4. Under the **Identity Provider** tab, click **Active Directory Domain**.
5. Click **Join AD**, enter the domain, optional organizational unit, and user name and password, and click **Join**.
6. Restart vCenter Server.

To attach users and groups from the joined Active Directory domain, add the joined domain as a vCenter Single Sign-On identity source. See [Add or Edit a vCenter Single Sign-On Identity Source](#).

vSphere Security Certificates

vSphere provides security by using certificates to encrypt communications, authenticate services, and sign tokens.

How Does vSphere Use Certificates

vSphere uses certificates to:

- Encrypt communications between two nodes, such as a vCenter Server and an ESXi host.
- Authenticate vSphere services.
- Perform internal actions such as signing tokens.

What Is the VMware Certificate Authority

vSphere's internal certificate authority, VMware Certificate Authority (VMCA), provides all the certificates necessary for vCenter Server and ESXi. VMCA is installed on every vCenter Server host, immediately securing the solution without any other modification. Keeping this default configuration provides the lowest operational overhead for certificate management. vSphere provides a mechanism to renew these certificates in the event they expire.

vSphere also provides a mechanism to replace certain certificates with your own certificates. However, replace only the SSL certificate that provides encryption between nodes, to keep your certificate management overhead low.

What Options Do You Have to Manage vSphere Certificates

The following options are recommended for managing certificates.

Table 102: Recommended Options for Managing vSphere Certificates

Mode	Description	Advantages
VMCA Default Certificates	VMCA provides all the certificates for vCenter Server and ESXi hosts.	Simplest and lowest overhead. VMCA can manage the certificate life cycle for vCenter Server and ESXi hosts.
VMCA Default Certificates with External SSL Certificates (Hybrid Mode)	You replace the vCenter Server SSL certificates, and allow VMCA to manage certificates for solution users and ESXi hosts. Optionally, for high-security conscious deployments, you can replace the ESXi host SSL certificates as well.	Simple and secure. VMCA manages internal certificates but you get the benefit of using your corporate-approved SSL certificates, and having those certificates trusted by your browsers.

What Tools Are Available to Replace vSphere Certificates

You can use the following options to replace the existing certificates.

Table 103: Different Approaches to vSphere Certificate Replacement

Option	See
Use the vSphere Client.	Managing Certificates Using the vSphere Client
Use the vSphere Automation API to manage the life cycle of certificates.	<i>VMware vSphere Automation SDKs Programming Guide</i>
Use the vSphere Certificate Manager utility from the command line.	Managing Certificates Using the vSphere Certificate Manager Utility
Use CLI commands for manual certificate replacement.	vSphere Certificates and Services CLI Command Reference

vSphere Certificate Requirements for Different Solution Paths

Certificate requirements depend on whether you use the VMware Certificate Authority (VMCA) as an intermediate certificate authority or you use custom certificates. Requirements are also different for machine certificates.

Before you begin to change certificates, ensure that all nodes in your vSphere environment are time synchronized.

Note: vSphere deploys only RSA certificates for server authentication and does not support generating ECDSA certificates. vSphere verifies ECDSA certificates presented by other servers. For example, if vSphere connects to a syslog server and the syslog server has an ECDSA certificate, vSphere supports verifying that certificate.

Requirements for All Imported vSphere Certificates

- Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded). The vSphere Client and API still accept a key size up to 16384 bits when generating the Certificate Signing Request.
Note: In vSphere 8.0, you can only generate CSRs with a minimum key length of 3072 bits when using the vSphere Client or the vSphere Certificate Manager. vCenter Server still does accept custom certificates bearing a key length of 2048 bits. In vSphere 8.0 Update 1 and later, you can use the vSphere Client to generate a CSR with a key length of 2048 bits.
Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 bits and 3072 bits.
- PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When you add keys to VECS, they are converted to PKCS8.
- x509 version 3

- SubjectAltName must contain DNS Name=*machine_FQDN*
- CRT format
- Exempting the vpxd-extension solution user certificate, Extended Key Usage can be either empty or contain Server Authentication.

vSphere does not support the following certificates.

- Certificates with wildcards.
- The algorithms md2WithRSAEncryption, md5WithRSAEncryption, RSASSA-PSS, dsaWithSHA1, ecdsa_with_SHA1, and sha1WithRSAEncryption are not supported.
- When creating a custom machine SSL certificate for vCenter Server, Server Authentication and Client Authentication are not supported, and must be removed when using the Microsoft Certificate Authority (CA) templates. For more information, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2112009>.

vSphere Certificate Compliance to RFC 2253

The certificate must be in compliance with RFC 2253.

If you do not generate CSRs using vSphere Certificate Manager, ensure that the CSR includes the following fields.

String	X.500 AttributeType
CN	commonName
L	localityName
ST	stateOrProvinceName
O	organizationName
OU	organizationalUnitName
C	countryName
STREET	streetAddress
DC	domainComponent
UID	userid

If you generate CSRs using vSphere Certificate Manager, you are prompted for the following information, and vSphere Certificate Manager adds the corresponding fields to the CSR file.

- The password of the administrator@vsphere.local user, or for the administrator of the vCenter Single Sign-On domain that you are connecting to.

- Information that vSphere Certificate Manager stores in the `certtool.cfg` file. For most fields, you can accept the default or provide site-specific values. The FQDN of the machine is required.
 - Password for `administrator@vsphere.local`
 - Two-letter country code
 - Company name
 - Organization name
 - Organization unit
 - State
 - Locality
 - IP address (optional)
 - Email
 - Host name, that is, the fully qualified domain name of the machine for which you want to replace the certificate. If the host name does not match the FQDN, certificate replacement does not complete correctly and your environment might end up in an unstable state.
 - IP address of the vCenter Server node on which you run vSphere Certificate Manager.

Note: The OU (organizationalUnitName) field is no longer mandatory.

Certificate Requirements When Using VMCA as an Intermediate Certificate Authority

When you use VMCA as an intermediate CA, the certificates must meet the following requirements.

Certificate Type	Certificate Requirements
Root certificate	- You can use vSphere Certificate Manager to create the CSR. See Generate CSR Using the Certificate Manager and Prepare Root Certificate (Intermediate CA). - If you prefer to create the CSR manually, the certificate that you send to be signed must meet the following requirements. - Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded) - PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When keys are added to VECS, they are converted to PKCS8. - x509 version 3 - The CA extension must be set to true for root certificates, and cert sign must be in the list of requirements. For example: <pre>basicConstraints = critical,CA:true keyUsage = critical,digitalSig- nature,keyCertSign</pre> - CRL signing must be enabled. - Extended Key Usage can be either empty or contain Server Authentication. - No explicit limit to the length of the certificate chain. VMCA uses the OpenSSL default, which is 10 certificates. - Certificates with wildcards or with more than one DNS name are not supported. - You cannot create subsidiary CAs of VMCA. - See the VMware knowledge base article at http://kb.vmware.com/kb/2112009, Creating a Microsoft Certificate Authority Template for SSL certificate creation in vSphere 6.x, for an example using Microsoft Certificate Authority.

Certificate Type	Certificate Requirements
Machine SSL certificate	You can use vSphere Certificate Manager to create the CSR or create the CSR manually. If you create the CSR manually, it must meet the requirements listed previously under <i>Requirements for all Imported vSphere Certificates</i>. You also have to specify the FQDN for the host.
Solution user certificate	You can use vSphere Certificate Manager to create the CSR or create the CSR manually. Note: You must use a different value for Name for each solution user. If you generate the certificate manually, this might show up as CN under Subject, depending on the tool you use. If you use vSphere Certificate Manager, the tool prompts you for certificate information for each solution user. vSphere Certificate Manager stores the information in <code>certool.cfg</code>. For the vpxd-extension solution user, you can either leave Extended Key Usage empty or use "TLS WWW client authentication".

Requirements When Using Custom Certificates

When you want to use custom certificates, the certificates must meet the following requirements.

Certificate Type	Certificate Requirements
Machine SSL certificate	The machine SSL certificate on each node must have a separate certificate from your third-party or enterprise CA. - You can generate the CSR using the vSphere Client or the vSphere Certificate Manager, or create the CSR manually. The CSR must meet the requirements listed previously under <i>Requirements for all Imported vSphere Certificates</i>. - For most fields, you can accept the default or provide site-specific values. The FQDN of the machine is required.
Solution user certificate	Each solution user on each node must have a separate certificate from your third-party or enterprise CA. - You can generate the CSRs using vSphere Certificate Manager or prepare the CSR yourself. The CSR must meet the requirements listed previously under <i>Requirements for all Imported vSphere Certificates</i>. - If you use vSphere Certificate Manager, the utility prompts you for certificate information for each solution user. vSphere Certificate Manager stores the information in <code>certool.cfg</code>. Note: You must use a different value for Name for each solution user. A manually generated certificate might show up as CN under Subject, depending on the tool you use. When later you replace solution user certificates with custom certificates, provide the complete signing certificate chain of the third-party CA. For the vpxd-extension solution user, you can either leave Extended Key Usage empty or use "TLS WWW client authentication".

vSphere Certificate Management

The work required for setting up or updating your vSphere certificate infrastructure depends on the requirements in your environment. You must consider whether you are performing a fresh install or an upgrade, and whether you are considering ESXi or vCenter Server.

Environments That Use VMware Certificate Authority Certificates

VMware Certificate Authority (VMCA) can handle all certificate management. VMCA provisions vCenter Server components and ESXi hosts with certificates that use VMCA as the root certificate authority. If you are upgrading to vSphere 6.0 or later from an earlier version of vSphere, all self-signed certificates are replaced with certificates that are signed by VMCA.

If you do not currently replace VMware certificates, your environment starts using VMCA-signed certificates instead of self-signed certificates.

Environments That Use Custom Certificates

If your company policy requires certificates that are signed by a third-party or enterprise certificate authority, or that require custom certificate information, you have several choices for a fresh installation.

- Have the VMCA root certificate signed by a third-party CA or enterprise CA. Replace the VMCA root certificate with that signed certificate. In this scenario, the VMCA certificate is an intermediate certificate. VMCA provisions vCenter Server components and ESXi hosts with certificates that include the full certificate chain.
- If your company policy does not allow intermediate certificates in the chain, you can replace certificates explicitly. You can use the vSphere Client, vSphere Certificate Manager utility, or perform manual certificate replacement using the certificate management CLIs.

When upgrading an environment that uses custom certificates, you can retain some of the certificates.

- ESXi hosts keep their custom certificates during upgrade. Make sure that the vCenter Server upgrade process adds all the relevant root certificates to the TRUSTED_ROOTS store in VMware Certificate Endpoint Store (VECS) on the vCenter Server.
After the upgrade to vSphere 6.0 or later, you can set the certificate mode to **Custom**. If the certificate mode is VMCA (the default), and you perform a certificate refresh from the vSphere Client, the VMCA-signed certificates replace the custom certificates.
- For an upgrade of a simple vCenter Server installation to an embedded deployment, vCenter Server retains custom certificates. After the upgrade, your environment works as before. The existing vCenter Server and vCenter Single Sign-On certificates are retained. The certificates are used as machine SSL certificates. In addition, VMCA assigns a VMCA-signed certificate to each solution user (collection of vCenter services). The solution user uses this certificate only to authenticate to vCenter Single Sign-On. VMware does not recommend replacing solution user certificates.

vSphere Certificate Interfaces

For vCenter Server, you can view and replace certificates with the following tools and interfaces.

Table 104: Interfaces for Managing vCenter Server Certificates

Interface	Use
vSphere Client	Perform common certificate tasks with a graphical user interface.
vSphere Automation API	See <i>VMware vSphere Automation SDKs Programming Guide</i> .
vSphere Certificate Manager utility	Perform common certificate replacement tasks from the command line of the vCenter Server installation.

Interface	Use
vSphere Certificate Management CLIs	Perform all certificate management tasks with <code>dir-cli</code> , <code>certool</code> , and <code>vecs-cli</code> .
<code>sso-config</code> utility	Perform STS certificate management from the command line of the vCenter Server installation.
PowerCLI 12.4 or later (also requires vSphere 7.0 or later)	Perform trusted certificate store management, manage vCenter Server Machine SSL certificates, and manage ESXi Machine SSL certificates.

For ESXi, you perform certificate management from the vSphere Client. VMCA provisions certificates and stores them locally on the ESXi host. VMCA does not store ESXi host certificates in VMDIR or in VECS. See the *vSphere Security* documentation.

Supported vCenter Server Certificates

For vCenter Server and related machines and services, the following certificates are supported:

- Certificates that are generated and signed by VMware Certificate Authority (VMCA).
- Custom certificates.
 - Enterprise certificates that are generated from your own internal PKI.
 - Third-party CA-signed certificates that are generated by an external PKI such as Verisign, GoDaddy, and so on.

Self-signed certificates that were created using OpenSSL in which no Root CA exists are not supported.

Replacing vSphere Certificates


You can perform different types of certificate replacement depending on your company policy and requirements for the system that you are configuring. You can perform certificate replacement from the vSphere Client, by using the vSphere Certificate Manager utility, or manually by using the CLIs included with your installation.

You can replace the default certificates. For vCenter Server components, you can use a set of command-line tools included in your installation. You have several options.

Note: If your vCenter Server is linked to NSX-T manager, and you replace vCenter Server certificates, you must update the thumbprint of the vCenter Server compute manager. See the topic titled "Add a Compute Manager" in *NSX-T Data Center Migration Coordinator Guide*.


Replacing Certificates with VMCA-Signed Certificates

If your VMCA certificate expires or you want to replace it for other reasons, you can use the certificate management CLIs to perform that process. By default, the VMCA root certificate expires after 10 years, and all certificates that VMCA signs expire when the root certificate expires, that is, after a maximum of 10 years.

Figure 53: Certificates Signed by VMCA Are Stored in VECS**Making VMCA an Intermediate Certificate Authority**

You can replace the VMCA root certificate with a certificate that is signed by an enterprise certificate authority (CA), or a third-party CA. VMCA signs the custom root certificate each time it provisions certificates, making VMCA an intermediate CA.


Note: If you perform a fresh install with a vCenter Server, replace the VMCA root certificate before you add ESXi hosts. If you do, VMCA signs the whole chain, and you do not have to generate new certificates.

Figure 54: Certificates Signed by a Third-Party or Enterprise CA Use VMCA as an Intermediate CA

Replacing VMCA-Signed Certificates with Custom Certificates

You can replace the existing VMCA-signed certificates with custom certificates. If you use that approach, you are responsible for all certificate provisioning and monitoring.

Figure 55: External Certificates Are Stored Directly in VECS



Using the Hybrid Approach to Certificate Deployment

In the hybrid approach, you can have VMCA supply some of the certificates, but use custom certificates for other parts of your infrastructure. For example, because solution user certificates are used only to authenticate to vCenter Single Sign-On, consider having VMCA provision those certificates. Replace the machine SSL certificates with custom certificates to secure all SSL traffic.

Company policy often does not allow intermediate CAs. For those cases, hybrid deployment is a good solution. It minimizes the number of certificates to replace, and secures all traffic. The hybrid deployment leaves only internal traffic, that is, solution user traffic, to use the default VMCA-signed certificates.

ESXi Certificate Replacement

For ESXi hosts, you can change certificate provisioning behavior from the vSphere Client. See the *vSphere Security* documentation for details.

Table 105: ESXi Certificate Replacement Options

Option	Description
VMware Certificate Authority mode (default)	When you renew certificates from the vSphere Client, VMCA issues the certificates for the hosts. If you changed the VMCA root certificate to include a certificate chain, the host certificates include the full chain.
Custom Certificate Authority mode	Allows you to update and use certificates manually that are not signed or issued by VMCA.

Option	Description
Thumbprint mode	Can be used to retain 5.5 certificates during refresh. Use this mode only temporarily in debugging situations.

Where vSphere Uses Certificates

The VMware Certificate Authority (VMCA) provisions your environment with certificates. Certificates include machine SSL certificates for secure connections, solution user certificates for authentication of services to vCenter Single Sign-On, and certificates for ESXi hosts.

The following certificates are in use.

Table 106: Certificates in vSphere

Certificate	Provisioned	Comments
ESXi certificates	VMCA (default)	Stored locally on ESXi host.
Machine SSL certificates	VMCA (default)	Stored in VMware Endpoint Certificate Store (VECS).
Solution user certificates	VMCA (default)	Stored in VECS.
vCenter Single Sign-On SSL signing certificate	Provisioned during installation.	Manage this certificate from the command line. Note: Do not change this certificate in the filesystem or unpredictable behavior results.
VMware Directory Service (VMDIR) SSL certificate	Provisioned during installation.	In vSphere 6.5 and later, the machine SSL certificate is used as the vmdir certificate.
SMS self-signed certificates	Provisioned during registration of the IOFilter Provider.	In vSphere 7.0 and later, SMS self-signed certificates are stored in <code>/etc/vmware/ssl/iofilter/vp_castore.pem</code> . Before vSphere 7.0, SMS self-signed certificates are stored in <code>/etc/vmware/ssl/castore.pem</code> . In addition, the SMS Store can also store VVOL VASA Provider's (version 4.0 and earlier) self-signed certificates when <code>retainVasaProviderCertificate=True</code> .

ESXi Certificates

ESXi certificates are stored locally on each host in the `/etc/vmware/ssl` directory. ESXi certificates are provisioned by VMCA by default, but you can use custom certificates instead. ESXi certificates are provisioned when the host is first added to vCenter Server and when the host reconnects. For more information, see the *vSphere Security* documentation.

Machine SSL Certificates

The machine SSL certificate for each node is used to create an SSL socket on the server side. SSL clients connect to the SSL socket. The certificate is used for server verification and for secure communication such as HTTPS or LDAPS.

Each vCenter Server node has its own machine SSL certificate. All services that are running on a vCenter Server node use the machine SSL certificate to expose their SSL endpoints.

The following services use the machine SSL certificate.

- The reverse proxy service. SSL connections to individual vCenter services always go to the reverse proxy. Traffic does not go to the services themselves.

- The vCenter Server service (vpxd).
- The VMware Directory Service (vmdir).

VMware products use standard X.509 version 3 (X.509v3) certificates to encrypt session information. Session information is sent over SSL between components.

Solution User Certificates

A solution user encapsulates one or more vCenter Server services. Each solution user must be authenticated to vCenter Single Sign-On. Solution users use certificates to authenticate to vCenter Single Sign-On through SAML token exchange.

A solution user presents the certificate to vCenter Single Sign-On when it first has to authenticate, after a reboot, and after a timeout has elapsed. The timeout (Holder-of-Key Timeout) can be set from the vSphere Client and defaults to 2592000 seconds (30 days).

For example, the vpxd solution user presents its certificate to vCenter Single Sign-On when it connects to vCenter Single Sign-On. The vpxd solution user receives a SAML token from vCenter Single Sign-On and can then use that token to authenticate to other solution users and services.

The following solution user certificate stores are included in VECS:

- `machine` : Used by the license server and the logging service.
Note: The machine solution user certificate has nothing to do with the machine SSL certificate. The machine solution user certificate is used for the SAML token exchange. The machine SSL certificate is used for secure SSL connections for a machine.
- `vpxd` : vCenter service daemon (vpxd) store. vpxd uses the solution user certificate that is stored in this store to authenticate to vCenter Single Sign-On.
- `vpxd-extension` : vCenter extensions store. Includes the Auto Deploy service, inventory service, and other services that are not part of other solution users.
- `vsphere-webclient` : vSphere Client store. Also includes some additional services such as the performance chart service.
- `wcp` : VMware vSphere® with VMware Tanzu™ store. Also used for vSphere Cluster Services.

Internal Certificates

vCenter Single Sign-On certificates are not stored in VECS and are not managed with certificate management tools. As a rule, changes are not necessary, but in special situations, you can replace these certificates.

vCenter Single Sign-On Signing Certificate

The vCenter Single Sign-On service includes an identity provider service that issues SAML tokens that are used for authentication throughout vSphere. A SAML token represents the user's identity, and also contains group membership information. When vCenter Single Sign-On issues SAML tokens, it signs each token with its signing certificate so that clients of vCenter Single Sign-On can verify that the SAML token comes from a trusted source.

You can replace this certificate from the CLI. See [Replace a vCenter Server STS Certificate Using the Command Line](#).

VMware Directory Service SSL Certificate

In vSphere 6.5 and later, the machine SSL certificate is used as the VMware directory certificate. For earlier versions of vSphere, see the corresponding documentation.

vSphere Virtual Machine Encryption Certificates

The vSphere Virtual Machine Encryption solution connects with a key server. Depending on how the solution authenticates to the key server, it might generate certificates and store them in VECS. See the *vSphere Security* documentation.

VMware Certificate Authority and VMware Core Identity Services

Core identity services are part of every vCenter Server system. VMware Certificate Authority (VMCA) is part of every VMware core identity services group. Use the management CLIs and the vSphere Client to interact with these services.

VMware core identity services include several components.

Table 107: Core Identity Services

Service	Description
VMware Directory Service (vmdir)	Identity source that handles SAML certificate management for authentication with vCenter Single Sign-On.
VMware Certificate Authority (VMCA)	Issues certificates for VMware solution users, machine certificates for machines on which services are running, and ESXi host certificates. VMCA can be used as is, or as an intermediary certificate authority. VMCA issues certificates only to clients that can authenticate to vCenter Single Sign-On in the same domain.
VMware Authentication Framework Daemon (VMAFD)	Includes the VMware Endpoint Certificate Store (VECS) and several other authentication services. VMware administrators interact with VECS. The other services are used internally.

VMware Endpoint Certificate Store

VMware Endpoint Certificate Store (VECS) serves as a local (client-side) repository for certificates, private keys, and other certificate information that can be stored in a keystore. You can decide not to use VMCA as your certificate authority and certificate signer, but you must use VECS to store all vCenter certificates, keys, and so on. ESXi certificates are stored locally on each host and not in VECS.

VECS runs as part of the VMware Authentication Framework Daemon (VMAFD). VECS runs on every vCenter Server node, and holds the keystores that contain the certificates and keys.

VECS polls VMware Directory Service (vmdir) periodically for updates to the trusted root store. You can also explicitly manage certificates and keys in VECS using `vecs-cli` commands. See [vecs-cli Command Reference](#).

VECS includes the following stores.

Table 108: Stores in VECS

Store	Description
Machine SSL store (MACHINE_SSL_CERT)	- Used by the reverse proxy service on every vSphere node. - Used by the VMware Directory Service (vmdir) on each vCenter Server node. All services in vSphere 6.0 and later communicate through a reverse proxy, which uses the machine SSL certificate. For backward compatibility, the 5.x services still use specific ports. As a result, some services such as <code>vpzd</code> still have their own port open.
Solution user stores <code>machine</code> <code>vpzd</code> <code>vpzd-extension</code> <code>vsphere-webclient</code> <code>wcp</code>	VECS includes one store for each solution user. The subject of each solution user certificate must be unique, for example, the machine certificate cannot have the same subject as the <code>vpzd</code> certificate.

Store	Description
	Solution user certificates are used for authentication with vCenter Single Sign-On. vCenter Single Sign-On checks that the certificate is valid, but does not check other certificate attributes. The following solution user certificate stores are included in VECS: <code>machine</code> : Used by the license server and the logging service. Note: The machine solution user certificate has nothing to do with the machine SSL certificate. The machine solution user certificate is used for the SAML token exchange. The machine SSL certificate is used for secure SSL connections for a machine. <code>vpzd</code> : vCenter service daemon (vpzd) store. vpzd uses the solution user certificate that is stored in this store to authenticate to vCenter Single Sign-On. <code>vpzd-extension</code> : vCenter extensions store. Includes the Auto Deploy service, inventory service, and other services that are not part of other solution users. <code>vsphere-webclient</code> : vSphere Client store. Also includes some additional services such as the performance chart service. <code>wcp</code> : VMware vSphere® with VMware Tanzu™ store. Also used for vSphere Cluster Services. Each vCenter Server node includes a <code>machine</code> certificate.
Trusted root store (TRUSTED_ROOTS)	Contains all trusted root certificates.
vSphere Certificate Manager Utility backup store (BACKUP_STORE)	Used by VMCA (VMware Certificate Manager) to support certificate revert. Only the most recent state is stored as a backup, you cannot go back more than one step.
Other stores	Other stores might be added by solutions. For example, the Virtual Volumes solution adds an SMS store. Do not modify the certificates in those stores unless VMware documentation or a VMware Knowledge Base article instructs you to do so. Note: Deleting the TRUSTED_ROOTS_CRLS store can damage your certificate infrastructure. Do not delete or modify the TRUSTED_ROOTS_CRLS store.

The vCenter Single Sign-On service stores the token signing certificate and its SSL certificate on disk. You can change the token signing certificate from the CLI.

Some certificates are stored on the file system, either temporarily during startup or permanently. Do not change the certificates on the file system.

Note: Do not change any certificate files on disk unless instructed by VMware documentation or Knowledge Base Articles. Unpredictable behavior might result otherwise.

Managing vSphere Certificate Revocation

If you suspect that one of your certificates has been compromised, replace all existing certificates, including the VMCA root certificate.

vSphere supports replacing certificates but does not enforce certificate revocation for ESXi hosts or for vCenter Server systems.

Remove revoked certificates from all nodes. If you do not remove revoked certificates, a man-in-the-middle attack might enable compromise through impersonation with the account's credentials.

Replacing vSphere Certificates in Large Deployments

When replacing certificates in deployments with large numbers of vCenter Server hosts, you can use the vSphere Certificate Management utility or replace certificates manually using the CLIs. Some best practices guide the process you choose.

Replacement of Machine SSL Certificates in Environments with Multiple vCenter Server Systems

If your environment includes multiple vCenter Server systems, you can replace machine SSL certificates with the vSphere Client or the vSphere Certificate Manager utility, or manually with CLI commands.

Using the vSphere Certificate Manager to Replace Machine SSL Certificates on Multiple vCenter Server Systems

You run vSphere Certificate Manager on each machine. Depending on the task you perform, you are also prompted for certificate information. See the following topics for details:

- [Replace VMCA Root Certificate with Custom Signing Certificate and Replace All Certificates Using the Certificate Manager](#)
- [Replace Machine SSL Certificate with VMCA Certificate \(Intermediate CA\) Using the Certificate Manager](#)
- [Replace Solution User Certificates with VMCA Certificates \(Intermediate CA\) Using the Certificate Manager](#)

Using the CLI to Manually Replace Machine SSL Certificates on Multiple vCenter Server Systems

For manual certificate replacement, you run the certificate replacement CLI commands on each machine. See the following topics for details:

- [Replace Machine SSL Certificates with VMCA-Signed Certificates Using the CLI](#)
- [Replace Machine SSL Certificates \(Intermediate CA\) Using the CLI](#)
- [Replace Machine SSL Certificates with Custom Certificates Using the CLI](#)

Replacement of Solution User Certificates in Environments with Multiple vCenter Server Systems in Enhanced Linked Mode

If your environment includes multiple vCenter Server systems in enhanced linked mode, follow these steps for replacing solution user certificates.

Note: When you list solution user certificates in large deployments, the output of `/usr/lib/vmware-vmafd/bin/dir-cli list` includes all solution users from all nodes. Run `/usr/lib/vmware-vmafd/bin/vmafd-cli get-machine-id --server-name localhost` to find the local machine ID for each host. Each solution user name includes the machine ID.

Using the vSphere Certificate Manager to Replace Machine SSL Certificates on vCenter Server Systems in ELM

You run vSphere Certificate Manager on each machine. Depending on the task you perform, you are also prompted for certificate information. See [Managing Certificates Using the vSphere Certificate Manager Utility](#).

Using the CLI to Manually Replace Machine SSL Certificates on vCenter Server Systems in ELM

The high-level steps to manually replace machine SSL certificates on vCenter Server in ELM include:

1. Generating or requesting a certificate.
You need the following certificates:

- A certificate for the machine solution user on each vCenter Server.
 - A certificate for each of the following solution users on each node:
 - `vpxd` solution user
 - `vpxd-extension` solution user
 - `vsphere-webclient` solution user
 - `wcp` solution user
2. Using the CLI commands to replace the certificates on each node.
- The precise process depends on the type of certificate replacement that you are performing. See the following topics for details:
- [Replace Solution User Certificates with New VMCA-Signed Certificates Using the CLI](#)
 - [Replace Solution User Certificates \(Intermediate CA\) Using the CLI](#)
 - [Replace Solution User Certificates with Custom Certificates Using the Certificate Manager](#)

Certificate Replacement in VMware Environments That Include External Solutions

Some solutions, such as VMware vCenter Site Recovery Manager or VMware vSphere Replication, are always installed on a different machine than the vCenter Server system. If you replace the default machine SSL certificate on the vCenter Server system, a connection error results if the solution attempts to connect to the vCenter Server system.

You can run the `ls_update_certs` script to resolve the issue. See the VMware knowledge base article at <https://kb.vmware.com/s/article/2109074>.

Managing Certificates Using the vSphere Client

You can view and manage certificates by using the vSphere Client.

The vSphere Client enables you to perform these management tasks.

- View the machine SSL, VMware Certificate Authority (VMCA) root, Trusted Root, and Security Token Service (STS) certificates.
- Add new Trusted Root certificates, and renew or replace existing machine SSL and STS certificates.
- Generate a custom Certificate Signing Request (CSR) for a machine SSL certificate and replace the certificate when the Certificate Authority returns it.

Most parts of the certificate replacement workflows are supported fully from the vSphere Client. Other certificate replacement workflows are supported by the vSphere Certificate Manager utility. See [Managing Certificates Using the vSphere Certificate Manager Utility](#).

To understand more about options for replacing the default certificates, see [Replacing vSphere Certificates](#).

Note:

If you use the VMCA as an intermediate CA, or use custom certificates, you might encounter significant complexity and the potential for a negative impact to your security, and an unnecessary increase in your operational risk.

Explore Certificate Stores Using the vSphere Client

A VMware Endpoint Certificate Store (VECS) instance is included on each vCenter Server node. You can explore the different stores inside the VMware Endpoint Certificate Store from the vSphere Client, including machine SSL, STS, and trusted root certificates.

For most management tasks, you must have the password for the administrator for the local domain account, `administrator@vsphere.local` or a different domain if you changed the domain during installation.

See [VMware Endpoint Certificate Store](#) for details on the different stores inside VECS.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for `administrator@vsphere.local` or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as `administrator@mydomain`.
3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Explore the certificates stored inside the VMware Endpoint Certificate Store (VECS).
[VMware Endpoint Certificate Store](#) explains what is in the individual stores.
6. To view details for a certificate, select the appropriate certificate tab, select the certificate, and expand the certificate to view details.

Set the Threshold for vCenter Certificate Expiration Warnings Using the vSphere Client

vCenter Server monitors all certificates in the VMware Endpoint Certificate Store (VECS) and issues an alarm when a certificate is 30 days or less from its expiration. You can use the vSphere Client to change how soon you are warned with the `vpdx.cert.threshold` advanced option.

1. Log in to the vSphere Client.
2. Select the vCenter Server object and click **Configure**.
3. Click **Advanced Settings**.
4. Click **Edit Settings** and filter for `threshold`.
5. Change the setting of `vpdx.cert.threshold` to the desired value and click **Save**.

Renew VMCA Certificates with New VMCA-Signed Certificates Using the vSphere Client

You can replace all VMCA-signed certificates with new VMCA-signed certificates. This process is called renewing certificates. You can renew selected certificates or all certificates in your environment from the vSphere Client.

For certificate management, you have to supply the password of the administrator of the local domain (`administrator@vsphere.local` by default). If you are renewing certificates for a vCenter Server system, you also have to supply the vCenter Single Sign-On credentials for a user with administrator privileges on the vCenter Server system.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for `administrator@vsphere.local` or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as `administrator@mydomain`.

3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Renew the VMCA-signed machine SSL certificate for the local system.
 - a) From the **Machine SSL** tab, select the desired certificate and click **Renew**.
 - b) Specify the duration of the certificate in days.
 - c) Click the checkbox to acknowledge that you have backed up vCenter Server and its databases.
 - d) Click **Renew**.

The system renews the certificate and displays a success message.
 - e) When the certificate has been changed message appears, click **Refresh** to refresh your browser.

Replace Certificates with Custom Certificates Using the vSphere Client

You can use the vSphere Client to replace the default certificates with custom certificates.

You can use the vSphere Client to generate CSRs for each machine, and replace certificates when you receive them from your internal or third-party Certificate Authority (CA). When you submit the CSRs to your internal or third-party CA, the CA returns signed certificates and the root certificate. You can upload both the root certificate and the signed certificates from the vSphere Client.

Generate Certificate Signing Request for Machine SSL Certificate Using the vSphere Client (Custom Certificates)

The machine SSL certificate is used by the reverse proxy service on every vCenter Server node. Each machine must have a machine SSL certificate for secure communication with other services. You can use the vSphere Client to generate a Certificate Signing Request (CSR) for the machine SSL certificate and to replace the certificate once it is ready.

The certificate must meet the following requirements:

- Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded).
- CRT format
- x509 version 3
- SubjectAltName must contain DNS Name=<machine_FQDN>.
- Contains the following Key Usages: Digital Signature, Key Encipherment

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 bits and 3072 bits.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as administrator@mydomain.

3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. Enter the credentials of your vCenter Server.
5. Generate the CSR.
 - a) Under the **Machine SSL** tab, select the desired certificate and click **Generate Certificate Signing Request (CSR)**.
 - b) Enter your certificate information and click **Next**.
2048 (bits) is the default value for the key size. Change this value as required.
Note: When you use vCenter Server to generate a CSR with a large key size, the generation takes a few minutes to complete because of the CPU-intensive nature of the operation.
 - c) Copy or download the CSR.
 - d) Click **Finish**.
 - e) Provide the CSR to your Certificate Authority.

When the Certificate Authority returns the certificate, replace the existing certificate in the certificate store. See [Add Custom Certificates Using the vSphere Client](#).

Add a Trusted Root Certificate to the Certificate Store Using the vSphere Client

If you want to use third-party certificates in your environment, you must add a trusted root certificate to the certificate store. You can do so using the vSphere Client.

Obtain the custom root certificate from your third-party or in-house certificate authority (CA).

vSphere accepts only valid CA certificates for import. To be valid, a CA certificate must have the CA bit and the keyCertSign bit set in the basic constraint and the key usage X.509 v3 certificate extensions respectively. This implies that the certificate is a CA and its purpose is for certificate signing. See <https://www.rfc-editor.org/rfc/rfc5280> for more information.

Ensure that the keyCertSign bit is set for all the certificates in the chain.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Under the **Trusted Root** tab, click **Add Trusted Root Certificate**.
6. Click **Browse** and select the location of the certificate chain.
You can use a file of type CER, PEM, or CRT.
7. Click **Add**.
The certificate is added to the store.

Note: In vSphere 8.0 Update 2 and later, the **Start Root certificate push to vCenter Hosts** check box is removed. vCenter Server pushes the root certificates to all connected hosts in the inventory when a certificate is added. When a host with different root certificates from vCenter Server is connected, vCenter Server pushes the root certificates to

correct this difference. In this case, vCenter Server root certificates overwrite the ones on the host, so administrators can ensure that any custom root certificates needed throughout the inventory are added to vCenter Server.

Add Custom Certificates Using the vSphere Client

You can use the vSphere Client to add custom Machine SSL certificates to the certificate store.

Generate certificate signing requests (CSRs) for each certificate that you want to replace. See [Generate Certificate Signing Request for Machine SSL Certificate Using the vSphere Client \(Custom Certificates\)](#). Place the certificate and private key in a location that the vCenter Server can access.

Usually, replacing the machine SSL certificate for each component is sufficient.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Under the **Machine SSL** tab, select the certificate then click **Import and Replace Certificate**.
6. Click the appropriate certificate replacement option and click **Next**.

Option	Description
Replace with VMCA certificate	Creates a VMCA-generated CSR to replace the current certificate.
Replace with external CA certificate where CSR is generated from vCenter Server (private key embedded)	Use a certificate signed using a vCenter Server generated CSR to replace the current certificate.
Replace with external CA certificate (requires private key)	Use a certificate signed by an external CA to replace the current certificate.

7. Enter the CSR information, or upload the appropriate certificates.
8. Click the checkbox to acknowledge that you have backed up vCenter Server and its databases.
9. Review the information and click **Finish**.
The system replaces the certificate and displays a success message.
10. When the certificate has been changed message appears, click **Refresh** to refresh your browser.

Generate a VMCA Leaf Certificate

You can generate a leaf certificate that is signed by the VMware Certificate Authority (VMCA) for use in your VMware infrastructure.

Generate a Certificate Signing Request (CSR) on the host in your VMware infrastructure where you want to install the leaf certificate.

In addition to VMware Certificate Authority (VMCA) handling all certificate management, it can generate leaf certificates. Leaf certificates are signed by VMCA and are used to identify other VMware resources. VMCA-generated leaf certificates are not stored in VECS. Also, vCenter Server does not track these leaf certificates for expiration.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Under the **Trusted Root** tab, select the VMCA root certificate and click **Issue New Leaf Certificate**.
6. Browse for the CSR that you previously generated, specify a duration, then click **Next**.
7. Click **Download Certificates** to save the Leaf and Root certificates.

The generated Leaf and Root certificates are created and downloaded to the specified location.

Import the Leaf and Root certificates to the target host in your VMware infrastructure.

Managing Certificates Using the vSphere Certificate Manager Utility

The vSphere Certificate Manager utility allows you to perform most certificate management tasks interactively from the command line. vSphere Certificate Manager prompts you for the task to perform, for certificate locations and other information as needed, and then stops and starts services and replaces certificates for you.

To understand more about options for replacing the default certificates, see [Replacing vSphere Certificates](#).

Note:

If you use the VMCA as an intermediate CA, or use custom certificates, you might encounter significant complexity and the potential for a negative impact to your security, and an unnecessary increase in your operational risk.

If you use vSphere Certificate Manager, you are not responsible for placing the certificates in VECS (VMware Endpoint Certificate Store) and you are not responsible for starting and stopping services.

You run vSphere Certificate Manager options in sequence to complete a workflow. Several options, for example, generating CSRs, are used in different workflows. Before you run vSphere Certificate Manager, be sure that you understand the replacement process and procure the certificates that you want to use.

CAUTION: vSphere Certificate Manager supports one level of revert. If you run vSphere Certificate Manager twice and notice that you unintentionally corrupted your environment, the tool cannot revert the first of the two runs.

vSphere Certificate Manager Utility Location

The vSphere Certificate Manager utility is located at:

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

Note:

When running vSphere Certificate Manager, some options prompt you as follows:

Enter proper value for VMCA 'Name':

Respond to this prompt by entering the fully qualified domain name of the machine on which the certificate configuration is running.

Workflows in the vSphere Certificate Manager Utility

The following table presents an overview of the certificate replacement workflows you can accomplish by using the vSphere Certificate Manager utility.

Table 109: Workflows in the vSphere Certificate Management Utility

Workflow	Description	See
Replacing VMCA root certificate with custom signing certificate and replacing all certificates	To generate the VMCA root certificate, and replace all certificates, use Option 4, Regenerate a new VMCA Root certificate and replace all Certificates.	Regenerate a New VMCA Root Certificate and Replace All Certificates Using the Certificate Manager
Making VMCA an intermediate certificate authority	To make VMCA an intermediate CA, you must run the vSphere Certificate Manager utility several times and use multiple options. This workflow gives the complete set of steps for replacing both machine SSL certificates and solution user certificates.	Make VMCA an Intermediate Certificate Authority Using the Certificate Manager
Replacing all certificates with custom certificates	To replace all certificates with custom certificates, you must run the vSphere Certificate Manager utility several times and use multiple options. This workflow gives the complete set of steps for replacing both machine SSL certificates and solution user certificates.	Replace All Certificates with a Custom Certificate Using the Certificate Manager
Reverting the last performed operation	To revert the last performed certificate operation and return to the previous state, use Option 7, Revert last performed operation by re-publishing old certificates.	Revert Last Performed Operation by Republishing Old Certificates Using the Certificate Manager
Resetting all certificates	To replace all existing vCenter certificates with certificates that are signed by VMCA, use Option 8, Reset all Certificates.	Reset All Certificates Using the Certificate Manager

Regenerate a New VMCA Root Certificate and Replace All Certificates Using the Certificate Manager

You can use the vSphere Certificate Manager utility to regenerate the VMCA root certificate, and replace the local machine SSL certificate and the local solution user certificates with VMCA-signed certificates. When multiple vCenter Server instances are connected in Enhanced Linked Mode configuration, you must replace certificates on each vCenter Server.

You must know the following information when you run vSphere Certificate Manager with this option.

- Password for administrator@vsphere.local.

- The FQDN of the machine for which you want to generate a new VMCA-signed certificate. All other properties default to the predefined values but can be changed.

When you replace the existing machine SSL certificate with a new VMCA-signed certificate, vSphere Certificate Manager prompts you for information and enters all values, except for the password and the IP address of the vCenter Server, into the `certool.cfg` file.

- Password for administrator@vsphere.local
- Two-letter country code
- Company name
- Organization name
- Organization unit
- State
- Locality
- IP address (optional)
- Email
- Host name, that is, the fully qualified domain name of the machine for which you want to replace the certificate. If the host name does not match the FQDN, certificate replacement does not complete correctly and your environment might end up in an unstable state.
- IP address of vCenter Server
- VMCA name, that is, the fully qualified domain name of the machine on which the certificate configuration is running.

Note: The OU (organizationalUnitName) field is no longer mandatory.

1. Log in to the vCenter Server shell and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select option 4, Regenerate a new VMCA Root Certificate and replace all certificates.

3. Enter the administrator user and password.

4. Respond to the prompts.

vSphere Certificate Manager generates a new VMCA root certificate based on your input and replaces all certificates on the system where you are running vSphere Certificate Manager. The replacement process is complete after vSphere Certificate Manager has restarted the services.

5. To replace the machine SSL certificate, run vSphere Certificate Manager with option 3, Replace Machine SSL certificate with VMCA Certificate.
6. To replace the solution user certificates, run Certificate Manager with option 6, Replace Solution user certificates with VMCA certificates.

Make VMCA an Intermediate Certificate Authority Using the Certificate Manager

You can use the vSphere Certificate Manager utility to make VMCA an Intermediate CA. After you complete the process, VMCA signs all new certificates with the full chain. If you want, you can use vSphere Certificate Manager to replace all existing certificates with new VMCA-signed certificates.

To make VMCA an intermediate CA, you must run vSphere Certificate Manager several times. The high-level steps for replacing both machine SSL certificates and solution user certificates include:

1. Launching the vSphere Certificate Manager utility.
2. Generating a CSR by running Option 2, Replace VMCA Root certificate with Custom Signing Certificate and replace all Certificates. You might have to provide some information about the certificate next. When prompted for an option again, select Option 1, Generate Certificate Signing Request(s) and Key(s) for VMCA Root Signing certificate.

3. Submitting the CSR to your external or enterprise CA. You receive a signed certificate and a root certificate from the CA.
4. Combining the VMCA root certificate with the CA root certificate and saving the file.
5. Replacing certificates by running Option 2, Replace VMCA Root certificate with Custom Signing Certificate and replace all Certificates, and following the prompts. This process replaces all certificates on the local machine.
6. (Optional) Replacing certificates on each node when multiple vCenter Server instances are connected in Enhanced Linked Mode configuration by:
 - a. First replacing the machine SSL certificate with the (new) VMCA certificate (Option 3, Replace Machine SSL certificate with VMCA Certificate).
 - b. Then replacing the solution user certificates with the (new) VMCA certificate (Option 6, Replace Solution user certificates with VMCA certificates).

Generate CSR Using the Certificate Manager and Prepare Root Certificate (Intermediate CA)

You can use the vSphere Certificate Manager utility to generate Certificate Signing Requests (CSRs). Submit those CSRs to your enterprise CA or to an external certificate authority for signing. You can use the signed certificates with the different supported certificate replacement processes.

vSphere Certificate Manager prompts you for information. The prompts depend on your environment and on the type of certificate that you want to replace.

For any CSR generation, you are prompted for the password of the administrator@vsphere.local user, or for the administrator of the vCenter Single Sign-On domain that you are connecting to.

- You can use vSphere Certificate Manager to create the CSR.

Note: In vSphere 8.0 and later, if you use the vSphere Certificate Manager to generate the CSR, the minimum key size is changed to 3072 bits from 2048. In vSphere 8.0 Update 1 and later, use the vSphere Client to generate a CSR with a key size of 2048 bits.

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 bits and 3072 bits.
- If you prefer to create the CSR manually, the certificate that you send to be signed must meet the following requirements.
 - Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
 - PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When keys are added to VECS, they are converted to PKCS8.
 - x509 version 3
 - The CA extension must be set to true for root certificates, and cert sign must be in the list of requirements. For example:


```
basicConstraints      = critical,CA:true
keyUsage              = critical,digitalSignature,keyCertSign
```
 - CRL signing must be enabled.
 - Extended Key Usage can be either empty or contain Server Authentication.
 - No explicit limit to the length of the certificate chain. VMCA uses the OpenSSL default, which is 10 certificates.
 - Certificates with wildcards or with more than one DNS name are not supported.
 - You cannot create subsidiary CAs of VMCA.

See the VMware knowledge base article at <http://kb.vmware.com/kb/2112009>, Creating a Microsoft Certificate Authority Template for SSL certificate creation in vSphere 6.x, for an example using Microsoft Certificate Authority.

1. Log in to the vCenter Server shell and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select Option 2, Replace VMCA Root certificate with Custom Signing Certificate and replace all Certificates. Initially, you use this option to generate the CSR, not to replace certificates.

3. Enter the administrator user and password.
4. Select Option 1, Generate Certificate Signing Request(s) and Key(s) for VMCA Root Signing certificate, to generate the CSR and answer the prompts.
As part of the process, you have to provide a directory. vSphere Certificate Manager places the certificate to be signed (*.csr file) and the corresponding key file (*.key file) in the directory.
5. Name the certificate signing request (CSR) `root_signing_cert.csr`.
6. Send the CSR to your enterprise or external CA for signing and name the resulting signed certificate `root_signing_cert.cer`.
7. In a text editor, combine the certificates as follows.

```
-----BEGIN CERTIFICATE-----
Signed VMCA root certificate
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
CA intermediate certificates
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
Root certificate of enterprise or external CA
-----END CERTIFICATE-----
```

8. Save the file as `root_signing_chain.cer`.

Replace the existing root certificate with the chained root certificate. See [Replace VMCA Root Certificate with Custom Signing Certificate and Replace All Certificates Using the Certificate Manager](#).

Replace VMCA Root Certificate with Custom Signing Certificate and Replace All Certificates Using the Certificate Manager

You can use the vSphere Certificate Manager utility to generate a CSR and send the CSR to an enterprise or third-party CA for signing. You can then replace the VMCA root certificate with a custom signing certificate and replace all existing certificates with certificates that are signed by the custom CA.

- Generate the certificate chain.
 - You can use vSphere Certificate Manager to create the CSR or create the CSR manually.
 - After you receive the signed certificate from your third-party or enterprise CA, combine it with the initial VMCA root certificate to create the full chain.
See [Generate CSR Using the Certificate Manager and Prepare Root Certificate \(Intermediate CA\)](#) for certificate requirements and the process of combining the certificates.
- Gather the information that you need.
 - Password for `administrator@vsphere.local`
 - Valid custom certificate for Root (.crt file)
 - Valid custom key for Root (.key file)

You run vSphere Certificate Manager on vCenter Server to replace the VMCA root certificate with a custom signing certificate.

1. Log in to the vCenter Server shell and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select Option 2, Replace VMCA Root certificate with Custom Signing Certificate and replace all Certificates.
3. Enter the administrator user and password.
4. Select Option 2, Import custom certificate(s) and key(s) to replace existing VMCA Root Signing certificate, and respond to the prompts.
 - a) Specify the full path to the root certificate when prompted.
 - b) If you are replacing certificates for the first time, you are prompted for information to be used for the machine SSL certificate.

This information includes the required FQDN of the machine and is stored in the `certtool.cfg` file.

Replace Machine SSL Certificate with VMCA Certificate (Intermediate CA) Using the Certificate Manager

When you use VMCA as an intermediate CA, you can replace the machine SSL certificate explicitly using the vSphere Certificate Manager utility. First you replace the VMCA root certificate on the vCenter Server, then you can replace the machine SSL certificate, which will be signed by the new root of the VMCA. You can also use this option to replace machine SSL certificates that are corrupt or about to expire.

- You must know the following information to run vSphere Certificate Manager with this option.
 - Password for administrator@vsphere.local.
 - The FQDN of the machine for which you want to generate a new VMCA-signed certificate. All other properties default to the predefined values but can be changed.
 - Host name or IP address of the vCenter Server system.

When you replace the existing machine SSL certificate with a new VMCA-signed certificate, vSphere Certificate Manager prompts you for information and enters all values, except for the password and the IP address of the vCenter Server, into the `certtool.cfg` file.

- Password for administrator@vsphere.local
- Two-letter country code
- Company name
- Organization name
- Organization unit
- State
- Locality
- IP address (optional)
- Email
- Host name, that is, the fully qualified domain name of the machine for which you want to replace the certificate. If the host name does not match the FQDN, certificate replacement does not complete correctly and your environment might end up in an unstable state.
- IP address of vCenter Server
- VMCA name, that is, the fully qualified domain name of the machine on which the certificate configuration is running.

Note: The OU (organizationalUnitName) field is no longer mandatory.

1. Log in to the vCenter Server shell and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select Option 3, Replace Machine SSL certificate with VMCA Certificate.
 3. Enter the administrator user and password.
 4. Respond to the prompts.
- vSphere Certificate Manager stores the information in the `certtool.cfg` file.

vSphere Certificate Manager replaces the machine SSL certificate.

Replace Solution User Certificates with VMCA Certificates (Intermediate CA) Using the Certificate Manager

When you use VMCA as an intermediate CA, you can replace the solution user certificate explicitly using the vSphere Certificate Manager utility. First you replace the VMCA root certificate on the vCenter Server, then you can replace the solution user certificate, which will be signed by the new root of the VMCA. You can also use this option to replace solution certificates that are corrupt or about to expire.

- Restart all vCenter Server nodes explicitly if you replaced the VMCA root certificate in a deployment consisting of multiple instances of vCenter Server in Enhanced Linked Mode configuration.
- You must know the following information to run vSphere Certificate Manager with this option.
 - Password for administrator@vsphere.local
 - Host name or IP address of the vCenter Server system

1. Log in to the vCenter Server shell and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select option 6, Replace Solution user certificates with VMCA certificates.

3. Enter the administrator user and password.

4. Respond to the prompts.

See the VMware knowledge base article at <https://kb.vmware.com/s/article/2112281> for more information.

vSphere Certificate Manager replaces all solution user certificates.

Replace All Certificates with a Custom Certificate Using the Certificate Manager

You can use the vSphere Certificate Manager utility to replace all certificates with custom certificates. Before you start the process, you must send CSRs to your certificate authority (CA). You can use Certificate Manager to generate the CSRs.

One option is to replace only the machine SSL certificate, and to use the solution user certificates that are provisioned by VMCA. Solution user certificates are used only for communication between vSphere components.

When you use custom certificates, you replace the VMCA-signed certificates with custom certificates. You can use the vSphere Client, the vSphere Certificate Manager utility, or CLIs for manual certificate replacement. Certificates are stored in VECS.

To replace all certificates with custom certificates, you must run the vSphere Certificate Manager utility several times. The high-level steps for replacing both machine SSL certificates and solution user certificates include:

1. Launching the vSphere Certificate Manager utility.
2. Generating certificate signing requests for the machine SSL certificate and the solution user certificates separately on each machine.
 - a. To generate CSRs for the machine SSL certificate, select Option 1, Replace Machine SSL certificate with Custom Certificate. When prompted for an option again, select Option 1, Generate Certificate Signing Request(s) and Key(s) for Machine SSL certificate.
 - b. If company policy does not allow a hybrid deployment, select Option 5, Replace Solution user certificates with Custom Certificate.
3. Submitting the CSR to your external or enterprise CA. You receive a signed certificate and a root certificate from the CA.
4. After receiving the signed certificates and the root certificate from your CA, replacing the machine SSL certificate on each machine by using Option 1, Replace Machine SSL certificate with Custom Certificate.

5. If you also want to replace the solution user certificates, select Option 5, Replace Solution user certificates with Custom Certificate.
6. Finally, when multiple vCenter Server instances are connected in Enhanced Linked Mode configuration, repeating the process on each node.

Generate Certificate Signing Requests Using the Certificate Manager (Custom Certificates)

You can use the vSphere Certificate Manager utility to generate Certificate Signing Requests (CSRs) that you can then use with your enterprise CA or send to an external certificate authority. You can use the certificates with the different supported certificate replacement processes.

vSphere Certificate Manager prompts you for information. The prompts depend on your environment and on the type of certificate you want to replace.

- For any CSR generation, you are prompted for the password of the administrator@vsphere.local user, or for the administrator of the vCenter Single Sign-On domain that you are connecting to.
- You are prompted for the host name or IP address of the vCenter Server.
- To generate a CSR for a machine SSL certificate, you are prompted for certificate properties, which are stored in the `certtool.cfg` file. For most fields, you can accept the default or provide site-specific values. The FQDN of the machine is required.

Note: In vSphere 8.0 and later, if you use the vSphere Certificate Manager to generate the CSR, the minimum key size is changed to 3072 bits from 2048. In vSphere 8.0 Update 1 and later, use the vSphere Client to generate a CSR with a key size of 2048 bits.

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 bits and 3072 bits.

1. Log in to each vCenter Server (the vCenter Server shell) in your environment and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Selection Option 1, Replace Machine SSL certificate with Custom Certificate.
3. Enter the administrator user and password.
4. Select Option 1, Generate Certificate Signing Request(s) and Key(s) for Machine SSL certificate, to generate the CSR, answer the prompts and exit vSphere Certificate Manager.
As part of the process, you have to provide a directory. vSphere Certificate Manager places the certificate and key files in the directory.
5. If you also want to replace all solution user certificates, restart vSphere Certificate Manager and selection Option 5, Replace Solution user certificates with Custom Certificate.
6. Supply the password and the vCenter Server IP address or host name if prompted.
7. Select Option 1, Generate Certificate Signing Request(s) and Key(s) for Solution User Certificates, to generate the CSRs, answer the prompts and exit vSphere Certificate Manager.
As part of the process, you have to provide a directory. Certificate Manager places the certificate and key files in the directory.

To perform certificate replacement, see [Replace Machine SSL Certificate with Custom Certificate Using the Certificate Manager](#).

Replace Machine SSL Certificate with Custom Certificate Using the Certificate Manager

You can use the vSphere Certificate Manager utility to replace the machine SSL certificate on each node with a custom certificate. The machine SSL certificate is used by the reverse proxy service on every vCenter Server node. Each machine must have a machine SSL certificate for secure communication with other services.

Before you start, you need a CSR for each machine in your environment. You can generate the CSR using vSphere Certificate Manager or explicitly.

1. To generate the CSR using vSphere Certificate Manager, see [Generate Certificate Signing Requests Using the Certificate Manager \(Custom Certificates\)](#).
2. To generate the CSR explicitly, request a certificate for each machine from your third-party or enterprise CA. The certificate must meet the following requirements:
 - Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
 - CRT format
 - x509 version 3
 - SubjectAltName must contain DNS Name=<machine_FQDN>.
 - Contains the following Key Usages: Digital Signature, Key Encipherment

See also the VMware knowledge base article at <https://kb.vmware.com/s/article/2112014>, Obtaining vSphere certificates from a Microsoft Certificate Authority.

1. Log in to vCenter Server and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Selection Option 1, Replace Machine SSL certificate with Custom Certificate.
3. Enter the administrator user and password.
4. Select Option 2, Import custom certificate(s) and key(s) to replace existing Machine SSL certificate, to start certificate replacement and respond to the prompts.

vSphere Certificate Manager prompts you for the following information:

- Password for administrator@vsphere.local
- Valid Machine SSL custom certificate (.crt file)
- Valid Machine SSL custom key (.key file)
- Valid signing certificate for the custom machine SSL certificate (.crt file)
- IP address of the vCenter Server

Replace Solution User Certificates with Custom Certificates Using the Certificate Manager

Many companies only require that you replace certificates of services that are accessible externally. However, the vSphere Certificate Manager utility also supports replacing solution user certificates. Solution users are collections of services, for example, all services that are associated with the vSphere Client.

Before you start, you need a CSR for each machine in your environment. You can generate the CSR using vSphere Certificate Manager or explicitly.

1. To generate the CSR using vSphere Certificate Manager, see [Generate Certificate Signing Requests Using the Certificate Manager \(Custom Certificates\)](#).
2. Request a certificate for each solution user on each node from your third-party or enterprise CA. You can generate the CSR using vSphere Certificate Manager or prepare it yourself. The CSR must meet the following requirements:
 - Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
 - CRT format
 - x509 version 3
 - SubjectAltName must contain DNS Name=<machine_FQDN>.
 - Each solution user certificate must have a different Subject. Consider, for example, including the solution user name (such as vpxd) or other unique identifier.
 - Contains the following Key Usages: Digital Signature, Key Encipherment

See also the VMware knowledge base article at <http://kb.vmware.com/kb/2112014>, Obtaining vSphere certificates from a Microsoft Certificate Authority.

When you are prompted for a solution user certificate, provide the complete signing certificate chain of the third-party CA.

The format looks similar to the following.

```
-----BEGIN CERTIFICATE-----
Signing certificate
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
CA intermediate certificates
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
Root certificate of enterprise or external CA
-----END CERTIFICATE-----
```

1. Log in to vCenter Server and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select Option 5, Replace Solution user certificates with Custom Certificate.
3. Enter the SSO user and password.
4. Select Option 2, Import custom certificate(s) and key(s) to replace existing Solution User Certificates, and respond to the prompts.

vSphere Certificate Manager prompts you for the following information:

- Password for administrator@vsphere.local
- Certificate and key for machine solution user
- The certificate and key (vpzd.crt and vpzd.key) for the machine solution user
- The full set of certificates and keys (vpzd.crt and vpzd.key) for all solution users

Revert Last Performed Operation by Republishing Old Certificates Using the Certificate Manager

When you perform a certificate management operation by using the vSphere Certificate Manager utility, the current certificate state is stored in the BACKUP_STORE store in VECS before certificates are replaced. You can revert the last performed operation and return to the previous state.

Note: The revert operation restores what is currently in the BACKUP_STORE. If you run vSphere Certificate Manager with two different options and you then attempt to revert, only the last operation is reverted.

1. Log in to the vCenter Server shell and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select Option 7, Revert last performed operation by re-publishing old certificates.
3. Enter the administrator user and password.
4. To continue, enter Y.

Reset All Certificates Using the Certificate Manager

You can use the vSphere Certificate Manager utility to replace all existing vCenter certificates with certificates that are signed by VMCA.

When you use this option, you overwrite all custom certificates that are currently in VMware Endpoint Certificate Store (VECS).

vSphere Certificate Manager can replace all certificates. Which certificates are replaced depends on which options you select.

1. Log in to vCenter Server shell and start the vSphere Certificate Manager.

```
/usr/lib/vmware-vmca/bin/certificate-manager
```

2. Select Option 8, Reset all Certificates.
3. Enter the administrator user and password.
4. When prompted, enter your certificate information.

After certificates are replaced and services are restarted, verify your certificate information.

Manual vSphere Certificate Replacement

For some special certificate replacement cases, you cannot use the vSphere Certificate Manager utility. Instead, you can use the CLIs included with your installation for certificate replacement.

Guidelines for Stopping and Starting of vCenter Server Services

For certain parts of manual certificate replacement, you must stop all vCenter Server services and then start only the services that manage the certificate infrastructure. If you stop services only when needed, you can minimize downtime.

You have to stop and start services as part of the certificate replacement process. You can use the `service-control` command for starting and stopping services. You can start and stop all services or individual services. See the command-line help for more information.

Follow these guidelines.

- Do not stop services to generate new public/private key pairs or new certificates.
- If you are the only administrator, you do not have to stop services when you add a new root certificate. The old root certificate remains available, and all services can still authenticate with that certificate.
- Stop services right before you delete a machine SSL certificate in VMware Endpoint Certificate Store (VECS).

Replace Existing VMCA-Signed Certificates with New VMCA-Signed Certificates Using the CLI

If the VMware Certificate Authority (VMCA) root certificate expires in the near future, or if you want to replace it for other reasons, you can use the CLI to generate a new root certificate and add it to the VMware Directory Service. You can then generate new machine SSL certificates and solution user certificates using the new root certificate.

Only `administrator@vsphere.local` or other users in the CAAdmins group can perform certificate management tasks. See [Add Members to a vCenter Single Sign-On Group](#).

Use the vSphere Certificate Manager utility to replace certificates for most cases.

If you need fine-grained control, this scenario gives detailed step-by-step instructions for replacing the complete set of certificates using CLI commands. You can instead replace only individual certificates using the procedure in the corresponding task.

Generate a New VMCA-Signed Root Certificate Using the CLI

You can generate new VMCA-signed certificates with the `certool` CLI and publish the certificates to `vmdir`.

1. On the vCenter Server, generate a new self-signed certificate and private key.

```
certool --genselfcacert --outprivkey <key_file_path> --outcert <cert_file_path> --config <config_file>
```

2. Replace the existing root certificate with the new certificate.

```
certool --rootca --cert <cert_file_path> --privkey <key_file_path>
```

The command generates the certificate, adds it to vmdir, and adds it to VECS.

3. Stop all services and start the services that handle certificate creation, propagation, and storage.

```
service-control --stop --all
service-control --start vmafd
service-control --start vmdird
service-control --start vmcad
```

4. Optional: Publish the new root certificate to vmdir.

```
dir-cli trustedcert publish --cert newRoot.crt
```

The command updates all instances of vmdir immediately. If you do not run the command, propagation of the new certificate to all nodes might take a while.

5. Restart all services.

```
service-control --start --all
```

Generate a New VMCA-Signed Root Certificate

The following example shows all the steps for verifying the current root CA information, and for regenerating the root certificate.

1. (Optional) On the vCenter Server, list the VMCA root certificate to make sure it is in the certificate store.

```
/usr/lib/vmware-vmca/bin/certool --getrootca
```

The output looks similar to this:

```
output:
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number:
      cf:2d:ff:49:88:50:e5:af
  ...
```

2. (Optional) List the VECS TRUSTED_ROOTS store and compare the certificate serial number there with the output from Step 1.

```
/usr/lib/vmware-vmca/bin/vecs-cli entry list --store TRUSTED_ROOTS --text
```

In the simplest case with only one root certificate, the output looks like this:

```
Number of entries in store :    1
Alias : 960d43f31eb95211ba3a2487ac840645a02894bd
Entry type :    Trusted Cert
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number:
      cf:2d:ff:49:88:50:e5:af
```

3. Generate a new VMCA root certificate. The command adds the certificate to the TRUSTED_ROOTS store in VECS and in vmdir (VMware Directory Service).

```
/usr/lib/vmware-vmca/bin/certool --selfca --config=/usr/lib/vmware-vmca/share/config/cer-
tool.cfg
```

Replace Machine SSL Certificates with VMCA-Signed Certificates Using the CLI

After you generate a new VMCA-signed root certificate, you can use the `vecs-cli` command to replace all machine SSL certificates in your environment.

Be prepared to stop all services and to start the services that handle certificate propagation and storage.

Each machine must have a machine SSL certificate for secure communication with other services. When multiple vCenter Server instances are connected in Enhanced Linked Mode configuration, you must run the Machine SSL certificate generation commands on each node.

1. Make one copy of `certtool.cfg` for each machine that needs a new certificate.

You can find the `certtool.cfg` file in the `/usr/lib/vmware-vmca/share/config/` directory.

2. Edit the custom configuration file for each machine to include that machine's FQDN.

Run `NSLookup` against the machine's IP address to see the DNS listing of the name, and use that name for the `Hostname` field in the file.

3. Generate a public/private key file pair and a certificate for each file, passing in the configuration file that you just customized.

For example:

```
certtool --genkey --privkey=machine1.priv --pubkey=machine1.pub
certtool --gencert --privkey=machine1.priv --cert machine1.crt --Name=Machine1_Cert --config machine1.cfg
```

4. Stop all services and start the services that handle certificate creation, propagation, and storage.

```
service-control --stop --all
service-control --start vmafd
service-control --start vmdird
service-control --start vmcad
```

5. Add the new certificate to VECS.

All machines need the new certificate in the local certificate store to communicate over SSL. You first delete the existing entry, then add the new entry.

```
vecs-cli entry delete --store MACHINE_SSL_CERT --alias __MACHINE_CERT
vecs-cli entry create --store MACHINE_SSL_CERT --alias __MACHINE_CERT --cert machine1.crt --key machine1.priv
```

6. Restart all services.

```
service-control --start --all
```

Replacing Machine Certificates with VMCA-Signed Certificates

1. Create a configuration file for the SSL certificate and save it as `ssl-config.cfg` in the current directory.

```
Country = US
Name = vmca-<FQDN-example>
Organization = <my_company>
OrgUnit = <my_company Engineering>
State = <my_state>
Locality = <mytown>
Hostname = <FQDN>
```

2. Generate a key pair for the machine SSL certificate. In a deployment of multiple vCenter Server instances connected in Enhanced Linked Mode configuration, run this command on each vCenter Server node.

```
/usr/lib/vmware-vmca/bin/certtool --genkey --privkey=ssl-key.priv --pubkey=ssl-key.pub
```

The `ssl-key.priv` and `ssl-key.pub` files are created in the current directory.

3. Generate the new machine SSL certificate. This certificate is signed by VMCA. If you replaced the VMCA root certificate with custom certificate, VMCA signs all certificates with the full chain.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vmca-ssl.crt --privkey=ssl-key.priv --
config=ssl-config.cfg
```

The new-vmca-ssl.crt file is created in the current directory.

4. (Optional) List the content of VECS.

```
/usr/lib/vmware-vmafd/bin/vecs-cli store list
```

- Sample output on vCenter Server:

```
output (on vCenter):
MACHINE_SSL_CERT
TRUSTED_ROOTS
TRUSTED_ROOT_CRLS
machine
vsphere-webclient
vpxd
vpxd-extension
hvc
data-encipherment
APPLMGMT_PASSWORD
SMS
wcp
KMS_ENCRYPTION
```

5. Replace the Machine SSL certificate in VECS with the new Machine SSL certificate. The --store and --alias values have to exactly match with the default names.

- On each vCenter Server, run the following commands to update the Machine SSL certificate in the MACHINE_SSL_CERT store. You must update the certificate for each machine separately because each has a different FQDN.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store MACHINE_SSL_CERT --alias
__MACHINE_CERT
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store MACHINE_SSL_CERT --alias
__MACHINE_CERT --cert new-vmca-ssl.crt --key ssl-key.priv
```

You can also replace the certificates for your ESXi hosts. See the *vSphere Security* publication.

Replace Solution User Certificates with New VMCA-Signed Certificates Using the CLI

After you replace the machine SSL certificates, you can use the `dir-cli` command to replace all solution user certificates. Solution user certificates must be valid, that is, not expired, but none of the other information in the certificate is used by the certificate infrastructure.

Be prepared to stop all services and to start the services that handle certificate propagation and storage.

Many VMware customers do not replace solution user certificates. They replace only the machine SSL certificates with custom certificates. This hybrid approach satisfies the requirements of their security teams.

- Certificates either sit behind a proxy, or they are custom certificates.
- No intermediate CAs are used.

You replace the machine solution user certificate and the solution user certificate on each vCenter Server system.

Note: When you list solution user certificates in large deployments, the output of `/usr/lib/vmware-vmafd/bin/dir-cli list` includes all solution users from all nodes. Run `/usr/lib/vmware-vmafd/bin/vmafd-cli get-machine-id --server-name localhost` to find the local machine ID for each host. Each solution user name includes the machine ID.

1. Make one copy of `certool.cfg`, remove the Name, IP address, DNS name, and email fields, and rename the file, for example, to `sol_usr.cfg`.

You can name the certificates from the command line as part of generation. The other information is not needed for solution users. If you leave the default information, the certificates that are generated are potentially confusing.

2. Generate a public/private key file pair and a certificate for each solution user, passing in the configuration file that you just customized.

For example:

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=vpzd.priv --pubkey=vpzd.pub
/usr/lib/vmware-vmca/bin/certool --gencert --privkey=vpzd.priv --cert vpzd.crt --Name=VPXD_1 --config
sol_usr.cfg
```

3. Find the name for each solution user.

```
/usr/lib/vmware-vmafd/bin/dir-cli service list
```

You can use the unique ID that is returned when you replace the certificates. The input and output might look as follows.

```
/usr/lib/vmware-vmafd/bin/dir-cli service list
Enter password for administrator@vsphere.local:
1. machine-623bef28-0311-436e-b21f-6e0d39aa5179
2. vsphere-webclient-623bef28-0311-436e-b21f-6e0d39aa5179
3. vpzd-623bef28-0311-436e-b21f-6e0d39aa5179
4. vpzd-extension-623bef28-0311-436e-b21f-6e0d39aa5179
5. hvc-623bef28-0311-436e-b21f-6e0d39aa5179
6. wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e
```

In a deployment of multiple vCenter Server instances connected in Enhanced Linked Mode configuration, the output of `/usr/lib/vmware-vmafd/bin/dir-cli service list` includes all solution users from all nodes. Run `/usr/lib/vmware-vmafd/bin/vmafd-cli get-machine-id --server-name localhost` to find the local machine ID for each host. Each solution user name includes the machine ID.

4. Stop all services and start the services that handle certificate creation, propagation, and storage.

```
service-control --stop --all
service-control --start vmafdd
service-control --start vmdir
service-control --start vmcad
```

5. For each solution user, replace the existing certificate in vmdir and then in VECS.

The following example shows how to replace the certificates for the vpzd service.

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name <vpzd-xxxx-xxx-7c7b769cd9f4> --cert ./vpzd.crt
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store vpzd --alias vpzd
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store vpzd --alias vpzd --cert vpzd.crt --key vpzd.priv
```

Note: Solution users cannot authenticate to vCenter Single Sign-On if you do not replace the certificate in vmdir.

6. Restart all services.

```
service-control --start --all
```

Using VMCA-Signed Solution User Certificates

1. Generate a public/private key pair for each solution user on each vCenter Server node in an Enhanced Linked Mode configuration. That includes a pair for the machine solution and a pair for each additional solution user (vpxd, vpxd-extension, vsphere-webclient, wcp).
 - a. Generate a key pair for the machine solution user.


```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=machine-key.priv --pubkey=ma-  
chine-key.pub
```
 - b. Generate a key pair for the vpxd solution user on each node.


```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=vpxd-key.priv --pubkey=vpxd-key.pub
```
 - c. Generate a key pair for the vpxd-extension solution user on each node.


```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=vpxd-extension-key.priv --pub-  
key=vpxd-extension-key.pub
```
 - d. Generate a key pair for the vsphere-webclient solution user on each node.


```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=vsphere-webclient-key.priv --pubkey=v-  
sphere-webclient-key.pub
```
 - e. Generate a key pair for the wcp solution user on each node.


```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=wcp-key.priv --pubkey=wcp-key.pub
```
2. Generate solution user certificates that are signed by the new VMCA root certificate for the machine solution user and for each additional solution user (vpxd, vpxd-extension, vsphere-webclient, wcp) on each vCenter Server node.

Note: The `--Name` parameter has to be unique. Including the name of the solution user store name makes it easy to see which certificate maps to which solution user. The example includes the name, for example vpxd or vpxd-extension in each case.

- a. Make one copy of the `/usr/lib/vmware-vmca/share/config/certool.cfg` file, then modify or remove the Name, IP address, DNS name, and email fields as required, and rename the file, for example, to `sol_usr.cfg`.
 - b. Generate a certificate for the machine solution user on each node.


```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-machine.crt --privkey=machine-key.priv  
--Name=machine --config sol_usr.cfg
```
 - c. Generate a certificate for the vpxd solution user on each node.


```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vpxd.crt --privkey=vpxd-key.priv --  
Name=vpxd --config sol_usr.cfg
```
 - d. Generate a certificate for the vpxd-extensions solution user on each node.


```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vpxd-extension.crt --privkey=vpxd-ex-  
tension-key.priv --Name=vpxd-extension --config sol_usr.cfg
```
 - e. Generate a certificate for the vsphere-webclient solution user on each node by running the following command.


```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vsphere-webclient.crt --privkey=v-  
sphere-webclient-key.priv --Name=vsphere-webclient --config sol_usr.cfg
```
 - f. Generate a certificate for the wcp solution user on each node by running the following command.


```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-wcp.crt --privkey=wcp-key.priv --  
Name=wcp --config sol_usr.cfg
```
3. Replace the solution user certificates in VECS with the new solution user certificates.

Note: The `--store` and `--alias` parameters have to exactly match the default names for services.

 - a. Replace the machine solution user certificate on each node:


```
/usr/lib/vmware-vmca/bin/vecs-cli entry delete --store machine --alias machine
```



```
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store machine --alias machine --cert new-machine.crt --key machine-key.priv
```

b. Replace the vpxd solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store vpxd --alias vpxd
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store vpxd --alias vpxd --cert new-vpxd.crt --key vpxd-key.priv
```

c. Replace the vpxd-extension solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store vpxd-extension --alias vpxd-extension
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store vpxd-extension --alias vpxd-extension --cert new-vpxd-extension.crt --key vpxd-extension-key.priv
```

d. Replace the vsphere-webclient solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store vsphere-webclient --alias vsphere-webclient
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store vsphere-webclient --alias vsphere-webclient --cert new-vsphere-webclient.crt --key vsphere-webclient-key.priv
```

e. Replace the wcp solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store wcp --alias wcp
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store wcp --alias wcp --cert new-wcp.crt --key wcp-key.priv
```

4. Update VMware Directory Service (vmdir) with the new solution user certificates. You are prompted for a vCenter Single Sign-On administrator password.

a. Run `/usr/lib/vmware-vmafd/bin/dir-cli service list` to get the unique service ID suffix for each solution user. You run this command on a vCenter Server system.

```
/usr/lib/vmware-vmafd/bin/dir-cli service list
output:
1. machine-623bef28-0311-436e-b21f-6e0d39aa5179
2. vsphere-webclient-623bef28-0311-436e-b21f-6e0d39aa5179
3. vpxd-623bef28-0311-436e-b21f-6e0d39aa5179
4. vpxd-extension-623bef28-0311-436e-b21f-6e0d39aa5179
5. hvc-623bef28-0311-436e-b21f-6e0d39aa5179
6. wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e
```

Note: When you list solution user certificates in large deployments, the output of `/usr/lib/vmware-vmafd/bin/dir-cli list` includes all solution users from all nodes. Run `/usr/lib/vmware-vmafd/`

`bin/vmafd-cli get-machine-id --server-name localhost` to find the local machine ID for each host. Each solution user name includes the machine ID.

- b. Replace the machine certificate in `vmdir` on each vCenter Server node. For example, if `machine-6fd7f140-60a9-11e4-9e28-005056895a69` is the machine solution user on the vCenter Server, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name machine-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-machine.crt
```

- c. Replace the `vpzd` solution user certificate in `vmdir` on each node. For example, if `vpzd-6fd7f140-60a9-11e4-9e28-005056895a69` is the `vpzd` solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name vpzd-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-vpzd.crt
```

- d. Replace the `vpzd-extension` solution user certificate in `vmdir` on each node. For example, if `vpzd-extension-6fd7f140-60a9-11e4-9e28-005056895a69` is the `vpzd-extension` solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name vpzd-extension-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-vpzd-extension.crt
```

- e. Replace the `vsphere-webclient` solution user certificate on each node. For example, if `vsphere-webclient-6fd7f140-60a9-11e4-9e28-005056895a69` is the `vsphere-webclient` solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name vsphere-webclient-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-vsphere-webclient.crt
```

- f. Replace the `wcp` solution user certificate on each node. For example, if `wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e` is the `wcp` solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e --cert new-wcp.crt
```

Restart all services on each vCenter Server node.

Make VMCA Into an Intermediate Certificate Authority Using the CLI

You can use the CLI to replace the VMCA root certificate with a third-party CA-signed certificate that includes VMCA in the certificate chain. Going forward, all certificates that VMCA generates include the full chain. You can replace existing certificates with newly generated certificates.

If you use VMCA as an intermediate CA, or use custom certificates, you might encounter significant complexity and the potential for a negative impact to your security, and an unnecessary increase in your operational risk.

Replace the Root Certificate (Intermediate CA) Using the CLI

The first step in replacing the VMCA certificates with custom certificates is generating a CSR, and sending the CSR to be signed. You then use the CLI to add the signed certificate to VMCA as a root certificate.

You can use the Certificate Manager utility or other tool to generate the CSR. The CSR must meet the following requirements:

- Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
- PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When keys are added to VECS, they are converted to PKCS8.
- x509 version 3
- The CA extension must be set to true for root certificates, and `cert sign` must be in the list of requirements. For example:

```
basicConstraints      = critical,CA:true
keyUsage              = critical,digitalSignature,keyCertSign
```

- CRL signing must be enabled.
- Extended Key Usage can be either empty or contain Server Authentication.
- No explicit limit to the length of the certificate chain. VMCA uses the OpenSSL default, which is 10 certificates.
- Certificates with wildcards or with more than one DNS name are not supported.
- You cannot create subsidiary CAs of VMCA.

See the VMware knowledge base article at <http://kb.vmware.com/kb/2112009>, Creating a Microsoft Certificate Authority Template for SSL certificate creation in vSphere 6.x, for an example using Microsoft Certificate Authority.

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 bits and 3072 bits.

VMCA validates the following certificate attributes when you replace the root certificate:

- Key size: 2048 bits (minimum) to 8192 bits (maximum).
- Key Usage: Cert Sign
- Basic Constraint: Subject Type CA

1. Generate a CSR and send it to your CA.

Follow your CA's instructions.

2. Prepare a certificate file that includes the signed VMCA certificate and the full CA chain of your third-party CA or enterprise CA. Save the file, for example as `rootcal.crt`.

You can accomplish this step by copying all CA certificates in PEM format into a single file. You start with the VMCA root certificate and end up with the root CA PEM certificate. For example:

```
-----BEGIN CERTIFICATE-----
<Certificate of VMCA>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Certificate of intermediary CA>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Certificate of Root CA>
-----END CERTIFICATE-----
```

3. Stop all services and start the services that handle certificate creation, propagation, and storage.

```
service-control --stop --all
service-control --start vmafd
service-control --start vmdir
service-control --start vmcad
```

4. Replace the existing VMCA root CA.

```
certtool --rootca --cert=rootcal.crt --privkey=root1.key
```

When you run this command, it:

- Adds the new custom root certificate to the certificate location in the file system.
- Appends the custom root certificate to the TRUSTED_ROOTS store in VECS (after a delay).
- Adds the custom root certificate to vmdir (after a delay).

- Optional: To propagate the change to all instances of vmdir (VMware Directory Service), publish the new root certificate to vmdir, supplying the full file path for each file.

For example, if the certificate has only one certificate in the chain:

```
dir-cli trustedcert publish --cert rootcal.crt
```

If the certificate has more than one certificate in the chain:

```
dir-cli trustedcert publish --cert rootcal.crt --chain
```

Replication between vmdir nodes happens every 30 seconds. You do not have to add the root certificate to VECS explicitly because VECS polls vmdir for new root certificate files every 5 minutes.

- Optional: If necessary, you can force a refresh of VECS.

```
vecs-cli force-refresh
```

- Restart all services.

```
service-control --start --all
```

Replacing the Root Certificate

Replace the VMCA root certificate with the custom CA root certificate using the `certtool` command with the `--rootca` option.

```
/usr/lib/vmware-vmca/bin/certtool --rootca --cert=<path>/root.pem --privkey=<path>/root.key
```

When you run this command, it:

- Adds the new custom root certificate to the certificate location in the file system.
- Appends the custom root certificate to the TRUSTED_ROOTS store in VECS.
- Adds the custom root certificate to vmdir.

You can remove the original VMCA root certificate from the certificate store if your company policy requires it. If you do, you have to replace the vCenter Single Sign-On Signing certificate. See [Replace a vCenter Server STS Certificate Using the Command Line](#).

Replace Machine SSL Certificates (Intermediate CA) Using the CLI

After you have received the signed certificate from the CA, you can use the CLI to make it the VMCA root certificate and replace all machine SSL certificates.

For each machine SSL certificate, the `SubjectAltName` must contain `DNS Name=<Machine FQDN>`.

These steps are essentially the same as the steps for replacing with a certificate that uses VMCA as the certificate authority. However, in this case, VMCA signs all certificates with the full chain.

Each machine must have a machine SSL certificate for secure communication with other services. When multiple vCenter Server instances are connected in Enhanced Linked Mode configuration, you must run the Machine SSL certificate generation commands on each node.

- Make one copy of `certtool.cfg` for each machine that needs a new certificate.

The `certtool.cfg` file is located in the `/usr/lib/vmware-vmca/share/config/` directory.

- Edit the custom configuration file for each machine to include that machine's FQDN.

Run `NSLookup` against the machine's IP address to see the DNS listing of the name, and use that name for the `Hostname` field in the file.

- Generate a public/private key file pair and a certificate for each machine, passing in the configuration file that you just customized.

For example:

```
certool --genkey --privkey=machine1.priv --pubkey=machine1.pub
```

```
certool --gencert --privkey=machine1.priv --cert machine42.crt --Name=Machine42_Cert --config machine1.cfg
```

4. Stop all services and start the services that handle certificate creation, propagation, and storage.

```
service-control --stop --all
service-control --start vmafd
service-control --start vmdird
service-control --start vmcad
```

5. Add the new certificate to VECS.

All machines need the new certificate in the local certificate store to communicate over SSL. You first delete the existing entry, then add the new entry.

```
vecs-cli entry delete --store MACHINE_SSL_CERT --alias __MACHINE_CERT
vecs-cli entry create --store MACHINE_SSL_CERT --alias __MACHINE_CERT --cert machine1.cert
--key machine1.priv
```

6. Restart all services.

```
service-control --start --all
```

Replacing Machine SSL Certificates (VMCA Is Intermediate CA)

1. Create a configuration file for the SSL certificate and save it as `ssl-config.cfg` in the current directory.

```
Country = US
Name = vmca-<FQDN-example>
Organization = VMware
OrgUnit = VMware Engineering
State = California
Locality = Palo Alto
Hostname = <FQDN>
```

2. Generate a key pair for the machine SSL certificate. In a deployment of multiple vCenter Server instances connected in Enhanced Linked Mode configuration, run this command on each vCenter Server node.

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=ssl-key.priv --pubkey=ssl-key.pub
```

The `ssl-key.priv` and `ssl-key.pub` files are created in the current directory.

3. Generate the new machine SSL certificate. This certificate is signed by VMCA. If you replaced the VMCA root certificate with custom certificate, VMCA signs all certificates with the full chain.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vmca-ssl.crt --privkey=ssl-key.priv --
config=ssl-config.cfg
```

The `new-vmca-ssl.crt` file is created in the current directory.

4. (Optional) List the content of VECS.

```
/usr/lib/vmware-vmca/bin/vecs-cli store list
```

– Sample output on vCenter Server:

```
output (on vCenter):
MACHINE_SSL_CERT
TRUSTED_ROOTS
TRUSTED_ROOT_CRLS
machine
vsphere-webclient
vpxd
vpxd-extension
hvc
data-encipherment
APPLMGMT_PASSWORD
SMS
```

```
wcp
KMS_ENCRYPTION
```

5. Replace the Machine SSL certificate in VECS with the new Machine SSL certificate. The `--store` and `--alias` values have to exactly match with the default names.
 - On each vCenter Server, run the following commands to update the Machine SSL certificate in the `MACHINE_SSL_CERT` store. You must update the certificate for each machine separately because each has a different FQDN.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store MACHINE_SSL_CERT --alias
__MACHINE_CERT
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store MACHINE_SSL_CERT --alias
__MACHINE_CERT --cert new-vmca-ssl.crt --key ssl-key.priv
```

Replace Solution User Certificates (Intermediate CA) Using the CLI

After you replace the machine SSL certificates, you can use the CLI to replace the solution user certificates.

Each solution user certificate must have a different `Subject`. Consider, for example, including the solution user name (such as `vpzd`) or other unique identifier.

Many VMware customers do not replace solution user certificates. They replace only the machine SSL certificates with custom certificates. This hybrid approach satisfies the requirements of their security teams.

- Certificates either sit behind a proxy, or they are custom certificates.
- No intermediate CAs are used.

You replace the machine solution user certificate and the solution user certificate on each vCenter Server system.

Note: When you list solution user certificates in large deployments, the output of `/usr/lib/vmware-vmafd/bin/dir-cli list` includes all solution users from all nodes. Run `/usr/lib/vmware-vmafd/bin/vmafd-cli get-machine-id --server-name localhost` to find the local machine ID for each host. Each solution user name includes the machine ID.

1. Make one copy of `certool.cfg`, remove the Name, IP address, DNS name, and email fields, and rename the file, for example, to `sol_usr.cfg`.

You can name the certificates from the command line as part of generation. The other information is not needed for solution users. If you leave the default information, the certificates that are generated are potentially confusing.

2. Generate a public/private key file pair and a certificate for each solution user, passing in the configuration file that you just customized.

For example:

```
certool --genkey --privkey=vpzd.priv --pubkey=vpzd.pub
certool --gencert --privkey=vpzd.priv --cert vpzd.crt --Name=VPXD_1 --config sol_usr.cfg
```

3. Find the name for each solution user.

```
/usr/lib/vmware-vmafd/bin/dir-cli service list
```

You can use the unique ID that is returned when you replace the certificates. The input and output might look as follows.

```
/usr/lib/vmware-vmafd/bin/dir-cli service list
Enter password for administrator@vsphere.local:
1. machine-623bef28-0311-436e-b21f-6e0d39aa5179
2. vsphere-webclient-623bef28-0311-436e-b21f-6e0d39aa5179
3. vpzd-623bef28-0311-436e-b21f-6e0d39aa5179
4. vpzd-extension-623bef28-0311-436e-b21f-6e0d39aa5179
5. hvc-623bef28-0311-436e-b21f-6e0d39aa5179
```

```
6. wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e
```

In a deployment of multiple vCenter Server instances connected in Enhanced Linked Mode configuration, the output of `/usr/lib/vmware-vmafd/bin/dir-cli service list` includes all solution users from all nodes. Run `/usr/lib/vmware-vmafd/bin/vmafd-cli get-machine-id --server-name localhost` to find the local machine ID for each host. Each solution user name includes the machine ID.

4. Stop all services and start the services that handle certificate creation, propagation, and storage.

```
service-control --stop --all
service-control --start vmafd
service-control --start vmdird
service-control --start vmcad
```

5. Replace the existing certificate in vmdir and then in VECS.

For solution users, you must add the certificates in that order. For example:

```
dir-cli service update --name <vpzd-xxxx-xxx-7c7b769cd9f4> --cert ./vpzd.crt
vecs-cli entry delete --store vpzd --alias vpzd
vecs-cli entry create --store vpzd --alias vpzd --cert vpzd.crt --key vpzd.priv
```

Note: Solution users cannot log in to vCenter Single Sign-On if you do not replace the certificate in vmdir.

6. Restart all services.

```
service-control --start --all
```

Replacing Solution User Certificates (Intermediate CA)

1. Generate a public/private key pair for each solution user on each vCenter Server node in an Enhanced Linked Mode configuration. That includes a pair for the machine solution and a pair for each additional solution user (vpzd, vpzd-extension, vsphere-webclient, wcp).

- a. Generate a key pair for the machine solution user.

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=machine-key.priv --pubkey=ma-
chine-key.pub
```

- b. Generate a key pair for the vpzd solution user on each node.

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=vpzd-key.priv --pubkey=vpzd-key.pub
```

- c. Generate a key pair for the vpzd-extension solution user on each node.

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=vpzd-extension-key.priv --pub-
key=vpzd-extension-key.pub
```

- d. Generate a key pair for the vsphere-webclient solution user on each node.

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=vsphere-webclient-key.priv --pubkey=v-
sphere-webclient-key.pub
```

- e. Generate a key pair for the wcp solution user on each node.

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=wcp-key.priv --pubkey=wcp-key.pub
```


2. Generate solution user certificates that are signed by the new VMCA root certificate for the machine solution user and for each additional solution user (vpxd, vpxd-extension, vsphere-webclient, wcp) on each vCenter Server node.

Note: The `--Name` parameter has to be unique. Including the name of the solution user store name makes it easy to see which certificate maps to which solution user. The example includes the name, for example vpxd or vpxd-extension in each case.

- a. Make one copy of the `/usr/lib/vmware-vmca/share/config/certool.cfg` file, then modify or remove the Name, IP address, DNS name, and email fields as required, and rename the file, for example, to `sol_usr.cfg`.
- b. Generate a certificate for the machine solution user on each node.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-machine.crt --privkey=machine-key.priv
--Name=machine --config sol_usr.cfg
```

- c. Generate a certificate for the vpxd solution user on each node.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vpxd.crt --privkey=vpxd-key.priv --
Name=vpxd --config sol_usr.cfg
```

- d. Generate a certificate for the vpxd-extensions solution user on each node.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vpxd-extension.crt --privkey=vpxd-ex-
tension-key.priv --Name=vpxd-extension --config sol_usr.cfg
```

- e. Generate a certificate for the vsphere-webclient solution user on each node by running the following command.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-vsphere-webclient.crt --privkey=v-
sphere-webclient-key.priv --Name=vsphere-webclient --config sol_usr.cfg
```

- f. Generate a certificate for the wcp solution user on each node by running the following command.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=new-wcp.crt --privkey=wcp-key.priv --
Name=wcp --config sol_usr.cfg
```

3. Replace the solution user certificates in VECS with the new solution user certificates.

Note: The `--store` and `--alias` parameters have to exactly match the default names for services.

- a. Replace the machine solution user certificate on each node:

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store machine --alias machine
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store machine --alias machine --cert new-
machine.crt --key machine-key.priv
```

- b. Replace the vpxd solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store vpxd --alias vpxd
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store vpxd --alias vpxd --cert new-
vpxd.crt --key vpxd-key.priv
```

- c. Replace the vpxd-extension solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store vpxd-extension --alias vpxd-exten-
sion
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store vpxd-extension --alias vpxd-exten-
sion --cert new-vpxd-extension.crt --key vpxd-extension-key.priv
```

- d. Replace the vsphere-webclient solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store vsphere-webclient --alias vsphere-
webclient
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store vsphere-webclient --alias vsphere-
webclient --cert new-vsphere-webclient.crt --key vsphere-webclient-key.priv
```

- e. Replace the wcp solution user certificate on each node.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store wcp --alias wcp
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store wcp --alias wcp --cert new-wcp.crt
--key wcp-key.priv
```

4. Update VMware Directory Service (vmdir) with the new solution user certificates. You are prompted for a vCenter Single Sign-On administrator password.

- a. Run `/usr/lib/vmware-vmafd/bin/dir-cli service list` to get the unique service ID suffix for each solution user. You run this command on a vCenter Server system.

```
/usr/lib/vmware-vmafd/bin/dir-cli service list
output:
1. machine-623bef28-0311-436e-b21f-6e0d39aa5179
2. vsphere-webclient-623bef28-0311-436e-b21f-6e0d39aa5179
3. vpxd-623bef28-0311-436e-b21f-6e0d39aa5179
4. vpxd-extension-623bef28-0311-436e-b21f-6e0d39aa5179
5. hvc-623bef28-0311-436e-b21f-6e0d39aa5179
6. wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e
```

Note: When you list solution user certificates in large deployments, the output of `/usr/lib/vmware-vmafd/bin/dir-cli list` includes all solution users from all nodes. Run `/usr/lib/vmware-vmafd/bin/vmafd-cli get-machine-id --server-name localhost` to find the local machine ID for each host. Each solution user name includes the machine ID.

- b. Replace the machine certificate in vmdir on each vCenter Server node. For example, if `machine-6fd7f140-60a9-11e4-9e28-005056895a69` is the machine solution user on the vCenter Server, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name machine-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-machine.crt
```

- c. Replace the vpxd solution user certificate in vmdir on each node. For example, if `vpxd-6fd7f140-60a9-11e4-9e28-005056895a69` is the vpxd solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name vpxd-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-vpxd.crt
```

- d. Replace the vpxd-extension solution user certificate in vmdir on each node. For example, if `vpxd-extension-6fd7f140-60a9-11e4-9e28-005056895a69` is the vpxd-extension solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name vpxd-extension-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-vpxd-extension.crt
```

- e. Replace the vsphere-webclient solution user certificate on each node. For example, if `vsphere-webclient-6fd7f140-60a9-11e4-9e28-005056895a69` is the vsphere-webclient solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name vsphere-webclient-6fd7f140-60a9-11e4-9e28-005056895a69 --cert new-vsphere-webclient.crt
```

- f. Replace the wcp solution user certificate on each node. For example, if `wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e` is the wcp solution user ID, run this command:

```
/usr/lib/vmware-vmafd/bin/dir-cli service update --name wcp-1cbe0a40-e4ce-4378-b5e7-9460e2b8200e --cert new-wcp.crt
```

Replace Certificates with Custom Certificates Using the CLI

If your company policy requires it, you can use the CLI to replace some or all certificates used in vSphere with certificates that are signed by a third-party or enterprise CA. If you do that, VMCA is not in your certificate chain. You are responsible for storing all vCenter certificates in VECS.

You can replace all certificates or use a hybrid solution. For example, consider replacing all certificates that are used for network traffic but leaving VMCA-signed solution user certificates. Solution user certificates are used only for authentication to vCenter Single Sign-On. vCenter Server uses solution user certificates for internal communication only. Solution user certificates are not used for external communication.

Note: If you do not want to use VMCA, you are responsible for replacing all certificates yourself, for provisioning new components with certificates, and for tracking certificate expiration.

Even if you decide to use custom certificates, you can still use the VMware Certificate Manager utility for certificate replacement. See [Replace All Certificates with a Custom Certificate Using the Certificate Manager](#).

If you encounter problems with vSphere Auto Deploy after replacing certificates, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2000988>.

Request Certificates and Import a Custom Root Certificate Using the CLI

You can use custom certificates from an enterprise or third-party CA. The first step is requesting the certificates from the certificate authority then using the CLI to import the root certificates into VMware Endpoint Certificate Store (VECS).

The certificate must meet the following requirements:

- Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
- PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When keys are added to VECS, they are converted to PKCS8.
- x509 version 3
- For root certificates, the CA extension must be set to true, and the cert sign must be in the list of requirements.
- SubjectAltName must contain DNS Name=<machine_FQDN>.
- CRT format
- Contains the following Key Usages: Digital Signature, Key Encipherment
- Start time of one day before the current time.
- CN (and SubjectAltName) set to the host name (or IP address) that the ESXi host has in the vCenter Server inventory.

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 bits and 3072 bits.

1. Send the Certificate Signing Requests (CSRs) for the following certificates to your enterprise or third-party certificate provider.
 - A machine SSL certificate for each machine. For the machine SSL certificate, the SubjectAltName field must contain the fully qualified domain name (DNS NAME=*machine_FQDN*).
 - Optionally, five solution user certificates for each node. Solution user certificates do not need to include IP address, host name, or email address. Each certificate must have a different certificate Subject.

Typically, the result is a PEM file for the trusted chain, plus the signed SSL certificates for each vCenter Server node.

2. List the TRUSTED_ROOTS and the machine SSL stores.

```
vecs-cli store list
```

- a) Ensure that the current root certificate and all machine SSL certificates are signed by VMCA.
- b) Note down the Serial number, issuer, and Subject CN fields.
- c) Optional: With a Web browser, open an HTTPS connection to a node where the certificate is to be replaced, view the certificate information, and ensure that it matches the machine SSL certificate.

3. Stop all services and start the services that handle certificate creation, propagation, and storage.

```
service-control --stop --all
service-control --start vmafd
service-control --start vmdir
service-control --start vmcad
```

4. Publish the custom root certificate.

```
dir-cli trustedcert publish --cert <my_custom_root>
```

If you do not specify a user name and password on the command line, you are prompted.

5. Restart all services.

```
service-control --start --all
```

You can remove the original VMCA root certificate from the certificate store if your company policy requires it. If you do, you have to refresh the vCenter Single Sign-On certificate. See [Replace a vCenter Server STS Certificate Using the Command Line](#).

Replace Machine SSL Certificates with Custom Certificates Using the CLI

After you receive the custom certificates, you can use the CLI to replace each machine certificate.

You must have received a certificate for each machine from your third-party or enterprise CA.

- Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
- CRT format
- x509 version 3
- SubjectAltName must contain DNS Name=<machine_FQDN>.
- Contains the following Key Usages: Digital Signature, Key Encipherment

Perform the steps on each vCenter Server host.

You must have the following information before you can start replacing the certificates:

- Password for administrator@vsphere.local
- Valid Machine SSL custom certificate (.crt file)
- Valid Machine SSL custom key (.key file)
- Valid custom certificate for Root (.crt file)

1. Back up the current machine SSL certificate.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry getcert --store MACHINE_SSL_CERT --alias __MACHINE_CERT > oldmachine.crt
/usr/lib/vmware-vmafd/bin/vecs-cli entry getkey --store MACHINE_SSL_CERT --alias __MACHINE_CERT > oldmachinekey.key
```

2. Log in to each host and add the new machine certificates that you received from the certificate authority to VECS.

All hosts need the new certificate in the local certificate store to communicate over SSL.

a) Delete the existing certificate.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry delete --store MACHINE_SSL_CERT --alias __MACHINE_CERT
```

b) Add the new certificate.

```
/usr/lib/vmware-vmafd/bin/vecs-cli entry create --store MACHINE_SSL_CERT --alias __MACHINE_CERT --cert <cert-file-path> --key <key-file-path>
```

3. Extract the hash of the old certificate to be replaced.

```
openssl x509 -in <path_to_old_machinessl_certificate> -noout -sha1 -fingerprint
```

Output similar to the following appears:

```
SHA1 Fingerprint=13:1E:60:93:E4:E6:59:31:55:EB:74:51:67:2A:99:F8:3F:04:83:88
```

4. Update the lookup service registration endpoint manually.

```
/usr/lib/vmware-lookupsvc/tools/ls_update_certs.py --url https://<vCenterServer_FQDN>/lookupservice/sdk
--certfile <cert-file-path> --user 'administrator@vsphere.local' --password '<password>' --fingerprint
<SHA1_hash_of_the_old_certificate_to_replace>
```

If you encounter problems running `ls_update_certs.py`, see the VMware Knowledge base article at <https://kb.vmware.com/s/article/95982>.

5. Restart all services.

```
service-control --stop --all && service-control --start --all
```

vSphere Certificates and Services CLI Command Reference

You can manage VMCA (VMware Certificate Authority), VECS (VMware Endpoint Certificate Store), VMware Directory Service (vmdir), and Security Token Service (STS) certificates by using a set of CLIs. The vSphere Certificate Manager utility supports many related tasks as well, but the CLIs are required for manual certificate management and for managing other services.

You normally access the CLI tools for managing certificates and associated services by using SSH to connect to the appliance shell. See the VMware knowledge base article at <https://kb.vmware.com/s/article/2100508> for more information.

[Manual vSphere Certificate Replacement](#) gives examples for replacing certificates using CLI commands.

Table 110: vSphere CLI Tools for Managing Certificates and Associated Services

CLI	Description	See
<code>certool</code>	Generate and manage certificates and keys. Part of VMCAD, the VMware Certificate Management service.	certool Initialization Commands Reference
<code>vecs-cli</code>	Manage the contents of VMware Certificate Store instances. Part of VMware Authentication Framework Daemon (VMAFD).	vecs-cli Command Reference
<code>dir-cli</code>	Create and update certificates in VMware Directory Service. Part of VMAFD.	dir-cli Command Reference
<code>sso-config.sh</code>	Manage STS certificates.	Command-line help. Entering <code>sso-config .sh</code> with no options displays the command-line help.
<code>service-control</code>	Start or stop services, for example as part of a certificate replacement workflow.	Run this command to stop services before running other CLI commands.

vSphere CLI Locations

By default, you find the CLIs in the following locations.

```
/usr/lib/vmware-vmafd/bin/vecs-cli
/usr/lib/vmware-vmafd/bin/dir-cli
/usr/lib/vmware-vmca/bin/certool
/opt/vmware/bin/sso-config.sh
```

Note: The `service-control` command does not require that you specify the path.

Required Privileges for Running vSphere CLIs

Required privileges depend on the CLI that you are using and on the command that you want to run. For example, for most certificate management operations, you have to be an administrator for the local vCenter Single Sign-On domain (vsphere.local by default). Some commands are available for all users.

dir-cli	You must be a member of the Administrators group in the local domain (vsphere.local by default) to run <code>dir-cli</code> commands. If you do not specify a user name and password, you are prompted for the password for the administrator of the local vCenter Single Sign-On domain, <code>administrator@vsphere.local</code> by default.
vecs-cli	Initially, only the store owner and users with blanket access privileges have access to a store. Users in the Administrators group have blanket access privileges. The <code>MACHINE_SSL_CERT</code> and <code>TRUSTED_ROOTS</code> stores are special stores. Only the root user or the administrator user, depending on the type of installation, has complete access.
certool	Most of the <code>certool</code> commands require that the user is in the Administrators group. All users can run the following commands. • <code>genselfcacert</code> • <code>initscr</code> • <code>getdc</code> • <code>waitVMDIR</code> • <code>waitVMCA</code> • <code>genkey</code> • <code>viewcert</code>

Changing the certool Configuration Options

When you run `certool --gencert` or certain other certificate initialization or management commands, the command reads all the values from a configuration file. You can edit the existing file, override the default configuration file with the `--config=<file name>` option, or override values on the command line.

The configuration file, `certool.cfg`, is located in the `/usr/lib/vmware-vmca/share/config/` directory by default.

The file has several fields with the following default values:

```
Country = US
Name= Acme
Organization = AcmeOrg
OrgUnit = AcmeOrg Engineering
State = California
Locality = Palo Alto
IPAddress = 127.0.0.1
Email = email@acme.com
Hostname = server.acme.com
```

Note: The OU (organizationalUnitName) field is no longer mandatory.

You can change the values by specifying a modified file on the command line, or by overriding individual values on the command line, as follows.

- Create a copy of the configuration file and edit the file. Use the `--config` command-line option to specify the file. Specify the full path to avoid path name issues.

- `/usr/lib/vmware-vmca/bin/certool --gencert --config /tmp/myconfig.cfg`
- Override individual values on the command line. For example, to override Locality, run this command:
`/usr/lib/vmware-vmca/bin/certool --gencert --privkey=private.key --Locality="Mountain View"`

Specify `--Name` to replace the CN field of the Subject name of the certificate.

- For solution user certificates, the name is `<sol_user name>@<domain>` by convention, but you can change the name if a different convention is used in your environment.
- For machine SSL certificates, the FQDN of the machine is used.
 VMCA allows only one DNSName (in the `Hostname` field) and no other Alias options. If the IP address is specified by the user, it is stored in `SubAltName` as well.

Use the `--Hostname` parameter to specify the DNSName of a certificate's `SubAltName`.

certool Initialization Commands Reference

The `certool` initialization commands allow you to generate certificate signing requests, view and generate certificates and keys that are signed by VMware Certificate Authority (VMCA), import root certificates, and perform other certificate management operations.

In many cases, you pass a configuration file in to a `certool` command. See [vSphere Certificates and Services CLI Command Reference](#). See [Replace Existing VMCA-Signed Certificates with New VMCA-Signed Certificates Using the CLI](#) for some usage examples. The command-line help provides details about the options.

certool --initcsr

Generates a Certificate Signing Request (CSR). The command generates a PKCS10 file and a private key.

Option	Description
<code>--gencsr</code>	Required for generating CSRs.
<code>--privkey <key_file></code>	Name of the private key file.
<code>--pubkey <key_file></code>	Name of the public key file.
<code>--csrfile <csr_file></code>	File name for the CSR file to be sent to the CA provider.
<code>--config <config_file></code>	Name of the configuration file. A sample configuration file is located at <code>/usr/lib/vmware-vmca/share/config/certool.cfg</code> . As a best practice, make a copy of the default configuration file and replace the required fields.

Example:

```
certool --gencsr --privkey=<filename> --pubkey=<filename> --csrfile=<filename>
```

certool --selfca

Creates a self-signed certificate and provisions the VMCA server with a self-signed root CA. Using this option is one of the simplest ways to provision the VMCA server. You can instead provision the VMCA server with a third-party root certificate so that VMCA is an intermediate CA. See [Make VMCA Into an Intermediate Certificate Authority Using the CLI](#).

This command generates a certificate that is predated by three days to avoid time zone conflicts.

Option	Description
<code>--selfca</code>	Required for generating a self-signed certificate.

Option	Description
--predate <number_of_minutes>	Allows you to set the Valid Not Before field of the root certificate to the specified number of minutes before the current time. This option can be helpful to account for potential time zone issues. The maximum is three days.
--config <config_file>	Name of the configuration file. A sample configuration file is located at <code>/usr/lib/vmware-vmca/share/config/certool.cfg</code> . As a best practice, make a copy of the default configuration file and replace the required fields.
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.

Example:

```
machine-70-59:/usr/lib/vmware-vmca/bin # ./certool --predate=2280 --selfca --server= 192.0.2.24 --srp-upn=administrator@vsphere.local
```

certool --rootca

Imports a root certificate. Adds the specified certificate and private key to VMCA. VMCA always uses the most recent root certificate for signing, but other root certificates remain trusted until you manually delete them. That means you can update your infrastructure one step at a time, and finally delete certificates that you no longer use.

Option	Description
--rootca	Required for importing a root CA.
--cert <certfile>	Name of the certificate file.
--privkey <key_file>	Name of the private key file. This file must be in PEM encoded format.
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.

Example:

```
certool --rootca --cert=root.cert --privkey=privatekey.pem
```

certool --getdc

Returns the default domain name that is used by vmdir.

Option	Description
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.
--port <port_num>	Optional port number. Defaults to port 389.

Example:

```
certool --getdc
```

certool --waitVMDIR

Wait until the VMware Directory Service is running or until the timeout specified by --wait has elapsed. Use this option along with other options to schedule certain tasks, for example returning the default domain name.

Option	Description
--wait	Optional number of minutes to wait. Defaults to 3.
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.
--port <port_num>	Optional port number. Defaults to port 389.

Example:

```
certool --waitVMDIR --wait 5
```

certool --waitVMCA

Wait until the VMCA service is running or until the specified timeout has elapsed. Use this option in conjunction with other options to schedule certain tasks, for example, generating a certificate.

Option	Description
--wait	Optional number of minutes to wait. Defaults to 3.
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.
--port <port_num>	Optional port number. Defaults to port 389.

Example:

```
certool --waitVMCA --selfca
```

certool --publish-roots

Forces an update of root certificates. This command requires administrative privileges.

Option	Description
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.

Example:

```
certool --publish-roots
```

certool Management Commands Reference

The `certool` management commands allow you to view, generate, and revoke certificates and to view information about certificates.

certool --genkey

Generates a private and public key pair. Those files can then be used to generate a certificate that is signed by VMCA.

Option	Description
--genkey	Required for generating a private and public key.
--privkey <keyfile>	Name of the private key file.
--pubkey <keyfile>	Name of the public key file.

Option	Description
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.

Example:

```
certool --genkey --privkey=<filename> --pubkey=<filename>
```

certool --gencert

Generates a certificate from the VMCA server. This command uses the information in `certool.cfg` or in the specified configuration file. You can use the certificate to provision machine certificates or solution user certificates.

Option	Description
--gencert	Required for generating a certificate.
--cert <certfile>	Name of the certificate file. This file must be in PEM encoded format.
--privkey <keyfile>	Name of the private key file. This file must be in PEM encoded format.
--config <config_file>	Name of the configuration file. A sample configuration file is located at <code>/usr/lib/vmware-vmca/share/config/certool.cfg</code> . As a best practice, make a copy of the default configuration file and replace the required fields.
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.

Example:

```
certool --gencert --privkey=<filename> --cert=<filename> --config=<config_file>
```

certool --getrootca

Prints the current root CA certificate in human-readable form. This output is not usable as a certificate, it is changed to be human readable.

Option	Description
--getrootca	Required for printing the root certificate.
--server <server>	Optional name of the VMCA server. By default, the command uses localhost.

Example:

```
certool --getrootca --server=remoteserver
```

certool --viewcert

Print all the fields in a certificate in human-readable form.

Option	Description
--viewcert	Required for viewing a certificate.

Option	Description
<code>--cert <certfile></code>	Name of the configuration file. A sample configuration file is located at <code>/usr/lib/vmware-vmca/share/config/certtool.cfg</code> . As a best practice, make a copy of the default configuration file and replace the required fields.

Example:

```
certool --viewcert --cert=<filename>
```

certool --enumcert

List all certificates that the VMCA server knows about. The required filter option lets you list all certificates or only revoked, active, or expired certificates.

Option	Description
<code>--enumcert</code>	Required for listing all certificates.
<code>--filter [all active]</code>	Required filter. Specify all or active. The revoked and expired options are not currently supported.

Example:

```
certool --enumcert --filter=active
```

certool --status

Sends a specified certificate to the VMCA server to check whether the certificate has been revoked. Prints

Certificate: REVOKED if the certificate is revoked, and Certificate: ACTIVE otherwise.

Option	Description
<code>--status</code>	Required to check the status of a certificate.
<code>--cert <certfile></code>	Name of the configuration file. A sample configuration file is located at <code>/usr/lib/vmware-vmca/share/config/certtool.cfg</code> . As a best practice, make a copy of the default configuration file and replace the required fields.
<code>--server <server></code>	Optional name of the VMCA server. By default, the command uses localhost.

Example:

```
certool --status --cert=<filename>
```

certool --genselfcacert

Generates a self-signed certificate based on the values in the configuration file. This command generates a certificate that is predated by three days to avoid time zone conflicts.

Option	Description
<code>--genselfcacert</code>	Required for generating a self-signed certificate.
<code>--outcert <cert_file></code>	Name of the certificate file. This file must be in PEM encoded format.

Option	Description
<code>--outprivkey <key_file></code>	Name of the private key file. This file must be in PEM encoded format.
<code>--config <config_file></code>	Name of the configuration file. A sample configuration file is located at <code>/usr/lib/vmware-vmca/share/config/certool.cfg</code> . As a best practice, make a copy of the default configuration file and replace the required fields.

Example:

```
certool --genselfcert --privkey=<filename> --cert=<filename> --config=<config_file>
```

vecs-cli Command Reference

The `vecs-cli` command set allows you to manage instances of VMware Certificate Store (VECS). Use these commands together with `dir-cli` and `certool` to manage your certificate infrastructure and authentication services.

vecs-cli store create

Creates a certificate store.

Option	Description
<code>--name <name></code>	Name of the certificate store.
<code>--server <server-name></code>	Used to specify a server name if you connect to a remote VECS instance.
<code>--upn <user-name></code>	User Principle Name that is used to log in to the server instance specified by <code>--server <server-name></code> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

Example:

```
vecs-cli store create --name <store>
```

vecs-cli store delete

Deletes a certificate store. You cannot delete the `MACHINE_SSL_CERT`, `TRUSTED_ROOTS` and `TRUSTED_ROOT_CRLS` system stores. Users with required privileges can delete solution user stores.

Option	Description
<code>--name <name></code>	Name of the certificate store to delete.
<code>--server <server-name></code>	Used to specify a server name if you connect to a remote VECS instance.
<code>--upn <user-name></code>	User Principle Name that is used to log in to the server instance specified by <code>--server <server-name></code> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

Example:

```
vecs-cli store delete --name <store>
```

vecs-cli store list

List certificate stores.

Option	Description
--server <server-name>	Used to specify a server name if you connect to a remote VECS instance.
--upn <user-name>	User Principle Name that is used to log in to the server instance specified by --server <server-name> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

VECS includes the following stores.

Table 111: Stores in VECS

Store	Description
Machine SSL store (MACHINE_SSL_CERT)	- Used by the reverse proxy service on every vSphere node. - Used by the VMware Directory Service (vmdir) on each vCenter Server node. All services in vSphere 6.0 and later communicate through a reverse proxy, which uses the machine SSL certificate. For backward compatibility, the 5.x services still use specific ports. As a result, some services such as vpxd still have their own port open.
Solution user stores - machine - vpxd - vpxd-extension - vsphere-webclient - wcp	VECS includes one store for each solution user. The subject of each solution user certificate must be unique, for example, the machine certificate cannot have the same subject as the vpxd certificate. Solution user certificates are used for authentication with vCenter Single Sign-On. vCenter Single Sign-On checks that the certificate is valid, but does not check other certificate attributes. The following solution user certificate stores are included in VECS: machine : Used by the license server and the logging service. Note: The machine solution user certificate has nothing to do with the machine SSL certificate. The machine solution user certificate is used for the SAML token exchange. The machine SSL certificate is used for secure SSL connections for a machine. vpxd : vCenter service daemon (vpxd) store. vpxd uses the solution user certificate that is stored in this store to authenticate to vCenter Single Sign-On. vpxd-extension : vCenter extensions store. Includes the Auto Deploy service, inventory service, and other services that are not part of other solution users. vsphere-webclient : vSphere Client store. Also includes some additional services such as the performance chart service. wcp : VMware vSphere® with VMware Tanzu™ store. Also used for vSphere Cluster Services.

Store	Description
	Each vCenter Server node includes a <code>machine</code> certificate.
Trusted root store (TRUSTED_ROOTS)	Contains all trusted root certificates.
vSphere Certificate Manager Utility backup store (BACKUP_STORE)	Used by VMCA (VMware Certificate Manager) to support certificate revert. Only the most recent state is stored as a backup, you cannot go back more than one step.
Other stores	Other stores might be added by solutions. For example, the Virtual Volumes solution adds an SMS store. Do not modify the certificates in those stores unless VMware documentation or a VMware Knowledge Base article instructs you to do so. Note: Deleting the TRUSTED_ROOTS_CRLS store can damage your certificate infrastructure. Do not delete or modify the TRUSTED_ROOTS_CRLS store.

Example:

```
vecs-cli store list
```

vecs-cli store permissions

Grants or revokes permissions to the store. Use either the `--grant` or the `--revoke` option.

The owner of the store can perform all operations, including granting and revoking permissions. The administrator of the local vCenter Single Sign-On domain, `administrator@vsphere.local` by default, has all privileges on all stores, including granting and revoking permissions.

You can use `vecs-cli get-permissions --name <store-name>` to retrieve the current settings for the store.

Option	Description
<code>--name <name></code>	Name of the certificate store.
<code>--user <username></code>	Unique name of the user who is granted permissions.
<code>--grant [read write]</code>	Permission to grant, either read or write.
<code>--revoke [read write]</code>	Permission to revoke, either read or write. Not currently supported.

vecs-cli store get-permissions

Retrieves the current permission settings for the store.

Option	Description
<code>--name <name></code>	Name of the certificate store.
<code>--server <server-name></code>	Used to specify a server name if you connect to a remote VECS instance.
<code>--upn <user-name></code>	User Principle Name that is used to log in to the server instance specified by <code>--server <server-name></code> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

vecs-cli entry create

Creates an entry in VECS. Use this command to add a private key or certificate to a store.

Note: Do not use this command to add root certificates to the TRUSTED_ROOTS store. Instead, use the `dir-cli` command to publish root certificates.

Option	Description
<code>--store <NameOfStore></code>	Name of the certificate store.
<code>--alias <Alias></code>	Optional alias for the certificate. This option is ignored for the trusted root store.
<code>--cert <certificate_file_path></code>	Full path of the certificate file.
<code>--key <key-file-path></code>	Full path of the key that corresponds to the certificate. Optional.
<code>--password <password></code>	Optional password for encrypting the private key.
<code>--server <server-name></code>	Used to specify a server name if you connect to a remote VECS instance.
<code>--upn <user-name></code>	User Principle Name that is used to log in to the server instance specified by <code>--server <server-name></code> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

vecs-cli entry list

Lists all entries in a specified store.

Option	Description
<code>--store <NameOfStore></code>	Name of the certificate store.

vecs-cli entry getcert

Retrieves a certificate from VECS. You can send the certificate to an output file or display it as human-readable text.

Option	Description
<code>--store <NameOfStore></code>	Name of the certificate store.
<code>--alias <Alias></code>	Alias of the certificate.
<code>--output <output_file_path></code>	File to write the certificate to.
<code>--text</code>	Displays a human-readable version of the certificate.
<code>--server <server-name></code>	Used to specify a server name if you connect to a remote VECS instance.
<code>--upn <user-name></code>	User Principle Name that is used to log in to the server instance specified by <code>--server <server-name></code> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

vecs-cli entry getkey

Retrieves a key that is stored in VECS. You can send the key to an output file or display it as human-readable text.

Option	Description
--store <NameOfStore>	Name of the certificate store.
--alias <Alias>	Alias for the key.
--output <output_file_path>	Output file to write the key to.
--text	Displays a human-readable version of the key.
--server <server-name>	Used to specify a server name if you connect to a remote VECS instance.
--upn <user-name>	User Principle Name that is used to log in to the server instance specified by --server <server-name> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

vecs-cli entry delete

Deletes an entry in a certificate store. If you delete an entry in VECS, you permanently remove it from VECS. The only exception is the current root certificate. VECS polls vmdir for a root certificate.

Option	Description
--store <NameOfStore>	Name of the certificate store.
--alias <Alias>	Alias for the entry you want to delete.
--server <server-name>	Used to specify a server name if you connect to a remote VECS instance.
--upn <user-name>	User Principle Name that is used to log in to the server instance specified by --server <server-name> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.
-y	Suppresses the confirmation prompt. For advanced users only.

vecs-cli force-refresh

Forces a refresh of VECS. By default, VECS polls vmdir for new root certificate files every 5 minutes. Use this command for an immediate update of VECS from vmdir.

Option	Description
--server <server-name>	Used to specify a server name if you connect to a remote VECS instance.
--upn <user-name>	User Principle Name that is used to log in to the server instance specified by --server <server-name> . When you create a store, it is created in the context of the current user. Therefore, the owner of the store is the current user context and not always the root user.

dir-cli Command Reference

The `dir-cli` utility supports creation and updates to solution users, account management, and management of certificates and passwords in VMware Directory Service (vmdir). You can use `dir-cli` to manage and query the domain functional level of vCenter Server instances.

dir-cli nodes list

Lists all the enhanced linked mode connected vCenter Server systems.

Option	Description
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.
--server <psc_ip_or_fqdn>	Use this option to connect to another vCenter Server to see its replication partners.

dir-cli computer password-reset

Enables you to reset the password of the machine account in the domain.

Option	Description
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.
--live-dc-hostname <server name>	Current name of the vCenter Server instance.

dir-cli service create

Creates a solution user. Primarily used by third-party solutions.

Option	Description
--name <name>	Name of the solution user to create
--cert <cert file>	Path to the certificate file. This can be a certificate signed by VMCA or a third-party certificate.
--ssogroups <comma-separated-groupnames>	Makes the solution user a member of the specified groups.
--wstrustrole <ActAsUser>	Makes the solution user a member of the built-in administrators or users group. In other words, determines whether the solution user has administrative privileges.
--ssoadminrole <Administrator/User>	Makes the solution user a member of the ActAsUser group. The ActAsUser role enables users to act on behalf of other users.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli service list

List the solution users that `dir-cli` knows about.

Option	Description
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.

Option	Description
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli service delete

Delete a solution user in vmdir. When you delete the solution user, all associated services become unavailable to all management nodes that use this instance of vmdir.

Option	Description
--name	Name of the solution user to delete.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli service update

Updates the certificate for a specified solution user, that is, collection of services. After running this command, update the solution user certificate entry in VECS by running the `vecs-cli entry create` command. See [vecs-cli Command Reference](#).

Option	Description
--name <name>	Name of the solution user to update .
--cert <cert_file>	Name of the certificate to assign to the service.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli user create

Creates a regular user inside vmdir. This command can be used for human users who authenticate to vCenter Single Sign-On with a user name and password. Use this command only during prototyping.

Option	Description
--account <name>	Name of the vCenter Single Sign-On user to create.
--user-password <password>	Initial password for the user.
--first-name <name>	First name for the user.
--last-name <name>	Last name for the user.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli user modify

Modifies the specified user inside vmdir.

Option	Description
--account <name>	Name of the vCenter Single Sign-On user to modify.
--password-never-expires	Set this option to true if you are modifying a user account for automated tasks that have to authenticate to vCenter Server, and you want to ensure that the tasks do not stop running because of password expiration. Use this option with care.
--password-expires	Set this option to true if you want to revert the --password-never-expires option.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli user delete

Deletes the specified user inside vmdir.

Option	Description
--account <name>	Name of the vCenter Single Sign-On user to delete.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli user find-by-name

Finds a user inside vmdir by name. The information that this command returns depends on what you specify in the --level option.

Option	Description
--account <name>	Name of the vCenter Single Sign-On user to find.
--level <info level 0 1 2>	Returns the following information: - Level 0 - Account and UPN - Level 1 - level 0 info + First and last name - Level 2 : level 0 + Account deactivated flag, Account locked flag, Password never expires flag, password expired flag and password expiry flag. The default level is 0.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli group modify

Adds a user or group to an existing group.

Option	Description
--name <name>	Name of the group in vmdir.
--add <user_or_group_name>	Name of the user or group to add.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli group list

Lists a specified vmdir group.

Option	Description
--name <name>	Optional name of the group in vmdir. This option allows you to check whether a specific group exists.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli ssogroup create

Create a group inside the local domain (vsphere.local by default).

Use this command if you want to create groups to manage user permissions for the vCenter Single Sign-On domain. For example, if you create a group and then add it to the Administrators group of the vCenter Single Sign-On domain, then all users that you add to that group have administrator permissions for the domain.

It is also possible to give permissions to vCenter inventory objects to groups in the vCenter Single Sign-On domain. See the *vSphere Security* documentation.

Option	Description
--name <name>	Name of the group in vmdir. Maximum length is 487 characters.
--description <description>	Optional description for the group.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli trustedcert publish

Publishes a trusted root certificate to vmdir. After running this command, VECS picks up the certificate change after one minute, or you can run the `vecs-cli force-refresh` command to sync the certificate immediately.

Note:

Starting in vSphere 8.0 Update 3, use either the vSphere Client or the API to publish a trusted root certificate and avoid having to restart services.

Option	Description
--cert <file>	Path to certificate file.
--crl <file>	This option is not supported by VMCA.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.
--chain	Specify this option if you are publishing a chained certificate. No option value is needed.

dir-cli trustedcert unpublish

Unpublishes a trusted root certificate currently in vmdir. Use this command, for example, if you added a different root certificate to vmdir that is now the root certificate for all other certificates in your environment. Unpublishing certificates that are no longer in use is part of hardening your environment.

Note:

Starting in vSphere 8.0 Update 3, use the API to unpublish a trusted root certificate and avoid having to restart services.

Option	Description
--cert-file <file>	Path to the certificate file to unpublish
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli trustedcert list

Lists all trusted root certificates and their corresponding IDs. You need the certificate IDs to retrieve a certificate with `dir-cli trustedcert get`.

Option	Description
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli trustedcert get

Retrieves a trusted root certificate from vmdir and writes it to a specified file.

Option	Description
--id <cert_ID>	ID of the certificate to retrieve. The <code>dir-cli trustedcert list</code> command shows the ID.
--outcert <path>	Path to write the certificate file to.
--outcrl <path>	Path to write the CRL file to. Not currently used.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.

Option	Description
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli password create

Creates a random password that meets the password requirements. This command can be used by third-party solution users.

Option	Description
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli password reset

Allows an administrator to reset a user's password. If you are a non-administrator user who wants to reset a password, use `dir-cli password change` instead.

Option	Description
--account	Name of the account to assign a new password to.
--new	New password for the specified user.
--login <admin_user_id>	The administrator of the local vCenter Single Sign-On domain, administrator@vsphere.local by default.
--password <admin_password>	Password of the administrator user. If you do not specify the password, you are prompted.

dir-cli password change

Allows a user to change their password. You must be the user who owns the account to make this change. Administrators can use `dir-cli password reset` to reset any password.

Option	Description
--account	Account name.
--current	Current password of the user who owns the account.
--new	New password of the user who owns the account.

vSphere Authentication with vCenter Single Sign-On

vCenter Single Sign-On is an authentication broker and security token exchange infrastructure. vCenter Single Sign-On issues a token when a user authenticates. The user can use the token to authenticate to vCenter Server services. The user can then perform the actions that user has privileges for.

Because traffic is encrypted for all communications, and because only authenticated users can perform the actions that they have privileges for, your environment is secure.

Users and service accounts authenticate with a token, or a user name and password. Solution users authenticate with a certificate.

The next step is authorizing the users who can authenticate to perform certain tasks. Usually, you assign vCenter Server privileges, typically by assigning the user to a group that has a role. vSphere includes other permission models such as global permissions. See the *vSphere Security* documentation.

How vCenter Single Sign-On Protects Your Environment

vCenter Single Sign-On allows vSphere components to communicate with each other through a secure token mechanism.

vCenter Single Sign-On uses the following services.

- Authentication of users through either external identity provider federation or the vCenter Server built-in identity provider. The built-in identity provider supports local accounts, Active Directory or OpenLDAP, Integrated Windows Authentication (IWA), and miscellaneous authentication mechanisms (smart card and RSA SecurID).
- Authentication of solution users through certificates.
- Security Token Service (STS).
- SSL for secure traffic.

vCenter Server Built-in Identity Provider

vCenter Server includes a built-in identity provider. By default, vCenter Server uses the vsphere.local domain as the identity source (but you can change the domain during installation). You can configure the vCenter Server built-in identity provider to use Active Directory (AD) as its identity source using LDAP/S, OpenLDAP/S, or Integrated Windows Authentication (IWA). Such configurations allow customers to log in to vCenter Server using their AD accounts.

vCenter Server and an External Identity Provider

In vSphere 7.0 and later, you can configure vCenter Server for an external identity provider using federated authentication. In such a configuration, you replace vCenter Server as the identity provider.

vSphere supports the following identity providers.


- vSphere 7.0 and later: Active Directory Federation Services (AD FS)
- vSphere 8.0 Update 1 and later: Okta
- vSphere 8.0 Update 2 and later: Microsoft Entra ID (formerly called Azure AD)
- Starting in vSphere 8.0 Update 3: PingFederate

When you configure vSphere to use an external identity provider, the external identity provider interacts with the identity sources on behalf of vCenter Server.

User Login with vCenter Server Identity Provider Federated Authentication

When you use an external identity provider to authenticate with vCenter Server, vCenter Server redirects the login request to the external identity provider. The external identity provider authenticates the user with its directory service then issues a token for vCenter Server to use to log in the user.

For example, the following figure shows a detailed look at the user login flow for vCenter Server Identity Provider Federation using AD FS.

Figure 56: vCenter Server User Login Using AD FS Identity Provider Federation

vCenter Server, AD FS, and Active Directory interact as follows:

1. The user starts on the vCenter Server landing page by entering a user name.
2. If the user name is for a federated domain, vCenter Server redirects the authentication request to AD FS.
3. If needed, AD FS prompts the user to log in with Active Directory credentials.
4. AD FS authenticates the user with Active Directory.
5. AD FS issues a security token with group information from Active Directory.
6. vCenter Server uses the token to log in the user.


The user is now authenticated, and can view and modify any objects that the user's role has privileges for.

In case the external identity provider is unreachable, the login process falls back to the vCenter Server landing page, showing an appropriate information message. Users can still log in using their local accounts in the vsphere.local identity source.

The interaction between vCenter Server and Okta, Microsoft Entra ID, or PingFederate is similar to that of AD FS, with the exception that vCenter Server uses VMware Identity Services. See [VMware Identity Services Authentication Process](#).

User Login with the vCenter Server Built-In Identity Provider

The following figure shows the user login flow when vCenter Server acts as the identity provider.

Figure 57: User Login with the vCenter Server Built-In Identity Provider

1. A user logs in to the vSphere Client with a user name and password to access the vCenter Server system or another vCenter service.
2. The vSphere Client passes the login information to the vCenter Single Sign-On service, which checks the SAML token of the vSphere Client. If the vSphere Client has a valid token, vCenter Single Sign-On then checks whether the user is in the configured identity source (for example, Active Directory).
 - If only the user name is used, vCenter Single Sign-On checks in the default domain.
 - If a domain name is included with the user name (*DOMAIN\user1* or *user1@DOMAIN*), vCenter Single Sign-On checks that domain.
3. If the user can authenticate to the identity source, vCenter Single Sign-On returns a token that represents the user to the vSphere Client.
4. The vSphere Client passes the token to the vCenter Server system.
5. vCenter Server checks with the vCenter Single Sign-On server that the token is valid and not expired.
6. The vCenter Single Sign-On server returns the token to the vCenter Server system, using the vCenter Server Authorization Framework to allow user access.

The user is now authenticated, and can view and modify any objects that the user's role has privileges for.

Supported Encryption in vSphere

AES encryption, which is the highest level of encryption, is supported. The supported encryption affects security when vCenter Single Sign-On uses Active Directory as an identity source.

vCenter Server Identity Provider Federation

In vSphere 7.0 and later, vCenter Server supports federated authentication to sign in to vCenter Server.

To enable federated authentication to vCenter Server, you configure a connection to an external identity provider. The identity provider instance that you configure replaces vCenter Server as the identity provider. Currently, vCenter Server supports Active Directory Federation Services (AD FS), Okta, Microsoft Entra ID (formerly called Azure AD), and PingFederate as external identity providers. vCenter Server supports AD FS in vSphere 7.0 and later, Okta in vSphere

8.0 Update 1 and later, Microsoft Entra ID in vSphere 8.0 Update 2 and later, and PingFederate starting in vSphere 8.0 Update 3.

Note: VMware encourages you to use federated authentication as vSphere moves towards token-based authentication. vCenter Server continues to have local accounts, for administrative access and error recovery.

How vCenter Server Identity Provider Federation Works

vCenter Server Identity Provider Federation enables you to configure an external identity provider for federated authentication. In this configuration, the external identity provider interacts with the identity source on behalf of vCenter Server.

vCenter Server Identity Provider Federation Basics

In vSphere 7.0 and later, vCenter Server supports federated authentication. In this scenario, when a user logs in to vCenter Server, vCenter Server redirects the user login to the external identity provider. The user credentials are no longer provided to vCenter Server directly. Instead, the user provides credentials to the external identity provider. vCenter Server trusts the external identity provider to perform the authentication. In the federation model, users never provide credentials directly to any service or application but only to the identity provider. As a result, you "federate" your applications and services, such as vCenter Server, with your identity provider.

vCenter Server External Identity Provider Support

vCenter Server supports the following external identity providers:

- AD FS (vSphere 7.0 and later)
- Okta (vSphere 8.0 Update 1 and later)
- Microsoft Entra ID, formerly called Azure AD (vSphere 8.0 Update 2 and later)
- PingFederate (starting in vSphere 8.0 Update 3)

vCenter Server Identity Provider Federation Benefits

vCenter Server Identity Provider Federation provides the following benefits.

- You can use Single Sign-On with existing federated infrastructure and applications.
- You can improve data center security because vCenter Server never handles the user's credentials.
- You can use the authentication mechanisms, such as multi-factor authentication, supported by the external identity provider.

vCenter Server Identity Provider Federation Architecture

To establish a relying party trust between vCenter Server and an external identity provider, you must establish the identifying information and a shared secret between them. vCenter Server uses the OpenID Connect (OIDC) protocol to receive an identity token that authenticates the user with vCenter Server.

The high-level steps to configure an external identity provider with vCenter Server include:

1. Establishing a relying party trust between vCenter Server and the external identity provider by creating an OIDC configuration. For AD FS, you create an Application Group, or Application. For Okta, Microsoft Entra ID, and PingFederate, you create a native application with OpenID Connect as the sign-on method. The OIDC configuration consists of a Server application and a Web API. The two components specify the information that vCenter Server uses to trust and communicate with the external identity provider.
2. Creating a corresponding identity provider in vCenter Server.
3. Configuring group memberships in vCenter Server to authorize logins from users in the external identity provider domain.

The identity provider administrator must provide the following information to create the vCenter Server identity provider configuration:

- **Client Identifier:** The UUID string that is generated in AD FS when you create the Application Group (or Application) and that identifies the Application Group (or Application), or that is generated in Okta, Microsoft Entra ID, or PingFederate when you create the OpenID Connect application.
- **Shared Secret:** The secret that is generated in AD FS when you create the Application Group (or Application), or that is generated in Okta, Microsoft Entra ID, or PingFederate when you create the OpenID Connect application, and that is used to authenticate vCenter Server with the external identity provider.
- **OpenID Address:** The OpenID Provider Discovery endpoint URL of the external identity provider server, specifying a well-known address that is typically the issuer endpoint concatenated with the path `/.well-known/openid-configuration`. Here is an example OpenID address for an AD FS configuration:

```
https://webserver.example.com/adfs/.well-known/openid-configuration
```

Likewise, here is an example OpenID address for an Okta configuration:

```
https://example.okta.com/oauth2/default/.well-known/openid-configuration
```

Here is an example OpenID address for a Microsoft Entra ID configuration:

```
https://login.microsoftonline.com/11111111-2222-3333-4444-555555555555/v2.0/.well-known/openid-configuration
```

Here is an example OpenID address for a PingFederate configuration:

```
https://pingfederate-fqdn-and-port/.well-known/openid-configuration
```

VMware Identity Services and Federated Authentication

In vSphere 8.0 Update 1 and later, VMware Identity Services provides integration with external identity providers as a federated identity provider. You can think of the VMware Identity Services as a "slimmed" down version of VMware Workspace ONE that is built-in to vSphere.

When you either install or upgrade to vSphere 8.0 Update 1 or later, VMware Identity Services are activated by default on vCenter Server. When you configure Okta, Microsoft Entra ID, or PingFederate as an external identity provider, vCenter Server uses VMware Identity Services to communicate with your Okta, Microsoft Entra ID, or PingFederate server.

vCenter Server supports Okta, Microsoft Entra ID, and PingFederate as an external identity provider in an Enhanced Linked Mode configuration. Even though, in an Enhanced Linked Mode configuration, multiple vCenter Server systems are running VMware Identity Services, only a single vCenter Server and, its VMware Identity Services communicate with your external identity provider server. For example, if you have an Enhanced Linked Mode configuration of three vCenter Server systems, A, B, and C, and you configure the Okta external identity provider on vCenter Server A, vCenter Server A is the only system to handle all the Okta logins. vCenter Server B and vCenter Server C do not directly communicate with the Okta server. To configure VMware Identity Services on other vCenter Server in the ELM configuration to interact with your external IDP server, see [Activation Process for External Identity Providers in Enhanced Linked Mode Configurations](#).

Note: When configuring Okta as an external identity provider, all the vCenter Server systems in an Enhanced Linked Mode configuration must run at least vSphere 8.0 Update 1. For Microsoft Entra ID, the requirement is at least vSphere 8.0 Update 2. For PingFederate the requirement is at least vSphere 8.0 Update 3.

Warning: When using an Enhanced Linked Mode configuration with Okta, Microsoft Entra ID, or PingFederate, you cannot remove the vCenter Server that runs the VMware Identity Services and that communicates with the identity provider from the ELM configuration.

VMware Identity Services Authentication Process

When you configure vCenter Server to use VMware Identity Services to communicate with your external identity provider, the following authentication process occurs:

1. A user logs in to vCenter Server with the vSphere Client.
2. vCenter Single Sign-On delegates the user authentication and redirects the user request to VMware Identity Services.

3. The VMware Identity Services process requests a token from the external identity provider to establish the user session.
4. The external identity provider authenticates the user (it can use Multifactor Authentication (MFA) or SSO credentials) and returns the token to VMware Identity Services.
The token contains the user claims.
5. The VMware Identity Services process validates the identity provider token, generates a corresponding VMware Identity Services token, and sends the VMware Identity Services token to vCenter Single Sign-On.
6. vCenter Single Sign-On validates the token and grants the login request.

Note: AD FS does not use VMware Identity Services for federated authentication.

How vCenter Server Interacts with Users and Groups Pushed by SCIM

When you configure your external identity provider, vCenter Server uses System for Cross-domain Identity Management (SCIM) for user and group management. SCIM is an open standard for automating the exchange of user identity information. A SCIM application that you create on your external IDP server manages the users and groups on external identity provider that you want to push to vCenter Server. vCenter Server also uses SCIM when searching for users and groups to assign permissions to vCenter Server objects.

Note: An AD FS configuration searches the Active Directory by using LDAP. It does not use SCIM.

vCenter Server Identity Provider Federation Components

The following components comprise a vCenter Server Identity Provider Federation configuration:

- A vCenter Server
 - For AD FS: vCenter Server 7.0 or later
 - For Okta: vCenter Server 8.0 Update 1 or later
 - For Microsoft Entra ID: vCenter Server 8.0 Update 2 or later
 - For PingFederate: vCenter Server 8.0 Update 3
- An identity provider service configured on the vCenter Server
- An external identity provider (AD FS, Okta, Microsoft Entra ID, or PingFederate)
- An OpenID Connect (OIDC) configuration:
 - For AD FS: An Application Group (also called an Application)
 - For Okta, Microsoft Entra ID, or PingFederate: An OpenID Connect application
- A System for Cross-domain Identity Management (SCIM) application for user and group management (only for Okta, Microsoft Entra ID, or PingFederate)
- External identity provider groups and users that you map to vCenter Server groups and users
- VMware Identity Services enabled on vCenter Server (only for Okta, Microsoft Entra ID, or PingFederate)
- Optionally, for PingFederate, the SSL certificate, or certificate chain, of the PingFederate server, if this certificate was not issued by a well-known, public Certificate Authority. You import the PingFederate SSL certificate to the vCenter Server.

vCenter Server Identity Provider Federation Caveats and Interoperability

vCenter Server Identity Provider Federation can interoperate with many other VMware features.

As you are planning your vCenter Server Identity Provider Federation strategy, consider possible interoperability limitations.

Authentication Mechanisms

In a vCenter Server Identity Provider Federation configuration, the external identity provider handles the authentication mechanisms (passwords, MFA, biometrics, and so on).

AD FS and Support for a Single Active Directory Domain

When you configure vCenter Server Identity Provider Federation for AD FS, the Configure Main Identity Provider wizard prompts you to enter LDAP information for the single AD domain that contains the users and groups you want to access vCenter Server. vCenter Server derives the AD domain to use for authorization and permissions from the User Base DN that you specify in the wizard. You can add permissions on vSphere objects only for users and groups from this AD domain. Users or groups from AD child domains or other domains in the AD forest are not supported by vCenter Server Identity Provider Federation.

Okta, Microsoft Entra ID, and PingFederate Support for Multiple Domains

When you configure vCenter Server Identity Provider Federation for Okta, Microsoft Entra ID, or PingFederate, the Configure Main Identity Provider wizard enables you to enter LDAP information for the multiple domains that contain the users and groups you want to access vCenter Server.

Password, Lockout, and Token Policies

When vCenter Server acts as the identity provider, you control vCenter Server password, lockout, and token policies for the default domain (vsphere.local, or the domain name you entered when installing vSphere). When using federated authentication with vCenter Server, the external identity provider controls the password, lockout, and token policies for the accounts stored in the identity source such as Active Directory.

Auditing and Compliance

When using vCenter Server Identity Provider Federation, vCenter Server continues to create log entries for successful user logins. However, the external identity provider is responsible for tracking and logging actions such as failed password-entry attempts and user account lockouts. vCenter Server does not log such events because they are no longer visible to vCenter Server. For example, when AD FS is the identity provider, AD FS tracks and logs errors for federated logins. When vCenter Server is the identity provider for local logins, vCenter Server tracks and logs errors for local logins. In a federated configuration, vCenter Server does continue to log user actions post-login.

Existing VMware Product Integration with External Identity Providers

VMware products that are integrated with vCenter Server (for example, VMware Aria Operations, vSAN, NSX, and so on) continue to work as before.

Products That Integrate Post-Login

Products that integrate post-login (that is, they do not require a separate login) continue to work as before.

Simple Authentication for API, SDK, and CLI Access

Existing scripts, products, and other functionality that rely on API, SDK, or CLI commands that use Simple Authentication (that is, user name and password) continue to work as before. Internally, the authentication occurs by passing the user name and password. This passing of the user name and password compromises some of the benefits of using identity federation, because it exposes the password to vCenter Server (and your scripts). Consider migrating to token-based authentication where possible.

Accessing the vCenter Server Management Interface

If the user is a member of the vCenter Server Administrators group, accessing the vCenter Server Management Interface (formerly called vCenter Server Appliance Management Interface or VAMI) is supported.

Entering User Name Text on the AD FS Login Page

The AD FS login page does not support passing text to pre-populate the user name text box. As a result, during federated logins with AD FS, after entering your user name on the vCenter Server landing page and redirecting to the AD FS login page, you must reenter your user name on the AD FS login page. The user name that you enter on the vCenter Server landing page is necessary to redirect the login to the appropriate identity provider, and the user name on the AD FS login page is necessary to authenticate with AD FS. This inability to pass the user name to the AD FS login page is a limitation of AD FS. You cannot configure or change this behavior directly from vCenter Server.

Support for IPv6 Addresses

AD FS, Microsoft Entra ID, and Ping Federate support IPv6 addresses. Okta does not support IPv6 addresses.

VMware Identity Services Single Instance Configuration

By default, when you either install or upgrade to vSphere 8.0 Update 1 or later, VMware Identity Services are enabled on vCenter Server. When you configure Okta, Microsoft Entra ID, or PingFederate in an Enhanced Link Mode configuration, you use VMware Identity Services on a single vCenter Server system. For example, if you use Okta in an Enhanced Mode Link configuration that consists of three vCenter Server systems, only one vCenter Server its instance of VMware Identity Services is used to communicate with the Okta server.

Warning:

In an ELM configuration that uses VMware Identity Services, if the vCenter Server system that communicates with the external identity provider becomes unavailable, you can configure VMware Identity Services on other vCenter Server in the ELM configuration to interact with your external IDP server. See [Activation Process for External Identity Providers in Enhanced Linked Mode Configurations](#).

Reconfiguring the Primary Network Identifier

Reconfiguring the Primary Network Identifier (PNID) of the vCenter Server requires that you update the external identity provider configuration as follows.

- AD FS: Add the new Redirect URIs to the AD FS server.
- Okta: Reconfigure Okta. See [Configure vCenter Server Identity Provider Federation for Okta](#) and follow the steps to create the identity provider on vCenter Server.
- Microsoft Entra ID: Reconfigure Entra ID. See [Configure vCenter Server Identity Provider Federation for Microsoft Entra ID](#) and follow the steps to create the identity provider on vCenter Server.
- PingFederate: Reconfigure PingFederate. See [Configure vCenter Server Identity Provider Federation for PingFederate](#) and follow the steps to create the identity provider on vCenter Server.

vCenter Server Identity Provider Federation Life Cycle

When managing the life cycle of vCenter Server Identity Provider Federation, there are some specific considerations.

You can manage your vCenter Server Identity Provider Federation life cycle in the following ways.

Migrating from Using Active Directory to an External Identity Provider

If you are using Active Directory as your identity source for vCenter Server, migrating to using an external identity provider is straight forward. If your Active Directory groups and roles match your identity provider groups and roles, you do not

need to take any additional action. When the groups and roles do not match, then you must perform some additional work. If vCenter Server is a domain member, consider removing it from the domain as it is not needed or used for identity federation.

Cross-Domain Repointing and Migration

vCenter Server Identity Provider Federation supports cross-domain repointing, that is, moving a vCenter Server from one vSphere SSO domain to another. The repointed vCenter Server receives the replicated identity provider configuration from the vCenter Server system, or systems, to which it was pointed.

In general, you do not need to perform any additional identity provider reconfiguration for a cross-domain repoint, unless one of the following is true.

1. The identity provider configuration of the repointed vCenter Server differs from the identity provider configuration of the vCenter Server to which it was pointed.
2. This is the first time the repointed vCenter Server is receiving an identity provider configuration.

In these cases, some additional work is required. For example, for AD FS, you must add the vCenter Server system's Redirect URIs to the corresponding Application Group on the AD FS server. For example, if vCenter Server 1 with AD FS Application Group A (or no AD FS configuration) is repointed to vCenter Server 2 with AD FS Application Group B, you must add the Redirect URIs of vCenter Server 1 to Application Group B.

Users and Group Synchronization and vCenter Server Backup and Restore

Depending on when you synchronize your users and groups with vCenter Server and when you back up your vCenter Server, if you must restore your vCenter Server, you might need to re-synchronize your users and groups that were pushed by SCIM.

To restore a deleted user or group, you cannot simply push the user or group from the external identity provider to vCenter Server. You must update the SCIM 2.0 application on the external identity provider with the missing user or group. See [Restore Deleted SCIM Users and Groups](#).

vCenter Server Identity Provider Federation and Enhanced Linked Mode


When you enable Identity Provider Federation in vCenter Server environments using enhanced linked mode, authentication and workflows continue to work as before.

If you use Enhanced Linked Mode configuration, note the following when logging in to vCenter Server using federated authentication.

- Users continue to see the same inventory, and can perform the same actions, based on the vCenter Server permissions and roles model.
- vCenter Server hosts in enhanced linked mode are not required to have access to each other's identity providers. For example, consider two vCenter Server systems A and B, and that use enhanced linked mode. After vCenter Server A authorizes a user, then the user is authorized on vCenter Server B as well.

Enhanced Link Mode and AD FS

The following figure shows the authentication workflow when using AD FS with enhanced linked mode.

Figure 58: Enhanced Linked Mode and AD FS Identity Provider Federation

1. Two vCenter Server nodes are deployed in Enhanced Linked Mode configuration.
2. The AD FS setup has been configured on vCenter Server A using the Change Identity Provider wizard in the vSphere Client. Group memberships and permissions have also been established for AD FS users or groups.
3. vCenter Server A replicates the AD FS configuration to vCenter Server B.
4. All Redirect URIs for both vCenter Server nodes are added to the OAuth Application Group in AD FS. Only one OAuth Application Group is created.
5. When a user logs into and is authorized by vCenter Server A, the user is also authorized on vCenter Server B. If the user logs in to vCenter Server B first, the same holds true.

Enhanced Linked Mode Configuration Scenarios with AD FS

vCenter Server enhanced linked mode supports the following configuration scenarios for AD FS. In this section, the terms "AD FS settings" and "AD FS configuration" refer to the settings that you configure in the vSphere Client using the Change Identity Provider wizard, and any group memberships or permissions that you have established for AD FS users or groups.

Enable AD FS on an Existing Enhanced Linked Mode Configuration

High-level steps:

1. Deploy N vCenter Server nodes in Enhanced Linked Mode configuration.
2. Configure AD FS on one of the linked vCenter Server nodes.
3. The AD FS configuration is replicated to all other (N-1) vCenter Server nodes.
4. Add all Redirect URIs for all N vCenter Server nodes to the configured OAuth Application Group in AD FS.

Link a New vCenter Server to an Existing Enhanced Linked Mode AD FS Configuration

High-level steps:

1. (Prerequisite) Set up AD FS on a vCenter Server N-node Enhanced Linked Mode configuration.
2. Deploy a new independent vCenter Server node.
3. Repoint the new vCenter Server to the N-node AD FS enhanced linked mode domain, using one of the N nodes as its replication partner.
4. All AD FS settings in the existing Enhanced Linked Mode configuration are replicated to the new vCenter Server. The AD FS settings that are in the N-node AD FS enhanced linked mode domain overwrite any existing AD FS settings on the newly linked vCenter Server.
5. Add all Redirect URIs for the new vCenter Server to the existing configured OAuth Application Group in AD FS.

Unlink a vCenter Server from an Enhanced Linked Mode AD FS Configuration


High-level steps:

1. (Prerequisite) Set up AD FS on an N-node vCenter Server Enhanced Linked Mode configuration.
2. Unregister one of the vCenter Server hosts in the N-node configuration and repoint it to a new domain to unlink it from the N-node configuration.
3. The domain repointing process does not preserve SSO settings, so all AD FS settings on the unlinked vCenter Server node are reverted and lost. To continue using AD FS on this vCenter Server unlinked node, you must reconfigure AD FS from the beginning or you must relink the vCenter Server to an Enhanced Linked Mode configuration where AD FS is already set up.

Enhanced Linked Mode and Okta, Microsoft Entra ID, or PingFederate Identity Provider Federation

The following figure shows the authentication workflow when using Okta, Microsoft Entra ID, or PingFederate with enhanced linked mode.

Figure 59: Enhanced Linked Mode and Okta, Microsoft Entra ID, or PingFederate Identity Provider Federation



Note: When configuring Okta, Microsoft Entra ID, or PingFederate as an external identity provider, all the vCenter Server systems in an Enhanced Linked Mode configuration must run at least vSphere 8.0 Update 1 for Okta, vSphere 8.0 Update 2 for Microsoft Entra ID, and vSphere 8.0 Update 3 for PingFederate.

1. Two vCenter Server nodes are deployed in Enhanced Linked Mode configuration.
2. The Okta, Microsoft Entra ID, or PingFederate setup has been configured on vCenter Server A using the Change Identity Provider wizard in the vSphere Client. Also, group memberships and permissions have also been established for Okta, Microsoft Entra ID, or PingFederate users or groups.
Note: Both vCenter Server A and B have VMware Identity Services enabled, but only the VMware Identity Services on vCenter Server A communicate with the identity provider server.
3. The VMware Identity Services running on vCenter Server A enable vCenter Server B to access its endpoint.
4. The Redirect URI for vCenter Server A is added to the OAuth Application in Okta, Microsoft Entra ID, or PingFederate. Only one OAuth Application is created.
5. When a user logs into and is authorized by vCenter Server A, the user is also authorized on vCenter Server B. If the user logs in to vCenter Server B first, the same holds true.

Enhanced Linked Mode Configuration Scenarios with Okta, Microsoft Entra ID, or PingFederate

vCenter Server enhanced linked mode supports the following configuration scenarios for Okta, Microsoft Entra ID, and PingFederate. In this section, the terms "Okta settings" and "Okta configuration," "Microsoft Entra ID settings" and "Microsoft Entra ID configuration," or "PingFederate settings" and "PingFederate configuration," refer to the settings that you configure in the vSphere Client using the Change Identity Provider wizard, and any group memberships or permissions that you have established for Okta, Microsoft Entra ID, or PingFederate users or groups.

Enable Okta, Microsoft Entra ID, or PingFederate on an Existing Enhanced Linked Mode Configuration

High-level steps:

1. Deploy N vCenter Server nodes in Enhanced Linked Mode configuration.
2. Configure Okta, Microsoft Entra ID, or PingFederate on one of the linked vCenter Server nodes.
3. The VMware Identity Services endpoint information is replicated to all other (N-1) vCenter Server nodes.
The Okta, Microsoft Entra ID, or PingFederate configuration (shared client ID, and so on) information and user/group information is not replicated.

Link a New vCenter Server to an Existing Enhanced Linked Mode Okta, Microsoft Entra ID, or PingFederate Configuration

High-level steps:

1. (Prerequisite) Set up Okta, Microsoft Entra ID, or PingFederate on a vCenter Server N-node Enhanced Linked Mode configuration.
2. Deploy a new independent vCenter Server node.
3. Repoint the new vCenter Server to the N-node Okta, Microsoft Entra ID, or PingFederate enhanced linked mode domain, using one of the *N* nodes as its replication partner.
4. The VMware Identity Services endpoint information is replicated to all other (N-1) vCenter Server nodes.
The Okta, Microsoft Entra ID, or PingFederate configuration (shared client ID, and so on) information and user/group information is not replicated.

Note: You can add a vCenter Server node with an existing VMware Identity Services configuration. In this scenario, the existing VMware Identity Services configuration is replaced with the VMware Identity Services Enhanced Link Mode configuration that it is joining.

You cannot add a vCenter Server node with an existing VMware Identity Services configuration to an ELM configuration that has not been configured with VMware Identity Services. In this scenario, first remove the existing VMware Identity Services configuration from the vCenter Server before adding it to the ELM configuration.

Unlink a vCenter Server from an Enhanced Linked Mode Okta, Microsoft Entra ID, or PingFederate Configuration

High-level steps:

1. (Prerequisite) Set up Okta, Microsoft Entra ID, or PingFederate on an N-node vCenter Server Enhanced Linked Mode configuration.
2. Unregister one of the vCenter Server hosts in the N-node configuration and repoint it to a new domain to unlink it from the N-node configuration.
3. The domain repointing process does not preserve SSO settings, so all Okta, Microsoft Entra ID, or PingFederate settings on the unlinked vCenter Server node are reverted and lost. To continue using Okta, Microsoft Entra ID, or PingFederate on this vCenter Server unlinked node, you must reconfigure Okta, Microsoft Entra ID, or PingFederate from the beginning or you must relink the vCenter Server to an Enhanced Linked Mode configuration where Okta, Microsoft Entra ID, or PingFederate is already set up.

Note: You cannot unlink a vCenter Server with an active VMware Identity Services configuration.

Activation Process for External Identity Providers in Enhanced Linked Mode Configurations

Learn more about the availability considerations in Enhanced Linked Mode configurations with Okta, Microsoft Entra ID, or PingFederate.

- Two or more vCenter Server systems in an Enhanced Linked Mode configuration. For example, the systems are labeled as VC_1, VC_2, VC_3, through VC_N, where N is the number of vCenter Server systems in the Enhanced Linked Mode configuration.
- For Okta and Microsoft Entra ID, all vCenter Server systems must run vSphere 8.0 Update 2 or later. For PingFederate, all vCenter Server systems must run at least vSphere 8.0 Update 3.
- Okta, Microsoft Entra ID, or PingFederate is configured as an external identity provider on one of the vCenter Server systems. For example, the system is labeled as VC_1.
- The external identity provider is configured with all required OAuth2 and SCIM applications.

1. To activate a given vCenter Server VC_i where i is between 2 and N:

a) Obtain local shell access to VC_i to run the activation script.

Note: To perform the following steps, the vCenter Server user account with administrative privileges can be given either on the command line or in the console prompts.

b) Execute 'status' from the activation script to get the current activation state of the vCenter Server.

```
python /usr/lib/vmware-trustmanagement/vmware_identity_services_activation.py status
```

c) If the 'status' command indicates that the vCenter Server is not activated, execute 'activate' from the activation script:

```
python /usr/lib/vmware-trustmanagement/vmware_identity_services_activation.py activate
```

d) If the 'status' command indicates that the vCenter Server is already activated, then execute the 'deactivate' option then execute the 'activate' option.

```
python /usr/lib/vmware-trustmanagement/vmware_identity_services_activation.py deactivate
```

- For example, execute the 'activate' option.
- Alternatively, you can specify the '--force-replace' option in the 'activate' command.

2. Open a browser to the vCenter Server VC_i and log in as an administrator to vCenter Server.
 - a) Navigate to **Home > Administration > Single Sign On > Configuration**.
 - b) Under **User Provisioning**, verify that the **Tenant URL** contains the FQDN of VC_i.
 - c) Copy the **Tenant URL** string and save this information for using it with the external identity provider.
 - d) Under **Secret Token**, click the **Generate**, copy the generated token string, and save this information for using it with the external identity provider.
 - e) Under **OpenID Connect**, verify that the **Redirect URI** contains the FQDN of VC_i.
 - f) Copy the **Redirect URI** string and save this information for using it with the external identity provider.
3. Open a browser to the administration page of the external identity provider.

Note: For more information, refer to the external identity provider specific details to perform the following steps.

 - a) Find the OAuth2 registration that was setup when the external identity provider was originally configured in VC_1.
 - b) Edit the OAuth2 registration and add the Redirect URI that was previously obtained for VC_i.
 - c) If the external identity provider supports SCIM Push configurations with multiple destinations, then:
 - Find the SCIM Push configuration that was setup when the external identity provider was originally configured in VC_1.
 - Edit the SCIM Push configuration and add the **Tenant URL** and **Secret Token** that were previously obtained for VC_i.
 - d) If the external identity provider supports SCIM Push configurations with only one destination:
 - Create a new SCIM Push configuration with the **Tenant URL** and **Secret Token** that were previously obtained for VC_i.
 - Ensure that the SCIM Push configuration is pushing the same user/group data as the SCIM Push configuration that was setup when the external identity provider was originally configured in VC_1.
 - e) Initiate a SCIM push operation to ensure that VC_i is populated with the latest user or group data.

Configuring vCenter Server Identity Provider Federation

After you deploy vCenter Server initially, you can configure an external identity provider for federated authentication.

vSphere 7.0 and later supports Active Directory Federation Services (AD FS). vSphere 8.0 Update 1 and later supports Okta. vSphere 8.0 Update 2 and later supports Microsoft Entra ID (formerly called Azure AD). Starting in vSphere 8.0 Update 3, vSphere supports PingFederate.

You configure vCenter Server Identity Provider Federation from the vSphere Client or the API. You also must perform some configuration on your external identity provider. To configure vCenter Server Identity Provider Federation, you must have vCenter Single Sign-On administrator privileges. Having vCenter Single Sign-On administrator privileges is different from having the Administrator role on vCenter Server or ESXi. In a new installation, only the vCenter Single Sign-On administrator (administrator@vsphere.local by default) can authenticate to vCenter Single Sign-On.


vCenter Server Identity Provider Federation Configuration Process Flow

To configure vCenter Server Identity Provider Federation effectively, you must understand the communication flows that take place.

You can configure vCenter Server Identity Provider Federation for AD FS, Microsoft Entra ID (formerly called Azure AD), Okta, or PingFederate.

vCenter Server Identity Provider Federation Configuration Process Flow for AD FS

The following figure shows the process flow that occurs when configuring vCenter Server Identity Provider Federation for AD FS.

Figure 60: vCenter Server Identity Provider Federation Configuration Process Flow for AD FS

vCenter Server, AD FS, and Active Directory interact as follows.

1. The AD FS administrator configures an AD FS OIDC Application for vCenter Server.
2. The vCenter Server administrator logs in to the vCenter Server using the vSphere Client.
3. The vCenter Server administrator adds an AD FS identity provider to vCenter Server, and also enters information about the Active Directory domain.
vCenter Server needs this information to make an LDAP connection to the Active Directory domain of the AD FS server. Using this connection, vCenter Server searches for users and groups, and adds them to vCenter Server local groups in the next step. See the section titled "Searching the Active Directory Domain" that follows for more information.
4. The vCenter Server administrator configures authorization permissions in vCenter Server for AD FS users.
5. The AD FS Provider queries the VcIdentityProviders API to get the LDAP connection information for the Active Directory source.
6. The AD FS Provider searches Active Directory for the queried users or groups to finish the Authorization configuration.

Searching the Active Directory Domain

You configure AD FS as the external identity provider in vCenter Server by using the Configure Main Identity Provider wizard in the vSphere Client. As part of the configuration process, you must enter information about your Active Directory domain, including user and group distinguished name information. Configuration of AD FS for authentication requires this Active Directory connection information. This connection is necessary to search for and map Active Directory user names and groups to roles and permissions in vCenter Server, whereas AD FS is used for the authentication of the user. This step of the Configure Main Identity Provider wizard does not create an Active Directory Over LDAP identity source. Rather, vCenter Server uses this information to establish a valid search-capable connection to your Active Directory domain to find users and groups there.

Consider an example using the following distinguished name entries:

- Base distinguished name for users: cn=Users,dc=corp,dc=local


- Base distinguished name for groups: dc=corp,dc=local
- User name: cn=Administrator,cn=Users,dc=corp,dc=local

If the AdfsUser@corp.local user is a member of the ADGroup@corp.local group, entering this information in the wizard allows a vCenter Server administrator to search for and find the ADGroup@corp.local group and add it to the vCenter Server Administrators@vsphere.local group. As a result, the AdfsUser@corp.local user is granted administrative privileges in vCenter Server upon logging in.

vCenter Server also uses this search process when you configure global permissions for Active Directory users and groups. In both cases, either configuring global permissions, or adding a user or group, you select the domain you entered for your AD FS identify provider from the **Domain** drop-down menu to search for and select users and groups from your Active Directory domain.

vCenter Server Identity Provider Federation Configuration Process Flow Using VMware Identity Services

To configure Okta, Microsoft Entra ID, and PingFederate, you use VMware Identity Services. The following figure shows the process flow that occurs when configuring vCenter Server Identity Provider Federation using VMware Identity Services.

Figure 61: vCenter Server Identity Provider Federation Configuration Process Flow using VMware Identity

Services

vCenter Server, VMware Identity Services, and Active Directory interact as follows.

1. The external IDP administrator configures an OIDC Application for vCenter Server.
2. The vCenter Server administrator logs in to the vCenter Server using the vSphere Client, adds an identity provider to vCenter Server, and also enters domain information.
3. The vCenter Server administrator gives the Redirect URI (obtained from the identity provider configuration page in the vSphere Client) to the identity provider administrator to add to the OIDC Application created in step 2.
4. The external IDP administrator configures a SCIM 2.0 Application.
5. The external IDP administrator assigns the users and groups to the SCIM 2.0 Application, and pushes users and groups to vCenter Server.
6. The vCenter Server administrator configures authorization permissions in vCenter Server for external IDP users.

External IDP Users and Groups

Because an external identity provider uses System for Cross-domain Identity Management (SCIM) for users and groups, those users and groups reside in your vCenter Server. When you search for users and groups in your external identity provider, for example, to assign permissions, the search happens locally on the vCenter Server.

vCenter Server also uses this search process when you configure global permissions for external IDP users and groups. In both cases, either configuring global permissions, or adding a user or group, you select the domain you entered for your identity provider from the **Domain** drop-down menu to search for and select users and groups from your domain.

Use the Trusted Root Certificates Store Instead of the JRE truststore

If you imported a root CA certificate issued by your own internal Certificate Authority to the JRE truststore in vSphere 7.0, starting in vSphere 7.0 Update 1, you can register the certificate to the Trusted Root Certificates Store.

To configure vCenter Server Identity Provider Federation in vSphere 7.0 with a root CA certificate that was issued by your own internal Certificate Authority, you had to import it to the JRE truststore. Starting in vSphere 7.0 Update 1, you can register the certificate to the Trusted Root Certificates Store. This change means that you should add the root CA certificate that was issued by your own internal Certificate Authority to the Trusted Root Certificates Store (also called the VMware Endpoint Certificate Store, or VECS). Certificates in the JRE truststore continue to function, however, vCenter Server is standardizing on using the Trusted Root Certificates Store.

1. Log in with the vSphere Client to the vCenter Server.

Note: To learn more, see [Add a Trusted Root Certificate to the Certificate Store Using the vSphere Client](#).

2. Navigate to **Administration > Certificates > Certificate Management**.
3. Next to **Trusted Root Certificates**, click **Add**.
4. Browse for the AD FS root certificate and click **Add**.

The certificate is added in a panel under **Trusted Root Certificates**.

Configure vCenter Server Identity Provider Federation for AD FS

After installing or upgrading to vSphere 7.0 or later, you can configure vCenter Server Identity Provider Federation for AD FS as an external identity provider.

Note: This process of configuring an AD FS identity provider requires that you have administrative access to both your vCenter Server and your AD FS server. During the configuration process, you enter information first in your vCenter Server, then in your AD FS server, then in your vCenter Server.

Active Directory Federation Services requirements:

- AD FS for Windows Server 2016 or later must already be deployed.
- AD FS must be connected to Active Directory.
- An Application Group for vCenter Server must be created in AD FS as part of the configuration process. See the VMware knowledge base article at <https://kb.vmware.com/s/article/78029>.
- An AD FS server certificate (or a CA or intermediate certificate that signed the AD FS server certificate) that you add to the Trusted Root Certificates Store.
- You have created a vCenter Server administrators group in AD FS that contains the users you want to grant vCenter Server administrator privileges to.

For more information about configuring AD FS, see the Microsoft documentation.

vCenter Server and other requirements:

- vSphere 7.0 or later

- vCenter Server must be able to connect to the AD FS discovery endpoint, and the authorization, token, logout, JWKS, and any other endpoints advertised in the discovery endpoint metadata.
- You need the **VcIdentityProviders > Manage** privilege to create, update, or delete a vCenter Server Identity Provider that is required for federated authentication. To limit a user to view the Identity Provider configuration information only, assign the **VcIdentityProviders > Read** privilege.

Note: These instructions are for vSphere 8.0 Update 1 and later. For vSphere 8.0, see the topic about configuring vCenter Server identity provider federation for AD FS in the *vSphere Authentication* documentation at <https://docs.vmware.com/en/VMware-vSphere/8.0/vsphere-documentation-80.zip>.

vCenter Server supports only one configured external identity provider (one source), and the vsphere.local identity source. You cannot use multiple external identity providers. vCenter Server Identity Provider Federation uses OpenID Connect (OIDC) for user login to vCenter Server.

This task describes how to add an AD FS group to the vSphere Administrators group as the way to control permissions. You can also configure privileges using AD FS Authorization through global or object permissions in vCenter Server. See the *vSphere Security* documentation for details about adding permissions.

CAUTION:

If you use an Active Directory identity source that you previously added to vCenter Server for your AD FS identity source, do not delete that existing identity source from vCenter Server. Doing so causes a regression with previously assigned roles and group memberships. Both the AD FS user with global permissions and users that were added to the Administrators group will not be able to log in.

Workaround: If you do not need the previously assigned roles and group memberships, and want to remove the previous Active Directory identity source, remove the identity source before creating the AD FS provider and configuring group memberships in vCenter Server.

1. Log in with the vSphere Client to the vCenter Server.
2. Add your AD FS server certificate (or a CA or intermediate certificate that signed the AD FS server certificate) to the Trusted Root Certificates Store.

Note: To learn more, see [Add a Trusted Root Certificate to the Certificate Store Using the vSphere Client](#).

- a) Navigate to **Administration > Certificates > Certificate Management**.
- b) Next to **Trusted Root Store**, click **Add**.
- c) Browse for the AD FS certificate and click **Add**.

The certificate is added in a panel under **Trusted Root Certificates**.

3. Start to create the identity provider on vCenter Server.
 - a) Use the vSphere Client to log in as an administrator to vCenter Server.
 - b) Navigate to **Home > Administration > Single Sign On > Configuration**.
 - c) Click **Change Provider** and select **ADFS**.

The **Configure Main Identity Provider** wizard opens.

- d) In the **Prerequisites** panel, review the AD FS and the vCenter Server requirements.
- e) Click **Run Prechecks**.

If the precheck finds errors, click **View Details** and take steps to resolve the errors as indicated.

- f) When the Precheck passes, click the confirmation check box then click **Next**.
- g) In the **Users and Groups** panel, enter user and group information for the Active Directory over LDAP connection to search for users and groups.

vCenter Server derives the AD domain to use for authorization and permissions from the Base Distinguished Name for users. You can add permissions on vSphere objects only for users and groups from this AD domain. Users

or groups from AD child domains or other domains in the AD forest are not supported by vCenter Server Identity Provider Federation.

Option	Description
Base distinguished name for users	Base Distinguished Name for users.
Base distinguished name for groups	The base Distinguished Name for groups.
Username	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups.
Password	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups.
Primary server URL	Primary domain controller LDAP server for the domain. Use the format <code>ldap://hostname:port</code> or <code>ldaps://hostname:port</code> . The port is typically 389 for LDAP connections and 636 for LDAPS connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for LDAP and 3269 for LDAPS. A certificate that establishes trust for the LDAPS endpoint of the Active Directory server is required when you use <code>ldaps://</code> in the primary or secondary LDAP URL.
Secondary server URL	Address of a secondary domain controller LDAP server that is used for failover.
SSL certificates	If you want to use LDAPS with your Active Directory LDAP Server or OpenLDAP Server identity source, click Browse to select a certificate.

h) Click **Next**.

i) In the **OpenID Connect** panel, copy the Redirect URI and the Log-out Redirect URI.

Leave the other fields blank for now. You return to the **OpenID Connect** panel after creating the OpenID Connection configuration in the next step.

4. Create an OpenID Connect configuration in AD FS and configure it for vCenter Server.

To establish a relying party trust between vCenter Server and an identity provider, you must establish the identifying information and a shared secret between them. In AD FS, you do so by creating an OpenID Connect configuration known as an Application Group, which consists of a Server application and a Web API. The two components specify the information that vCenter Server uses to trust and communicate with the AD FS server. To enable OpenID Connect in AD FS, see the VMware knowledge base article at <https://kb.vmware.com/s/article/78029>.

Note the following when you create the AD FS Application Group.

- You need the two vCenter Server Redirect URIs obtained in the preceding step.
- Copy the following information from the AD FS Application Group to a file or write it down for use when finishing the creation of the vCenter Server identity provider in the next step.
 - Client Identifier
 - Shared secret
 - OpenID address of the AD FS server

Note: If necessary, obtain the OpenID address of your AD FS server by running the following PowerShell command as an AD FS administrator.

```
Get-AdfsEndpoint | Select FullUrl | Select-String openid-configuration
```

Copy the URL that is returned (select only the URL itself, not the closing bracket or the initial "`@{FullUrl=`" part).

5. In the vCenter Server **OpenID Connect** panel:
 - a) Enter the following information, which you obtained from the previous step when creating the AD FS Application Group:
 - Client Identifier
 - Shared Secret
 - OpenID Address

The Identity Provider Name is automatically filled in as Microsoft ADFS.
 - b) Click **Next**.
6. Review the information and click **Finish**.
vCenter Server creates the AD FS identity provider and displays the configuration information.
7. Configure group membership in vCenter Server for AD FS Authorization.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
 - c) Click the **Groups** tab.
 - d) Click the **Administrators** group and click **Add Members**.
 - e) Select the domain from the drop-down menu.
 - f) In the text box below the drop-down menu, enter the first few characters of AD FS group that you want to add then wait for the drop-down selection to appear.
It might take several seconds for the selection to appear as vCenter Server establishes the connection to and searches Active Directory.
 - g) Select the AD FS group and add it to the Administrators group.
 - h) Click **Save**.
8. Verify logging in to vCenter Server with an Active Directory user.

Configure vCenter Server Identity Provider Federation for Okta

After installing or upgrading to vSphere 8.0 Update 1 or later, you can configure vCenter Server Identity Provider Federation for Okta as an external identity provider.

Okta requirements:

- You are using Okta and have a dedicated domain space, for example, <https://your-company.okta.com>.
- To perform OIDC logins and manage user and group permissions, you must create the following Okta applications.
 - An Okta native application with OpenID Connect as the sign-on method. The native application must include the grant types of authorization code, refresh token, and resource owner password.
 - A System for Cross-domain Identity Management (SCIM) 2.0 application with an OAuth 2.0 Bearer Token to perform user and group synchronization between the Okta server and the vCenter Server.

See the VMware knowledge base article at <https://kb.vmware.com/s/article/90835>.

- You have identified the Okta users and groups that you want to share with vCenter Server. This sharing is a SCIM operation (not an OIDC operation).

Okta connectivity requirements:

- vCenter Server must be able to connect to the Okta discovery endpoint, and the authorization, token, JWKS, and any other endpoints advertised in the discovery endpoint metadata.
- Okta must also be able to connect with vCenter Server to send user and group data for the SCIM provisioning.

vCenter Server requirements:

- vSphere 8.0 Update 1 or later

- On the vCenter Server where you want to create the Okta identity source, verify that the VMware Identity Services are activated.

Note: When you install or upgrade to vSphere 8.0 Update 1 or later, VMware Identity Servers are activated by default. You can use the vCenter Server Management Interface to confirm the status of the VMware Identity Services. See [Stop and Start the VMware Identity Services](#).

vSphere privileges requirements:

- You must have the **VcIdentityProviders > Manage** privilege to create, update, or delete a vCenter Server Identity Provider that is required for federated authentication. To limit a user to view the Identity Provider configuration information only, assign the **VcIdentityProviders > Read** privilege.

Enhanced Linked Mode requirements:

- You can configure vCenter Server Identity Provider Federation for Okta in an Enhanced Linked Mode configuration. When you configure Okta in an Enhanced Link Mode configuration, you configure the Okta identity provider to use VMware Identity Services on a single vCenter Server system. For example, if your Enhanced Mode Link configuration consists of two vCenter Server systems, only one vCenter Server and its instance of VMware Identity Services is used to communicate with the Okta server. If this vCenter Server system becomes unavailable, you can configure VMware Identity Services on other vCenter Server in the ELM configuration to interact with your Okta server. For more information, see [Activation Process for External Identity Providers in Enhanced Linked Mode Configurations](#).
- When configuring Okta as an external identity provider, all the vCenter Server systems in an Enhanced Linked Mode configuration must run at least vSphere 8.0 Update 1.

Networking requirements:

- If your network is not publicly available, you must create a network tunnel between your vCenter Server system and your Okta server, then use the appropriate publicly accessible URL as the Base Uri.

vCenter Server supports only one configured external identity provider (one source), and the vsphere.local identity source (local source). You cannot use multiple external identity providers. vCenter Server Identity Provider Federation uses OpenID Connect (OIDC) for user login to vCenter Server.

You can configure privileges using Okta groups and users through global or object permissions in vCenter Server. See the *vSphere Security* documentation for details about adding permissions.

1. Create an OpenID Connect application in Okta and assign groups and users to the OpenID Connect application. To create the OpenID Connect application and assign groups and users, see the VMware knowledge base article at <https://kb.vmware.com/s/article/90835>. Follow the steps in the section titled "Create the OpenID Connect Application." After you create the Okta OpenID Connect application, copy the following information from the Okta OpenID Connect application to a file for use when configuring the vCenter Server identity provider in the next step.
 - Client Identifier
 - Client secret (shown as Shared secret in the vSphere Client)
 - Active Directory domain information, or Okta domain information if you are not running Active Directory

2. To create the identity provider on vCenter Server:

- a) Use the vSphere Client to log in as an administrator to vCenter Server.
- b) Navigate to **Home > Administration > Single Sign On > Configuration**.
- c) Click **Change Provider** and select **Okta**.

The **Configure Main Identity Provider** wizard opens.

- d) In the **Prerequisites** panel, review the Okta and the vCenter Server requirements.
- e) Click **Run Prechecks**.

If the precheck finds errors, click **View Details** and take steps to resolve the errors as indicated.

- f) When the Precheck passes, click the confirmation checkbox then click **Next**.
- g) In the **Directory Information** panel, enter the following information.

- **Directory Name:** Name of the local directory to create on vCenter Server that stores the users and groups pushed from Okta. For example, `vcenter-okta-directory`.
- **Domain Name(s):** Enter the Okta domain names that contain the Okta users and groups you want to synchronize with vCenter Server.

After you enter your Okta domain name, click the Plus icon (+) to add it. If you enter multiple domain names, specify the default domain.

- h) Click **Next**.
- i) In the **OpenID Connect** panel, enter the following information.
 - **Redirect UI:** Filled in automatically. You give the redirect UI to your Okta administrator for use in creating the OpenID Connect application.
 - **Identity Provider Name:** Filled in automatically as Okta.
 - **Client Identifier:** Obtained when you created the OpenID Connect application in Okta in step 1. (Okta refers to Client Identifier as the Client ID.)
 - **Shared Secret:** Obtained when you created the OpenID Connect application in Okta in step 1. (Okta refers to Shared Secret as the Client Secret.)
 - **OpenID Address:** Takes the form `https://Okta domain space/oauth2/default/.well-known/openid-configuration`.
For example, if your Okta domain space is `example.okta.com`, then the OpenID Address is: `https://example.okta.com/oauth2/default/.well-known/openid-configuration`
See <https://developer.okta.com/docs/reference/api/oidc/#well-known-openid-configuration> for more information.
- j) Click **Next**.
- k) Review the information and click **Finish**.

vCenter Server creates the Okta identity provider and displays the configuration information.

- l) If necessary, scroll down and click the **Copy** icon for the Redirect URI and save it to a file.

You use the Redirect URI in the Okta OpenID Connection application.

- m) Click the **Copy** icon for the Tenant URL and save it to a file.

Note: If your network is not publicly available, you must create a network tunnel between your vCenter Server system and your Okta server. After creating the network tunnel, use the appropriate publicly accessible URL as the Base Uri.

- n) Under **User Provisioning**, click **Generate** to create the secret token, select the token lifespan from the drop-down, then click **Copy to Clipboard**. Save the token to a secure location.

You use the Tenant URL and the token in the Okta SCIM 2.0 application. The Okta SCIM 2.0 application uses the token to synchronize the Okta users and groups into VMware Identity Services. This information is necessary to push Okta users and groups from Okta to vCenter Server.

3. Return to the VMware knowledge base article at <https://kb.vmware.com/s/article/90835> to update the Okta Redirect URI.
Follow the steps in the section titled "Update the Okta Redirect URI."
4. To create the SCIM 2.0 application, remain in the VMware knowledge base article at <https://kb.vmware.com/s/article/90835>.
Follow the steps in the section titled "Create the SCIM 2.0 Application and Push Users and Groups to vCenter Server." When done creating the SCIM 2.0 application as described in the knowledge base article, continue with the next step.
5. Configure vCenter Server for Okta Authorization.
You can either assign Okta users to a vCenter Server group or assign inventory-level and global permissions to Okta users. The minimum permission required for logging in is Read-Only.
To assign Okta users to a group, see [Add Members to a vCenter Single Sign-On Group](#). To assign inventory-level and global permissions to Okta users, see the topic about managing permissions for vCenter Server components in the *vSphere Security* documentation.
6. Verify logging in to vCenter Server with an Okta user.

Configure vCenter Server Identity Provider Federation for Microsoft Entra ID

After installing or upgrading to vSphere 8.0 Update 2 or later, you can configure vCenter Server Identity Provider Federation for Microsoft Entra ID (formerly called Azure AD) as an external identity provider.

Microsoft Entra ID requirements:

- You are customer of Microsoft and have a Microsoft Entra ID account.

Microsoft Entra ID Connectivity requirements:

- You have created an enterprise (non-gallery) application with OpenID Connect as a sign-on method.
- Add authorization code, refresh token and resource owner password as grant types in the created application.
- For user and group sync, you need to configure VMware Identity Services Gallery Application for SCIM 2.0 provisioning in Microsoft Entra ID with OAuth 2.0 Bearer Token.

vCenter Server requirements:

- vSphere 8.0 Update 2 or later, with the VMware Identity Services activated (they are activated by default).
- On the vCenter Server where you want to create the Microsoft Entra ID identity source, verify that the VMware Identity services are activated.
- The users and groups from the identity provider are provisioned in your vCenter Server.

vSphere privileges requirements:

- You must have the **VcIdentityProviders.Manage** privilege to create, update, or delete a vCenter Server Identity Provider that is required for federated authentication. To limit a user to view the Identity Provider configuration information only, assign the **VcIdentityProviders.Read** privilege.

Enhanced Linked Mode requirements:

- You can configure vCenter Server Identity Provider Federation for Microsoft Entra ID in an Enhanced Linked Mode configuration. When you configure Microsoft Entra ID in an Enhanced Link Mode configuration, you configure the Microsoft Entra ID identity provider to use VMware Identity Services on a single vCenter Server system. For example, if your Enhanced Mode Link configuration consists of two vCenter Server systems, only one vCenter Server and its instance of VMware Identity Services is used to communicate with the Microsoft Entra ID server. If this vCenter Server system becomes unavailable, you can configure VMware Identity Services on other vCenter Server systems in the ELM configuration to interact with your Microsoft Entra ID server. For more information, see [Activation Process for External Identity Providers in Enhanced Linked Mode Configurations](#).

- When configuring Microsoft Entra ID as an external identity provider, all the vCenter Server systems in an Enhanced Linked Mode configuration must run at least vSphere 8.0 Update 2.

Networking requirements:

- If your network is not publicly available, you must create a network tunnel between your vCenter Server system and your Microsoft Entra ID server, then use the appropriate publicly accessible URL as the Base Uri.

vCenter Server supports only one configured external identity provider (one source), and the vsphere.local identity source (local source). You cannot use multiple external identity providers. vCenter Server Identity Provider Federation uses OpenID Connect (OIDC) for user login to vCenter Server.

You can configure privileges using Microsoft Entra ID groups and users through global or object permissions in vCenter Server. See the *vSphere Security* documentation for details about adding permissions.

For a walk-through of the configuration process, see the following video:

[vCenter Authentication: AzureAD/Entra ID integration | vSphere 8 Update 2](#)

1. Create an OpenID Connect application in Microsoft Entra ID and assign groups and users to the OpenID Connect application.
To create the OpenID Connect application and assign groups and users, see the VMware knowledge base article at <https://kb.vmware.com/s/article/94182>. Follow the steps in the section titled "Create the OpenID Connect Application." After you create the OpenID Connect application, copy the following information from the Microsoft Entra ID OpenID Connect application to a file for use when configuring the vCenter Server identity provider in the next step.
 - Client Identifier
 - Client secret (shown as Shared secret in the vSphere Client).
 - Active Directory domain information, or Microsoft Entra ID domain information if you are not running Active Directory.
2. To create the identity provider on vCenter Server:
 - a) Use the vSphere Client to log in as an administrator to vCenter Server.
 - b) Navigate to **Home > Administration > Single Sign On > Configuration**.
 - c) Click **Change Provider** and select **Microsoft Entra ID**.
The **Configure Main Identity Provider** wizard opens.
 - d) In the **Prerequisites** panel, review the Microsoft Entra ID and the vCenter Server requirements.
 - e) Click **Run Prechecks**.
If the precheck finds errors, click **View Details** and take steps to resolve the errors as indicated.
 - f) When the Precheck passes, click the confirmation check box then click **Next**.
 - g) In the **Directory Information** panel, enter the following information.
 - Directory Name: Name of the local directory to create on vCenter Server that stores the users and groups pushed from Microsoft Entra ID. For example, `vcenter-entraid-directory`.
 - Domain Names: Enter the Microsoft Entra ID domain names that contain the Microsoft Entra ID users and groups you want to synchronize with vCenter Server.
After you enter your Microsoft Entra ID domain name, click the Plus icon (+) to add it. If you enter multiple domain names, specify the default domain.
 - h) Click **Next**.
 - i) In the **OpenID Connect** panel, enter the following information.
 - Redirect UI: Filled in automatically. You give the redirect UI to your Microsoft Entra ID administrator for use in creating the OpenID Connect application.
 - Identity Provider Name: Filled in automatically as Microsoft Entra ID.
 - Client Identifier: Obtained when you created the OpenID Connect application in Microsoft Entra ID in step 1. (Microsoft Entra ID refers to Client Identifier as the Client ID.)

- **Shared Secret:** Obtained when you created the OpenID Connect application in Microsoft Entra ID in step 1. (Microsoft Entra ID refers to Shared Secret as the Client Secret.)
- **OpenID Address:** Takes the form `https://Microsoft Entra ID domain space/oauth2/default/.well-known/openid-configuration`.

For example, if your Microsoft Entra ID domain space is `example.EntraID.com`, then the OpenID Address is:
`https://example.EntraID.com/oauth2/default/.well-known/openid-configuration`

j) Click **Next**.

k) Review the information and click **Finish**.

vCenter Server creates the Microsoft Entra ID identity provider and displays the configuration information.

l) If necessary, scroll down and click the **Copy** icon for the Redirect URI and save it to a file.

You use the Redirect URI in the Microsoft Entra ID OpenID Connection application.

m) Click the **Copy** icon for the Tenant URL and save it to a file.

Note: If your network is not publicly available, you must create a network tunnel between your vCenter Server system and your Microsoft Entra ID server. After creating the network tunnel, use the appropriate publicly accessible URL as the Base Uri.

n) Under **User Provisioning**, click **Generate** to create the secret token, select the token lifespan from the drop-down, then click **Copy to Clipboard**. Save the token to a secure location.

You use the Tenant URL and the token in the Microsoft Entra ID SCIM 2.0 application. The Microsoft Entra ID SCIM 2.0 application uses the token to synchronize the Microsoft Entra ID users and groups into VMware Identity Services. This information is necessary to push Microsoft Entra ID users and groups from Microsoft Entra ID to vCenter Server.

3. Return to the VMware knowledge base article at <https://kb.vmware.com/s/article/94182> to update the Microsoft Entra ID Redirect URI.

Follow the steps in the section titled "Update the Azure AD Redirect URI."

4. To create the SCIM 2.0 application, remain in the VMware knowledge base article at <https://kb.vmware.com/s/article/94182>.

Follow the steps in the section titled "Create the SCIM 2.0 Application and Push Users and Groups to vCenter Server." When done creating the SCIM 2.0 application as described in the knowledge base article, continue with the next step.

5. Configure the group membership in vCenter Server for Microsoft Entra ID Authorization.
You must configure group membership before Microsoft Entra ID users can log in to vCenter Server.
 - a) In the vSphere Client, while logged in as a local administrator, go to **Administration > Single Sign On > Users and Groups**.
 - b) Click the **Groups** tab.
 - c) Click the **Administrators** group and click **Add Members**.
 - d) Select the domain name of the Microsoft Entra ID group that you want to add from the drop-down menu.
 - e) In the text box below the drop-down menu, enter the first few characters of Microsoft Entra ID group that you want to add then wait for the drop-down selection to appear.
 - f) Select the Microsoft Entra ID group and add it to the Administrators group.
 - g) Click **Save**.
6. Verify logging in to vCenter Server with a Microsoft Entra ID user.
7. To assign inventory-level and global permissions to Microsoft Entra ID users, see the topic about managing permissions for vCenter Server components in the *vSphere Security* documentation.

Configuring vCenter Server Identity Provider for PingFederate

After installing or upgrading to vSphere 8.0 Update 3, you can configure vCenter Server Identity Provider Federation for PingFederate as an external identity provider.

High-Level Steps to Configure vCenter Server Identity Provider for PingFederate

Configuring vCenter Server for PingFederate involves the following high-level steps:

1. On PingFederate, creating the vCenter Server/VMware Identity Services specific configuration, including the scopes and common configuration for PingFederate workflows.
2. On PingFederate, creating global items, including the Password Grant Flow configuration and Authorization Code Flow configuration.
3. On PingFederate, installing the SCIM Provisioner.
4. On vCenter Server, creating the PingFederate identity provider.
5. On PingFederate, creating the SCIM Application (SP Connection).
6. On vCenter Server, authorizing PingFederate users.

Note: The instructions in this documentation create a typical setup for your PingFederate server. Your environment might differ, and thus you might make different selections.

Prerequisites to Configure vCenter Server Identity Provider for PingFederate

PingFederate requirements:

- You have installed an on-premises PingFederate server.
- From the vCenter Server on which you configure the PingFederate identity provider, you must obtain the trusted root certificates and import them into the PingFederate server.
- Optionally, you might need to import the PingFederate SSL certificate, or certificate chain, to the vCenter Server, if that certificate is self-signed (that is, not issued by a well-known, public Certificate Authority). If the PingFederate SSL certificate or one of the certificates in the chain was issued by a well-known Certificate Authority, vCenter Server automatically trusts that certificate, and you do not need to import it. If you are using one or more intermediate signing authorities for your PingFederate server SSL certificate, include the entire chain of certificates.
To export the PingFederate SSL certificate, in the PingFederate Admin Console, go to **Security > SSL Server Certificates**, select the default certificate, and select **Export** from the **Select Action** drop-down.

You import the PingFederate SSL certificate using the vSphere Client in the **OpenID Connect** panel as part of the configure identity provider workflow.

- To perform OIDC logins and manage user and group permissions, you must create the following PingFederate applications.
 - A PingFederate native application with OpenID Connect as the sign-on method. The native application must include the grant types of authorization code, refresh token, and resource owner password.
 - A System for Cross-domain Identity Management (SCIM) 2.0 application (in PingFederate, this is called the SP Connection), with an OAuth 2.0 Bearer Token to perform user and group synchronization between the PingFederate server and the vCenter Server.
- You have identified the PingFederate users and groups that you want to share with vCenter Server. This sharing is a SCIM operation (not an OIDC operation).

PingFederate connectivity requirements:

- vCenter Server must be able to connect to the PingFederate discovery endpoint, and the authorization, token, JWKS, and any other endpoints advertised in the discovery endpoint metadata.
- PingFederate must also be able to connect with vCenter Server to send user and group data for the SCIM provisioning.

vCenter Server requirements:

- vSphere 8.0 Update 3
- On the vCenter Server where you want to create the PingFederate identity source, verify that the VMware Identity Services are activated.

Note: When you install or upgrade to vSphere 8.0 Update 1 or later, VMware Identity Servers are activated by default. You can use the vCenter Server Management Interface to confirm the status of the VMware Identity Services. See [Stop and Start the VMware Identity Services](#).

vSphere privileges requirements:

- You must have the **VcIdentityProviders > Manage** privilege to create, update, or delete a vCenter Server Identity Provider that is required for federated authentication. To limit a user to view the Identity Provider configuration information only, assign the **VcIdentityProviders > Read** privilege.

Enhanced Linked Mode requirements:

- You can configure vCenter Server Identity Provider Federation for PingFederate in an Enhanced Linked Mode configuration. When you configure PingFederate in an Enhanced Link Mode configuration, you configure the PingFederate identity provider to use VMware Identity Services on a single vCenter Server system. For example, if your Enhanced Mode Link configuration consists of two vCenter Server systems, only one vCenter Server and its instance of VMware Identity Services is used to communicate with the PingFederate server. If this vCenter Server system becomes unavailable, you can configure VMware Identity Services on other vCenter Server in the ELM configuration to interact with your PingFederate server. For more information, see [Activation Process for External Identity Providers in Enhanced Linked Mode Configurations](#).
- When configuring PingFederate as an external identity provider, all the vCenter Server systems in an Enhanced Linked Mode configuration must run at least vSphere 8.0 Update 3.

Create the Scopes

PingFederate supports the use of scopes to constrain and define access privileges.

Review the [Prerequisites to Configure vCenter Server Identity Provider for PingFederate](#).

Log in to the PingFederate Admin console with an Administrator Account.

1. Go to **System > OAuth Settings > Scope Management**.
2. On the **Common Scopes** tab, add the following **Scope Values** along with a description. Click **Add** after entering each value and description.
 - **openid**
 - **profile**
 - **email**
3. Skip the **Exclusive Scopes** tab.
4. On the **Default Scope** tab, enter a description for **Default Scope**.
A description is required. If **Default Scope Description** is empty, PingFederate logs the following error:
The requested scope is invalid, unknown, malformed, or exceeds that which the client is permitted to request.
5. Click **Save**.

Continue with [Create a Common Configuration for PingFederate Workflows](#).

Create a Common Configuration for PingFederate Workflows

Creating the common configuration for PingFederate includes creating the Access Token Manager, the objectID Attribute, the OpenID Connect Policy, and the OAuth Client Application.

Complete the following task:

- [Create the Scopes](#)

Log in to the PingFederate Admin console with an Administrator Account.

1. Create the Access Token Manager.
 - a) Go to **Applications > OAuth > Access Token Management**.
 - b) Click **Create New Instance**.
 - c) On the **Type** tab:
 - **Instance Name**: Enter an instance name. For example, vIDB Access Token Manager.
 - **Instance ID**: Enter the instance ID. For example, vIDB.
 - **Type**: Select **JSON Web Tokens**.
 - **Parent Instance**: Leave the default, **None**.
 - d) On the **Instance Configuration** tab:
 - **Use Centralized Signing Key**: Select the checkbox.
Leaving this checkbox unselected causes PingFederate to expect "Active Signing Certificate Key ID" to be configured.
 - **JWS Algorithm**: Select an algorithm. For example, **RSA using SHA-256**.
 - At the bottom of the screen, click **Show Advanced Fields**.
 - **JWT ID Claim Length**: Add a number greater than zero (0). For example, 24. If you do not enter a value, the JTI claim is omitted in the access token.

- e) Click **Next**.
 - f) On the **Access Token Attribute Contract** tab:
 - In the **Extend the Contract** text box, add the following claims to be generated in the Ping access token. Click **Add** after entering each claim.
 - **aud**
 - **iss**
 - **exp**
 - **iat**
 - **userName**
 - **Subject Attribute Name:** Select one claim to be used for auditing purposes. For example, **iss**.
 - g) Click **Next** twice to skip the **Resource URIs** and **Access Control** tabs.
 - h) Click **Save**.
2. Add the objectGUID Attribute.
- a) Go to **System > Data Stores > Your Data Store > LDAP Configuration**.
 - b) On the **LDAP Configuration** tab, click **Advanced** at the bottom.
 - c) On the **LDAP Binary Attributes** tab, in the **Binary Attribute** name field, use **objectGUID** and click **Add**.
 - d) Click **Save**.
3. Create the OpenID Connect Policy.
- a) Go to **Applications > OAuth > OpenID Connect Policy Management**.
 - b) Click **Add Policy**.
 - c) On the **Manage Policy** tab:
 - **Policy ID:** Enter a policy ID. For example, OIDC.
 - **Name:** Enter a policy name. For example, OIDC Policy.
 - **Access Token Manager:** Select the access token manager that you previously created. For example, vIDB Access Token Manager.
 - d) Click **Next**.
 - e) On the **Attribute Contract** tab:
 - Click **Delete** to remove all the attributes except **sub**. Otherwise, you must map the attributes to a value in the **Contract Fulfillment** tab later.
 - f) Click **Next**, then click **Next** again to skip the **Attributes Scope** tab.
 - g) On the **Attribute Sources & User Lookup** tab, click **Add Attribute Source**.
After entering the information on each tab that follows, click **Next** to advance.
 - **Data Store:**
 - **Attribute Source ID:** Enter an attribute source ID. For example, vIDBLDAP.
 - **Attribute Source Description:** Enter a description. For example: vIDBLDAP.
 - **Active Data Store:** Select your Active Directory or OpenLDAP domain name from the drop-down.
 - **LDAP Directory Search:**
 - **Base DN:** Enter your base DN to find your users and groups.
 - **Search Scope:** Leave the default, **Subtree**.
 - **Attributes to return from search:** Select **<Show All Attributes>** and select **objectGUID**.
Click **Add Attribute**.
 - **LDAP Binary Attribute Encoding Types:**

- **ObjectGUID**: Select **Hex** for the **Attribute Encoding Type**.
- **LDAP Filter**:
 - **Filter**: Enter a filter. For example, `userPrincipalName=${userName}`.
- h) On the **Summary** page, click **Done**.
- i) Click **Next** to advance, and on the **Contract Fulfillment** tab, map the **Attribute Contract** for the ID token:

Attribute Contract	Source	Value
sub	Select the Attribute Source ID previously created. In this documentation, the example used is vIDBLDAP.	objectGUID

- j) Click **Next**, then click **Next** again to skip the **Insurance Criteria** tab.
 - k) Click **Save**.
4. Create the OAuth Client Application.
- a) Go to **Applications > OAuth > Clients**.
 - b) Click **Add Client**.
 - c) On the **Clients | Client** page:
 - **Client ID**: Enter the client ID. For example, vIDB.
Note: Copy and save the client ID for use later when creating the vCenter Server identity provider for PingFederate.
 - **Name**: Enter a name. For example, vIDB.
 - **Client Authentication**: Select **Client Secret**.
 - **Client Secret**: You can input your own client secret or generate a secret. After you leave this page, you cannot view the secret again. You only have the option to change the secret.
Note: Copy and save the secret for use later when creating the vCenter Server identity provider.
 - **Redirection URIs**: Enter the redirection URI(s) in the form: `https://vCenter_Server_FQDN:port/federation/t/CUSTOMER/auth/response/oauth2`.
 - Click **Add**.
 - **Allowed Grant Types**: Check **Authorization Code**, **Refresh Token**, **Client Credentials**, and **Resource Owner Password Credentials**.
 - **Default Access Token Manager**: Select the access token manager you previously created. For example, the one used in this documentation is vIDB Access Token Manager.
 - **OpenID Connect**: For **Policy**, select the one you previously created. For example, the one used in this documentation is OIDC.
 - d) Click **Save**.

Continue with [Create the Password Grant Flow Configuration](#).

Create the Password Grant Flow Configuration

For PingFederate to authenticate with vCenter Server, you set up the password grant flow.

Complete the following tasks:

- [Create the Scopes](#)
- [Create a Common Configuration for PingFederate Workflows](#)

Log in to the PingFederate Admin console with an Administrator Account.

1. Create the Password Credential Validator.
 - a) Go to **System > Data & Credential Stores > Password Credential Validators**.
 - b) Click **Create New Instance**.
 - c) In the **Password Credential Validators | Create New Instance** page, enter the information as follows for each tab, then click **Next** to advance.
 - On the **Type** tab:
 - **Instance Name**: Enter the instance name. For example, vIDB Validator.
 - **Instance ID**: Enter the instance ID. For example, vIDB.
 - **Type**: Select **LDAP Username Password Credential Validator**.
 - On the **Instance Configuration** tab:
 - **LDAP Datastore**: Select the data store you are using.
 - **Search Base**: Enter your base DN to find your users and groups.
 - **Search Filter**: Enter a filter. For example, `userPrincipalName=${username}`.
 - **Scope of Search**: Select **Subtree**.
 - On the **Extended Contract** tab:
 - By default, the following are added:
 - **DN**
 - **mail**
 - **givenName**
 - **username**
 - d) Click **Next** then click **Save**.
2. Map the Validator in the Authorization Server Settings.
 - a) Go to **System > OAuth Settings > Authorization Server Settings**.
 - b) In **Password Credentials Validator**, select the one you previously created. For example, this documentation uses vIDB Validator.
 - c) Click **Save**.
3. Create the Resource Owner Credentials Grant Mapping.
 - a) Go to **Authentication > OAuth > Resource Owner Credentials Grant Mapping**.
 - b) In the **Resource Owner Credentials Grant Mapping** window:
 - **Source Password Validator Instance**: Select the one you previously created and click **Add Mapping**.
 - c) In the **Resource Owner Credentials Grant Mapping | Resource Owner Credentials Mapping** page, click **Next** to skip the **Attribute Sources & User Lookup** tab.
 - d) On the **Contract Fulfillment** tab:
 - For **USER_KEY**, select **Password Credentials Validator**, and for **Value**, select **username**.
 - e) Click **Next** to skip the **Insurance Criteria** tab then click **Save**.
4. Create the Access Token Mapping - Map the Password Credentials Validator to the Access Token Manager.

This mapping is needed for the Password Grant workflow. If the mapping does not exist, PingFederate logs the following error:

There are no access token managers available for the selected client and authentication context.

 - a) Go to **Applications > OAuth > Access Token Mapping**.
 - **Context**: Select the one you previously created. For example, this documentation uses vIDB Validator.

- **Access Token Manager:** Select the one previously created. For example, this documentation uses vIDB Access Token Manager.
- b) Click **Add Mapping**.
- c) Click **Next** to skip the **Attribute Sources & User Lookup** tab.
- d) On the **Contract Fulfilment** tab, use the following table.

Contract	Source	Value
aud	Context	The client ID previously created. For example, the ID used in this documentation is vIDB.
exp	No Mapping	-
iat	Expression	Enter the following: @org.jose4j.jwt.NumericDate@now().getValue()
iss	Expression (If not visible, see the PingFederate documentation at https://docs.pingidentity.com/r/en-us/pingfederate-120/pf_enable_disable_express .)	Enter the following: #tmp=#this.get("context.HttpRequest").getObjectValue().getRequestURL().toString(), #url=new java.net.URL(#tmp), #protocol=#url.getProtocol(), #host=#url.getHost(), #port=#url.getPort(), #result=(#port != -1) ? @java.lang.String@format("%s://%s:%d", #protocol, #host, #port) : @java.lang.String@format("%s://%s", #protocol, #host, #port)
userName	No Mapping	- This contract is used later in the LDAP Filter for the OIDC Policy for Authorization Code workflow. For the PingFederate workflow, it is not needed.

- e) Click **Next** to skip the **Insurance Criteria** tab then click **Save**.

Continue with [Create the Authorization Code Flow Configuration](#).

Create the Authorization Code Flow Configuration

Creating the authorization code flow in PingFederate involves creating and configuring an IdP adapter.

Complete the following tasks:

- [Create the Scopes](#)
- [Create a Common Configuration for PingFederate Workflows](#)
- [Create the Password Grant Flow Configuration](#)

Log in to the PingFederate Admin console with an Administrator Account.

1. Create the IdP Adapter.

- a) Go to **Authentication > Integration > IdP Adapters**.
- b) Create **Create New Instance**.
- c) On the **Type** tab:
 - **Instance Name**: Enter a name, for example, HTML Form Auth Adapter.
 - **Instance ID**: Enter an ID, for example, HTMLFormAuthAdapter.
 - **Type**: Select HTML Form IdP Adapter.
 - **Parent Instance**: Select **None**.
- d) Click **Next**.
- e) On the **IdP Adapter** tab:

Under **Password Credential Validator Instance**, click **Add a new row to 'Credential Validators.'** then select a validator (in this documentation, vIDB Validator is used), and click **Update**.
- f) Click **Next**.
- g) Click **Next** to skip the **Extended Contract** tab.
- h) On the **Adapter Attributes** tab:
 - **Unique User Key Attribute**: Select **username** and check **Pseudonym**.
- i) Click **Next** to skip the **Adapter Contract Mapping** tab then click **Save**.

2. Create IdP Adapter Grant Mapping.

- a) Go to **Authentication > OAuth > IdP Adapter Grant Mapping**.
- b) **Source Adapter Instance**: Select the Adapter instance you just created and click **Add Mapping**.
- c) On the **Attribute Sources & User Lookup** page, click **Add Attribute Source**.
- d) Enter the information as follows for each tab, then click **Next** to advance.
 - On the **Data Store** tab:
 - **Attribute Source ID**: Enter an ID with alphanumeric values.
 - **Attribute Source Description**: Enter a description.
 - **Active Data Store**: Select the active directory in use.
 - On the **LDAP Directory Search** tab:
 - **Base DN**: Enter your base DN to find your users and groups.
 - **Search Scope**: Use the default, **Subtree**.
 - **Attributes to return from search**: Select **<Show All Attributes>**, then, once loaded, from the attribute list, select **userPrincipalName**.
- e) Click **Add Attribute** then click **Next**.
- f) On the **LDAP Filter** tab:
 - **Filter**: Enter the filter. For example, `userPrincipalName=${username}`.
- g) Click **Next** then click **Save**.
- h) In the **IdP Adapter Grant Mapping | IdP Adapter Mapping** page, finish creating the IdP Grant Mapping.

On the **Contract Fulfillment Lookup** tab, use the following table.

Contract	Source	Value
USER_KEY	Select the source previously created.	Subject DN

Contract	Source	Value
USER_NAME	Select the source previously created.	userPrincipalName

- i) Click **Next** then click **Save**.

The created IdP Adapter Grant Mapping displays as '*Adapter Name*' to **Persistent Grant Contract**.

3. Map the IdP Adapter to Access Token Manager.

- a) Go to **Applications > OAuth > Access Token Mappings**.

- **Context:** Select **IdP Adapter:Adapter Name**.
- **Access Token Manager:** Select the access token manager instance previously created. For example, in this documentation, it is vIDB Access Token Manager.

- b) Click **Add Mapping**.

If you do not perform this mapping, PingFederate generates the following log file message:

There are no mapped authentication sources to choose from. Please map an IdP Adapter or IdP connection first.

- c) Skip the **Attribute Sources & User Lookup** tab, and on the **Contract Fulfillment** tab, use the following table.

Contract	Source	Value
aud	No Mapping	-
exp	No Mapping	-
iat	No Mapping	-
iss	No Mapping	-
userName	Adapter	username

- d) Click **Next** to skip the **Insurance Criteria** tab then click **Save**.

Continue with [Install the SCIM Provisioner](#).

Install the SCIM Provisioner

You create a System for Cross-domain Identity Management (SCIM) application that uses a token to synchronize the PingFederate users and groups into VMware Identity Services.

Complete the following tasks:

- [Create the Scopes](#)
- [Create a Common Configuration for PingFederate Workflows](#)
- [Create the Password Grant Flow Configuration](#)
- [Create the Authorization Code Flow Configuration](#)

On the PingFederate server, you must install the SCIM Provisioner to enable the provisioning of users and groups using SCIM.

Note: If you are using an existing PingFederate environment, you might have already installed the SCIM Provisioner.

1. Download the SCIM Provisioner from <https://support.pingidentity.com/s/marketplace-integration/a7i1W0000004IDNQA2/scim-provisioner>.

You must log in to the PingIdentity portal.

2. Copy the `pf-scim-quickconnection-1.4.jar` file to the folder that is mounted to the `/opt/out` folder of your PingFederate server.
For example, place the file in the `/opt/out/instance/server/default/deploy` folder.

3. View the `/opt/out/instance/bin/run.properties` file and ensure that this setting is present:

```
pf.provisioner.mode=STANDALONE
```

According to the PingFederate documentation:

STANDALONE - This server is a standalone instance that runs both the UI console and protocol engine (default).

4. If your PingFederate server instance is running as a container image, and you updated the `run.properties` file, you might need to restart the server. For example:
 - a) Connect to the PingFederate server using SSH.
 - b) Change to the `/root/ping` directory.
 - c) Run the following commands:

```
docker-compose down
docker-compose up
```

The SCIM Connector displays as an option when configuring the user provisioning in [Create the SCIM Application \(SP Connection\)](#).

Continue with [Configure vCenter Server Identity Provider Federation for PingFederate](#).

Configure vCenter Server Identity Provider Federation for PingFederate

After installing or upgrading to vSphere 8.0 Update 3, you can configure vCenter Server Identity Provider Federation for PingFederate as an external identity provider.

Complete the following tasks:

- [Create the Scopes](#)
- [Create a Common Configuration for PingFederate Workflows](#)
- [Create the Password Grant Flow Configuration](#)
- [Create the Authorization Code Flow Configuration](#)
- [Install the SCIM Provisioner](#)

Ensure that you have the following information from the PingFederate OpenID Connect application:

- Client Identifier
- Client secret (shown as Shared secret in the vSphere Client)
- Active Directory domain information, or PingFederate domain information if you are not running Active Directory

vCenter Server supports only one configured external identity provider (one source), and the `vsphere.local` identity source (local source). You cannot use multiple external identity providers. vCenter Server Identity Provider Federation uses OpenID Connect (OIDC) for user login to vCenter Server.

You can configure privileges using PingFederate groups and users through global or object permissions in vCenter Server. See the *vSphere Security* documentation for details about adding permissions.

1. To create the identity provider on vCenter Server:
 - a) Use the vSphere Client to log in as an administrator to vCenter Server.
 - b) Go to **Home > Administration > Single Sign On > Configuration**.
 - c) Click **Change Provider** and select **PingFederate**.
The **Configure Main Identity Provider** wizard opens.
 - d) In the **Prerequisites** panel, review the PingFederate and the vCenter Server and other requirements.
 - e) Click **Run Prechecks**.

If the precheck finds errors, click **View Details** and take steps to resolve the errors as indicated.

- f) When the Precheck passes, click the confirmation checkbox then click **Next**.
 - g) In the **Directory Information** panel, enter the following information.
 - Directory Name: Name of the local directory to create on vCenter Server that stores the users and groups pushed from PingFederate. For example, `vcenter-PingFederate-directory`.
 - Domain Name(s): Enter the PingFederate domain names that contain the PingFederate users and groups you want to synchronize with vCenter Server.
After you enter your PingFederate domain name, click the Plus icon (+) to add it. If you enter multiple domain names, specify the default domain.
 - h) Click **Next**.
 - i) In the **OpenID Connect** panel, enter the following information.
 - Redirect UI: Filled in automatically. This redirect UI must match what you use in creating the OpenID Connect application in PingFederate.
 - Identity Provider Name: Filled in automatically as PingFederate.
 - Client Identifier: Obtained when you created the OpenID Connect application. (PingFederate refers to Client Identifier as the Client ID.)
 - Shared Secret: Obtained when you created the OpenID Connect application in PingFederate. (PingFederate refers to Shared Secret as the Client Secret.)
 - OpenID Address: Takes the form `https://PingFederate_domain_space/idp/.well-known/openid-configuration`.
For example, if your PingFederate domain space is `example.PingFederate.com`, then the OpenID Address is: <https://example.PingFederate.com:9031/.well-known/openid-configuration>
 - SSL Certificate: Optionally, browse for the PingFederate SSL certificate, or certificate chain, if this certificate was not issued by a well-known, public Certificate Authority, to upload to vCenter Server. To export the PingFederate SSL certificate, in the Admin Console, go to **Security > SSL Server Certificates**, select the default certificate, and select **Export** from the **Select Action** drop-down. For more information, see the article, Exporting a certificate, at <https://docs.pingidentity.com/r/en-us/pingfederate-111/nfv1585678806463>. You can export the PingFederate SSL certificate without the private key, as it is not needed for the vCenter Server configuration.
 - j) Click **Next**.
 - k) Review the information and click **Finish**.
vCenter Server creates the PingFederate identity provider and displays the configuration information.
2. Under **User Provisioning**, click **Generate** to create the secret token, select the token lifespan from the drop-down, then click **Copy to Clipboard**. Save the token to a secure location.
When you create the PingFederate SP Connection (SCIM application), you use the token to synchronize the PingFederate users and groups into VMware Identity Services.

Continue with [Create the SCIM Application \(SP Connection\)](#).

Create the SCIM Application (SP Connection)

Creating a System for Cross-domain Identity Management (SCIM) 2.0 application is required, so that you can specify which PingFederate users and groups to push to vCenter Server.

Complete the following tasks:

- [Create the Scopes](#)
- [Create a Common Configuration for PingFederate Workflows](#)
- [Create the Password Grant Flow Configuration](#)
- [Create the Authorization Code Flow Configuration](#)
- [Install the SCIM Provisioner](#)

- [Configure vCenter Server Identity Provider Federation for PingFederate](#)

1. Add the vCenter Server Trusted Root Certificate to the PingFederate Server.

Before you begin, export the trusted root certificate(s) from vCenter Server. You can obtain the certificate from the file system of the vCenter Server at `/var/lib/vmware/vmca/root.cer`. Or, see the Knowledge Base article at <https://kb.vmware.com/s/article/2108294>.

- a) Log in to the PingFederate Admin console with an Administrator Account.
- b) Go to **Security > Certificate & Key Management**.
- c) Select **Trusted CAs** then click **Import** to add the SSL certificate of the vCenter Server.
- d) If your PingFederate server instance is running as a container image, you might need to restart the server to add the certificate to the trust store. For example:
 1. Connect to the PingFederate server using SSH.
 2. Change to the `/root/ping` directory.
 3. Run the following commands:

```
docker-compose down
docker-compose up
```

2. Create the SP Connection.

- a) Log in to the PingFederate Admin console with an Administrator Account.
- b) Go to **Applications > Integration > SP Connections**.
- c) Click **Create Connection**.
- d) Select **Use a template for this connection** then select **SCIM Connector** from the drop-down.
If the SCIM Connector option does not appear in the drop-down, check that you placed the SCIM Connector .jar file in the correct folder (the /opt/out folder of your PingFederate server).
- e) Click **Next**.
- f) Select only **Outbound Provisioning** then click **Next**.
- g) On the **General Info** tab:
 - **Partner's Entity ID (Connection ID)**: Update **SCIM Connector** to a name of your choice.
 - **Connection Name**: Enter a name.
 - **Base URL**: Enter the HTTPS address of the vCenter Server where you are configuring the PingFederate external identity provider, for example: `https://vcenter1.example.com`.
- h) Click **Next**.
- i) Click **Configure Provisioning**.
On the **Target** tab:
 - **SCIM URL**: Enter the Usergroup endpoint.
This is the Tenant URL obtained under **User Provisioning** on the **Configuration** page of the vCenter Server.
For example: `https://vcenter1.example.com/usergroup/t/CUSTOMER/scim/v2`
 - **Authentication Method**: Select **OAuth 2 Bearer Token** from the drop-down.
 - **Access Token**: Paste the Secret Token that was generated from vCenter Server, and that you should have saved previously. See Step 2 in [Configure vCenter Server Identity Provider Federation for PingFederate](#).
 - **Unique User Identifier**: Select **userName** from the drop-down.
 - **Filter Expression**: Copy the following expression into the text box: `externalId eq "%s"`
- j) Accept the rest of the default configuration setting values and click **Next**.
 - **Provisioning Options**: **User Create**, **User Update**, and **User Disable/ Delete** are checked.
 - **Remove User Action**: **Disable** is selected.
Note: With **Disable** selected, when users are deleted from the Active Directory, they are not automatically shown as "disabled" in VMware Identity Services. This is expected behavior.
 - In Active Directory, instead of deleting the user on the next provisioning cycle, the user's property becomes "active"="false".
 - The user does not show as "disabled" in VMware Identity Services until another user is created or updated in Active Directory on the next provisioning cycle, and you follow the workaround at <https://>

support.pingidentity.com/s/article/After-deleting-an-AD-user-account-SaaS-provisioner-does-not-remove-the-user-in-the-next-provisioning-cycle-when-Group-DN-is-specified.

- **Group Name Source:** **Common name** is selected.
- k) On the **Manage Channels** tab, click **Create**.
- On the **Channel Info** tab:
 - **Channel Name:** Enter a name.
 - Accept the **Max threads** and **Timeout (Secs)** default values.
- l) Click **Next**.
- On the **Source** tab:
 - **Active Data Store:** Choose your Active Directory domain.
- m) Click **Next**.
- On the **Source Location** tab:
 - **Base DN:** Enter your base DN to find your users and groups.
 - **Users:** Customize to your environment. For example:
 - **Group DN:** Do not use.
 - **Filter:** Enter `(| (objectClass=person) (objectClass=organizationalPerson) (objectClass=user) )`.
 - **Groups:** Customize to your environment. For example:
 - **Group DN:** Do not use.
 - **Filter:** Enter `(objectClass=group)`.
- n) Click **Next**.
- o) Accept the defaults on the **Attribute Mapping** tab.
- p) Click **Next**.
- On the **Activation & Summary** tab:
- **Channel Status:** Select **Active**.
- q) Click **Done**.
- The SP Connection is created and the SP Connections screen is displayed.
- r) Click **Done**.
- s) On the **Outbound Provisioning** tab, click **Next**.
- t) Review the summary then click **Save**.
- u) To make the connection active, toggle the **Enabled** slider.

PingFederate now pushes users and groups from the configured data store to vCenter Server. Allow some time for the push to occur. You can view the pushed users and groups in the vSphere Client. Go to **Administration > Single Sign On > Users and Groups**, and select the PingFederate domain.

Continue with [Configure vCenter Server for PingFederate Authorization](#).

Configure vCenter Server for PingFederate Authorization

You can either assign PingFederate users to a vCenter Server group or assign inventory-level and global permissions to PingFederate users.

Complete the following tasks:

- [Create the Scopes](#)
- [Create a Common Configuration for PingFederate Workflows](#)
- [Create the Password Grant Flow Configuration](#)
- [Create the Authorization Code Flow Configuration](#)

- [Install the SCIM Provisioner](#)
- [Configure vCenter Server Identity Provider Federation for PingFederate](#)
- [Create the SCIM Application \(SP Connection\)](#)

The minimum permission required for a PingFederate user to log in is Read-Only.

1. To assign PingFederate users to a group, see [Add Members to a vCenter Single Sign-On Group](#).
2. To assign inventory-level and global permissions to PingFederate users, see the topic about managing permissions for vCenter Server components in the *vSphere Security* documentation.
3. After assigning a PingFederate user permissions, verify that the user can log in.

Configure VMware Single Sign-On

After installing or upgrading to vSphere 8.0 Update 3, you can configure vCenter Server hosts for VMware Single Sign-On. When you configure VMware Single Sign-On, you use an external identity provider to sign into your vCenter Server hosts.

VMware Single Sign-On requirements:

- The vCenter Server on which you configure VMware Single Sign-On runs vSphere 8.0 Update 3.
- The vCenter Server hosts you want to connect to run at least vSphere 8.0 Update 1.
- You have configured one of the following external identity providers:
 - Microsoft Entra ID
 - Okta
 - PingFederate
- You must add the trusted root certificate from the vCenter Server host on which the external identity provider is configured to the vCenter Server host on which you configure VMware Single Sign-On.

VMware Single Sign-On enables you to connect vCenter Server hosts in a non-Enhanced Linked Mode configuration. That is, as long as you configure an external identity provider, you can leverage that configuration for single sign-on to other vCenter Server hosts. The vCenter Server host on which the external identity provider is configured acts as the identity provider for the other vCenter Server hosts.

You can configure multiple vCenter Server hosts to perform VMware Single Sign-On. To do so, you must configure each vCenter Server host to point to the vCenter Server host configured with an external identity provider.

After performing the VMware Single Sign-On configuration, you can still log in to your vCenter Server hosts with a local account.

Note: VMware Single Sign-On does not share inventories between vCenter Server hosts as occurs in Enhanced Linked Mode.

1. Download the trusted root certificate from the vCenter Server host on which the external identity provider is configured. For example, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2108294>.
2. Upload that trusted root certificate to the vCenter Server host on which you are configuring VMware SSO. See [Add a Trusted Root Certificate to the Certificate Store Using the vSphere Client](#).
3. Use the vSphere Client to log in as an administrator to the vCenter Server host on which you are configuring VMware SSO.
4. Navigate to **Home > Administration > Single Sign On > Configuration**.
5. Click **Change Provider** and select **VMware SSO**.
The **Configure Main Identity Provider** wizard opens.

6. In the **Prerequisites** panel, review the vCenter Server requirements.
7. Click **Run Prechecks**.
If the precheck finds errors, click **View Details** and take steps to resolve the errors as indicated.
8. When the Precheck passes, click the confirmation checkbox then click **Next**.
9. In the **OpenID Connect** panel, enter the following information.
 - Identity Provider Name: Filled in with VMware SSO.
 - vCenter Server FQDN: Enter the FQDN of the vCenter Server host where the external identity provider is configured.
 - Port Number: Accept the default of 443, or change to the port you want to use.
 - User Name and Password: Enter the user name and password for an administrator account on this vCenter Server host where the external identity provider is configured.
10. Click **Next**.
11. Review the information and click **Finish**.
vCenter Server creates the VMware SSO provider and displays the configuration information. This vCenter Server host now contains the same external identity provider configuration as the host on which the configuration was created. For example, when you compare the OpenID configurations between the two hosts, it is the same.
12. Configure this vCenter Server to use the external identity provider for authorization.
You can either assign the external identity provider users to a vCenter Server group or assign inventory-level and global permissions to the users. The minimum permission required for logging in is Read-Only.
To assign the external identity provider users to a group, see [Add Members to a vCenter Single Sign-On Group](#). To assign inventory-level and global permissions to the users, see the topic about managing permissions for vCenter Server components in the *vSphere Security* documentation.
13. Verify logging in to this vCenter Server host with an external identity provider user.
When you launch the vSphere Client, you see the Welcome to VMware vSphere screen, with the **Sign in with SSO** button. When you click this button, you are redirected to the external identity provider's sign-in screen.

Managing VMware Identity Services

You can stop and start VMware Identity Services, regenerate a SCIM token, and restore deleted SCIM users and groups. Depending on the task, you use either the vSphere Client or the external identity provider's administration console.

Stop and Start the VMware Identity Services

To configure and run Okta, Microsoft Entra ID (formerly called Azure AD), or PingFederate as an external identity provider, VMware Identity Services must be started on vCenter Server. By default, when you either install or upgrade to vSphere 8.0 Update 1 or later, VMware Identity Services are started. You use the vCenter Server Management Interface to manage VMware Identity Services.

When you install or upgrade to vSphere 8.0 Update 1 or later, VMware Identity Services are automatically started. When you configure Okta, Microsoft Entra ID, or PingFederate as an external identity provider, you do not need to start VMware Identity Services, as they are already running. To start or stop VMware Identity Services, you must be root.

You configure the external identity provider on a single vCenter Server only. That vCenter Server, through its instance of VMware Identity Services, communicates to the identity provider. The other vCenter Server systems in the Enhanced

Linked Mode configuration also have VMware Identity Services running, however, they do not communicate directly to the identity provider.

Starting in version 8.0 Update 1, vSphere includes the VMware Identity Services to support authenticating to Okta. Starting in version 8.0 Update 2, VMware Identity Services support authenticating to Microsoft Entra ID. Starting in version 8.0 Update 3, VMware Identity Services support authenticating to PingFederate.

1. In a Web browser, go to the vCenter Server Management Interface at `https://vcenter-IP-address-or-FQDN:5480`.
2. Log in as root.
The default root password is the password that you set while deploying the vCenter Server.
3. Select **Services**.
4. View the status of VMware Identity Services.
5. To stop or start the service, select **VMware Identity Services**, then click either **Stop** or **Start**.
After starting VMware Identity Services, no vCenter Server reboot is necessary.

Regenerate the SCIM Token in vCenter Server

In vCenter Server, you can regenerate a System for Cross-Domain Identity Management (SCIM) token for an external identity provider.

You must have created an external identity provider in vCenter Server.

If you generate another token, it becomes active immediately, and the previous token is revoked.

1. Log in as an administrator with the vSphere Client to the vCenter Server.
2. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
3. On the **Configuration** page, under **User Provisioning/Secret Token**, click **Regenerate** to regenerate the secret token, select the token lifespan from the drop-down, then click **Copy to Clipboard**. Save the token to a secure location.
4. The copied token is available for you to update your external identity provider configuration.

Restore Deleted SCIM Users and Groups

If the SCIM-pushed users and groups on your vCenter Server become out-of-sync with your external identity provider, you can take steps to fix the problem.

When you want to restore an SCIM-pushed user or group that you deleted from your vCenter Server, you cannot simply push the user or group from your identity provider. Because of the way vCenter Server uses System for Cross-domain

Identity Management (SCIM) for user and group management, you must update the SCIM 2.0 application itself with the missing user or group.

1. Log in to the external IDP Admin console.
2. Navigate to the SCIM 2.0 application.
3. Assign the deleted or missing user or group.
4. Select the appropriate action to delete the pushed group or user to unlink the pushed group or user.
5. Select the appropriate action to push the group.
6. Verify on your vCenter Server that external IDP synchronized the group or user.

vCenter Single Sign-On

If you are not using an external identity provider, you must understand the underlying architecture of the built-in identity provider, vCenter Single Sign-On, and how it affects installation and upgrades.

vCenter Single Sign-On Components

vCenter Single Sign-On includes the Security Token Service (STS), an administration server, the vCenter Lookup Service, and the VMware Directory Service (vmdir). The VMware Directory Service is also used for certificate management.

During installation, the following components are deployed as part of a vCenter Server deployment.

STS (Security Token Service)

The STS service issues Security Assertion Markup Language (SAML) tokens. These security tokens represent the identity of a user in one of the identity source types supported by vCenter Server. The SAML tokens allow interactive, scripted, and service users (including solution users) who authenticate successfully to vCenter Single Sign-On to use any vCenter service that vCenter Single Sign-On supports without authenticating again to each service.

The vCenter Single Sign-On service signs all tokens with a signing certificate, and stores the token signing certificate on disk. The certificate for the service itself is also stored on disk.

Administration server

The administration server allows users with administrator privileges to vCenter Single Sign-On to configure the vCenter Single Sign-On server and manage users and groups from the vSphere Client. Initially, only the user `administrator@your_domain_name` has these privileges. You can change the vSphere domain when you install vCenter Server. Do not name the domain name with your Microsoft Active Directory or OpenLDAP domain name.

VMware Directory Service (vmdir)

A VMware Directory Service (vmdir) is associated with the domain you specify during installation and is included in each vCenter Server deployment. This service is a multi-tenanted, peer-replicating directory service that makes an LDAP directory available on port 389. It also stores and manages vCenter Single Sign-On user accounts and passwords, which are secured by the SHA-512 hashing algorithm.

If your environment includes multiple instances of vCenter Server configured in linked mode, an update of vmdir content in one vmdir instance is propagated to all other instances of vmdir.

The VMware Directory Service stores not only vCenter Single Sign-On information but also certificate information.

Identity Management Service

Handles identity sources and STS authentication requests.

Using vCenter Single Sign-On with vSphere

When a user logs in to a vSphere component, or when a vCenter Server solution user accesses another vCenter Server service, vCenter Single Sign-On performs authentication. Users must be authenticated with vCenter Single Sign-On and have the necessary privileges for interacting with vSphere objects.

vCenter Single Sign-On authenticates both solution users and other users.

- Solution users represent a set of services in your vSphere environment. During installation, VMCA assigns a certificate to each solution user by default. The solution user uses that certificate to authenticate to vCenter Single Sign-On. vCenter Single Sign-On gives the solution user a SAML token, and the solution user can then interact with other services in the environment.
- When other users log in to the environment, for example, from the vSphere Client, vCenter Single Sign-On prompts for a user name and password. If vCenter Single Sign-On finds a user with those credentials in the corresponding identity source, it assigns the user a SAML token. The user can now access other services in the environment without being prompted to authenticate again.

Which objects the user can view, and what a user can do, is usually determined by vCenter Server permission settings. vCenter Server administrators assign those permissions from the **Permissions** interface in the vSphere Client, not through vCenter Single Sign-On. See the *vSphere Security* documentation.

vCenter Single Sign-On and vCenter Server Users

Users authenticate to vCenter Single Sign-On by entering their credentials on the login page. After connecting to vCenter Server, authenticated users can view all vCenter Server instances or other vSphere objects for which their role gives them privileges. No further authentication is required.

vCenter Single Sign-On Administrator Users

The vCenter Single Sign-On administrative interface is accessible from the vSphere Client.

To configure vCenter Single Sign-On and manage vCenter Single Sign-On users and groups, the user administrator@vsphere.local or a user in the vCenter Single Sign-On Administrators group must log in to the vSphere Client. Upon authentication, that user can access the vCenter Single Sign-On administration interface from the vSphere Client and manage identity sources and default domains, specify password policies, and perform other administrative tasks.

Groups in the vCenter Single Sign-On Domain

The vCenter Single Sign-On domain (vsphere.local by default) includes several predefined groups. Add users to one of those groups to permit them to perform the corresponding actions.

See [Managing vCenter Single Sign-On Users and Groups](#).

For all objects in the vCenter Server hierarchy, you can assign permissions by pairing a user and a role with the object. For example, you can select a resource pool and give a group of users read privileges to that resource pool object by giving them the corresponding role.

For some services that are not managed by vCenter Server directly, membership in one of the vCenter Single Sign-On groups determines the privileges. For example, a user who is a member of the Administrators group can manage vCenter Single Sign-On. A user who is a member of the CAAadmins group can manage the VMware Certificate Authority, and a user who is in the LicenseService.Administrators group can manage licenses.

The following groups are predefined in vsphere.local. Many of these groups are internal to vsphere.local or give users high-level administrative privileges. Add users to any of these groups only after careful consideration of the risks.

CAUTION: Do not delete any of the predefined groups in the vsphere.local domain. If you do, errors with authentication or certificate provisioning might result.

Table 112: Groups in the vsphere.local Domain

Privilege	Description
Users	Users in the vCenter Single Sign-On domain (vsphere.local by default).
SolutionUsers	Solution users group for vCenter services. Each solution user authenticates individually to vCenter Single Sign-On with a certificate. By default, VMCA provisions solution users with certificates. Do not add members to this group explicitly.
CAAdmins	Members of the CAAdmins group have administrator privileges for VMCA. Do not add members to this group unless you have compelling reasons.
DCAdmins	Members of the DCAdmins group can perform Domain Controller Administrator actions on VMware Directory Service. Note: Do not manage the domain controller directly. Instead, use the <code>vmdir</code> CLI or the vSphere Client to perform corresponding tasks.
SystemConfiguration.BashShellAdministrators	A user in this group has full access to all the Appliance Management APIs. By default, a user who connects to the vCenter Server with SSH can access only commands in the restricted shell, but users in this group have Bash Shell Access over SSH and gain full privileges similar to the root user.
ActAsUsers	Members of Act-As Users are allowed to get Act-As tokens from vCenter Single Sign-On.
ExternalIDPUsers	This internal group is not used by vSphere. VMware vCloud Air requires this group.
SystemConfiguration.Administrators	Members of the SystemConfiguration.Administrators group can view and manage the system configuration in the vCenter Server Management Interface running on port 5480. These users can view services, start and restart services, and troubleshoot services. These users can also access Appliance Management APIs except for those APIs that modify critical system configurations.
DCClients	This group is used internally to allow the management node access to data in VMware Directory Service. Note: Do not modify this group. Any changes might compromise your certificate infrastructure.
ComponentManager.Administrators	Members of the ComponentManager.Administrators group can invoke component manager APIs that register or unregister services, that is, modify services. Membership in this group is not necessary for read access on the services.
LicenseService.Administrators	Members of LicenseService.Administrators have full write access to all licensing-related data and can add, remove, assign, and unassign serial keys for all product assets registered in the licensing service.
Administrators	Administrators of the VMware Directory Service (vmdir). Members of this group can perform vCenter Single Sign-On administration tasks. Do not add members to this group unless you have compelling reasons and understand the consequences.
TrustedAdmins	Members of this group can perform VMware®vSphere Trust Authority™ configuration and administration tasks. By default, this group does not contain any members. You must add a member to this group so that you can perform vSphere Trust Authority tasks.

Privilege	Description
AutoUpdate	This group is used internally for vCenter Cloud Gateway.
SyncUsers	This group is used internally for vCenter Cloud Gateway.
vSphereClientSolutionUsers	This group is used internally for the vSphere Client.
ServiceProviderUsers	Members of this group can manage the vSphere with Tanzu and VMware Cloud on AWS infrastructure.
NsxAdministrators	This group is used for VMware NSX.
WorkloadStorage	Workload storage group.
RegistryAdministrators	Members of this group can manage the registry.
NsxAuditors	This group is used for VMware NSX.
NsxViAdministrators	This group is used for VMware NSX.
SystemConfiguration.SupportUsers	Members of the SystemConfiguration.SupportUsers group can access the support bundle API.
SystemConfiguration.ReadOnly	Members of this group can access vCenter Server Appliance read-only operations under Appliance Management.
VCLSAdmin	Members of this group have administrative privileges for vSphere Cluster Services (VCLS).
AnalyticsService.Administrators	This group is used for the VMware Analytics Service APIs.
vStatsGroup	This group is used for vStats collection.

Configuring vCenter Single Sign-On Identity Sources

When a user logs in with just a user name, vCenter Single Sign-On checks in the default identity source whether that user can authenticate. When a user logs in and includes the domain name in the login screen, vCenter Single Sign-On checks the specified domain if that domain has been added as an identity source. You can add identity sources, remove identity sources, and change the default.

Identity Sources for vCenter Server with vCenter Single Sign-On

You can use identity sources to attach one or more domains to vCenter Single Sign-On. A domain is a repository for users and groups that the vCenter Single Sign-On server can use for user authentication.

Note: At any time, only one default domain exists. If a user from a non-default domain logs in, that user must add the domain name to authenticate successfully. The domain name is in the form:

`DOMAIN\user`

The following identity sources are available.

- Active Directory over LDAP. vCenter Single Sign-On supports multiple Active Directory over LDAP identity sources.
- OpenLDAP versions 2.4 and later. vCenter Single Sign-On supports multiple OpenLDAP identity sources.

Note: A future update to Microsoft Windows will change the default behavior of Active Directory to require strong authentication and encryption. This change will impact how vCenter Server authenticates to Active Directory. If you use Active Directory as your identity source for vCenter Server, you must plan to enable LDAPS. For more information about this Microsoft security update, see <https://portal.msrc.microsoft.com/en-US/security-guidance/advisory/ADV190023> and <https://blogs.vmware.com/vsphere/2020/01/microsoft-ldap-vsphere-channel-binding-signing-adv190023.html>.

Set the Default Domain for vCenter Single Sign-On

Each vCenter Single Sign-On identity source is associated with a domain. vCenter Single Sign-On uses the default domain to authenticate a user who logs in without a domain name. Users who belong to a domain that is not the default domain must include the domain name when they log in.

When a user logs in to a vCenter Server system from the vSphere Client, the login behavior depends on whether the user is in the domain that is set as the default identity source.

- Users who are in the default domain can log in with their user name and password.
- Users who are in a domain that has been added to vCenter Single Sign-On as an identity source but is not the default domain can log in to vCenter Server but must specify the domain in one of the following ways.
 - Including a domain name prefix, for example, MYDOMAIN\user1
 - Including the domain, for example, user1@mydomain.com
- Users who are in a domain that is not a vCenter Single Sign-On identity source cannot log in to vCenter Server. If the domain that you add to vCenter Single Sign-On is part of a domain hierarchy, Active Directory determines whether users of other domains in the hierarchy are authenticated or not.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as administrator@mydomain.

3. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
4. Under the **Identity Provider** tab, click **Identity Sources**, select an identity source, and click **Set as Default**.
5. Click **OK**.

In the domain display, the default domain shows (default) in the Type column.

Add or Edit a vCenter Single Sign-On Identity Source

Users can log in to vCenter Server only if they are in a domain that has been added as a vCenter Single Sign-On identity source. vCenter Single Sign-On administrator users can add identity sources, or change the settings for identity sources that they added.

If you are adding an Active Directory (Integrated Windows Authentication) identity source, the vCenter Server must be in the Active Directory domain. See [Add a vCenter Server to an Active Directory Domain](#).

An identity source can be an Active Directory over LDAP, a native Active Directory (Integrated Windows Authentication) domain, or an OpenLDAP directory service. See [Identity Sources for vCenter Server with vCenter Single Sign-On](#).

Immediately after installation, the vsphere.local domain (or the domain you specified during installation) with the vCenter Single Sign-On internal users is available.

Note:

If you have updated or replaced your Active Directory SSL certificate, you must remove and re-add the identity source in vCenter Server.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as administrator@mydomain.

3. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
4. Click **Add**.

Initially, each user is assigned the No Access role. A vCenter Server administrator must assign the user at least to the Read Only role before the user can log in. See the topic on using roles to assign privileges in the *vSphere Security* documentation.

Active Directory over LDAP and OpenLDAP Server Identity Source Settings

The Active Directory over LDAP identity source is preferred over the Active Directory (Integrated Windows Authentication) option. The OpenLDAP Server identity source is available for environments that use OpenLDAP.

If you are configuring an OpenLDAP identity source, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2064977> for additional requirements.

Important: Groups in AD-over-LDAP identity sources cannot use users in different domains even if you create an additional identity source for each domain.

Groups in LDAP identity sources only recognize those users that exist in the specified user base DN. This can lead to unexpected problems in large Active Directory environments with child domains. For example, consider the following scenario:

1. An Active Directory forest with two child domains, ChildA and ChildB.
2. A vCenter Server configured with two AD-over-LDAP identity sources, one for child domain ChildA and one for child domain ChildB.
3. ChildA contains two users named UserA1 and UserA2.
4. ChildB contains two users named UserB1 and UserB2.

The vCenter Server administrator creates a group in ChildA named TestGroup that contains UserA1, UserA2, UserB1, and UserB2. The vCenter Server administrator grants login (or any) privileges to TestGroup. Unfortunately, UserB1 and UserB2 cannot log in because they reside in a different domain than the group.

As a workaround, perform the following:

1. Create another group named SecondTestGroup in ChildB.
2. Remove UserB1 and UserB2 from TestGroup.
3. Add UserB1 and UserB2 to SecondTestGroup.
4. In vCenter Server, assign SecondTestGroup the same privileges that were granted to TestGroup.

Note: Microsoft Windows changed the default behavior of Active Directory to require strong authentication and encryption. This change impacts how vCenter Server authenticates to Active Directory. If you use Active Directory as your identity source for vCenter Server, you must enable LDAPS. For more information about this Microsoft security update, see <https://portal.msrm.microsoft.com/en-US/security-guidance/advisory/ADV190023> and <https://blogs.vmware.com/vsphere/2020/01/microsoft-ldap-vsphere-channel-binding-signing-adv190023.html>.

Table 113: Active Directory over LDAP and OpenLDAP Server Settings

Option	Description
Name	Name of the identity source.
Base DN for users	Base Distinguished Name for users. Enter the DN from which to start user searches. For example, cn=Users,dc=myCorp,dc=com.

Option	Description
Base DN for groups	The Base Distinguished Name for groups. Enter the DN from which to start group searches. For example, cn=Groups,dc=myCorp,dc=com.
Domain name	The FQDN of the domain.
Domain alias	For Active Directory identity sources, the domain's NetBIOS name. Add the NetBIOS name of the Active Directory domain as an alias of the identity source if you are using SSPI authentications. For OpenLDAP identity sources, the domain name in capital letters is added if you do not specify an alias.
User name	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups. The ID can be in any of these formats: • UPN (user@domain.com) • NetBIOS (DOMAIN\user) • DN (cn=user,cn=Users,dc=domain,dc=com) The user name must be fully-qualified. An entry of "user" does not work.
Password	Password of the user who is specified by Username .
Connect to	Domain controller to connect to. Can be any domain controller in the domain, or specific controllers.
Primary Server URL	Primary domain controller LDAP server for the domain. You can use either the host name or the IP address. Use the format <code>ldap://hostname_or_IPaddress:port</code> or <code>ldaps://hostname_or_IPaddress:port</code> . The port is typically 389 for LDAP connections and 636 for LDAPS connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for LDAP and 3269 for LDAPS. A certificate that establishes trust for the LDAPS endpoint of the Active Directory server is required when you use <code>ldaps://</code> in the primary or the secondary LDAP URL.
Secondary server URL	Address of a secondary domain controller LDAP server that is used when the primary domain controller is unavailable. You can use either the host name or the IP address. For every LDAP operation, vCenter Server always tries the primary domain controller before falling back to the secondary domain controller. This can lead to Active Directory logins taking some time, and even failing, when the primary domain controller is unavailable. Note: When the primary domain controller fails, the secondary domain controller might not take over automatically.
Certificates (for LDAPS)	If you want to use LDAPS with your Active Directory LDAP Server or OpenLDAP Server identity source, click Browse to select a certificate that was exported from the domain controller specified in the LDAPS URL. (Note that the certificate used here is not a root CA certificate.) To export the certificate from Active Directory, consult the Microsoft documentation. You can browse for and select multiple certificates. Tip: When browsing for and selecting multiple certificates, they must be located in the same directory.

Option	Description
	vCenter Server only trusts certificates directly signed by a registered and trusted certificate authority. vCenter Server does not trace a path up to a registered CA certificate and only checks if the certificate is signed by a registered and trusted certificate authority. As long as your certificate is signed by a publicly trusted certificate authority, or is self-signed, no further action is necessary. However, if you create your own internal certificates (that is, you use a private certificate authority), you might need to include those certificates. For example, if your organization uses Microsoft Enterprise Root Certificate Authority to generate the LDAPS certificate, you must also select the Enterprise Root Certificate to add it to vCenter Server. In addition, if you use intermediate certificate authorities between the LDAPS certificate and the Enterprise Root certificate, you must also select those intermediate certificates to add them to vCenter Server.

Active Directory Identity Source Settings

If you select the Active Directory (Integrated Windows Authentication) identity source type, you can use the local machine account as your SPN (Service Principal Name) or specify an SPN explicitly. You can use this option only if the vCenter Single Sign-On server is joined to an Active Directory domain.

Prerequisites for Using an Active Directory (Integrated Windows Authentication) Identity Source

You can set up vCenter Single Sign-On to use an Active Directory (Integrated Windows Authentication) identity source only if that identity source is available. Follow the instructions in the *vCenter Server Configuration* documentation.

Note: Active Directory (Integrated Windows Authentication) always uses the root of the Active Directory domain forest. To configure your Integrated Windows Authentication identity source with a child domain within your Active Directory forest, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2070433>.

Select **Use machine account** to speed up configuration. If you expect to rename the local machine on which vCenter Single Sign-On runs, specifying an SPN explicitly is preferable.

If you have enabled diagnostic event logging in your Active Directory to identify where hardening might be needed, you might see a log event with Event ID 2889 on that directory server. Event ID 2889 is generated as an anomaly rather than a security risk when using Integrated Windows Authentication. For more information about Event ID 2889, see the VMware knowledge base article at <https://kb.vmware.com/s/article/78644>.

Table 114: Add Identity Source Settings

Text Box	Description
Domain name	FQDN of the domain name, for example, mydomain.com. Do not provide an IP address. This domain name must be DNS-resolvable by the vCenter Server system.
Use machine account	Select this option to use the local machine account as the SPN. When you select this option, you specify only the domain name. Do not select this option if you expect to rename this machine.
Use Service Principal Name (SPN)	Select this option if you expect to rename the local machine. You must specify an SPN, a user who can authenticate with the identity source, and a password for the user.

Text Box	Description
Service Principal Name (SPN)	SPN that helps Kerberos to identify the Active Directory service. Include the domain in the name, for example, STS/example.com. The SPN must be unique across the domain. Running the <code>setspn -S</code> command checks that no duplicate is created. See the Microsoft documentation for information on <code>setspn</code> .
User Principal Name (UPN) Password	Name and password of a user who can authenticate with this identity source. Use the email address format, for example, jchin@mydomain.com. You can verify the User Principal Name with the Active Directory Service Interfaces Editor (ADSI Edit).

Add or Remove an Identity Source Using the CLI

You can use the `sso-config` utility to add or remove an identity source.

If you are adding an Active Directory identity source, the vCenter Server must be in the Active Directory domain. See [Add a vCenter Server to an Active Directory Domain](#).

Enable SSH login. See [Manage vCenter Server Using the vCenter Server Shell](#).

An identity source can be a native Active Directory (Integrated Windows Authentication) domain, AD over LDAP, AD over LDAP using LDAPS (LDAP over SSL), or OpenLDAP. See [Identity Sources for vCenter Server with vCenter Single Sign-On](#). You also use the `sso-config` utility to set up smart card and RSA SecurID authentication.

1. Use SSH or another remote console connection to start a session on the vCenter Server system.
2. Log in as root.
3. Change to the directory where the `sso-config` utility is located.

```
cd /opt/vmware/bin
```

4. Refer to the `sso-config` help by running `sso-config.sh -help`, or see the VMware knowledge base article at <https://kb.vmware.com/s/article/67304> for usage examples.

Managing the vCenter Server Security Token Service

The vCenter Server Security Token Service (STS) is a Web service that issues, validates, and renews security tokens.

As a token issuer, the Security Token Service (STS) uses a private key to sign the tokens and publishes the public certificates for services to verify the token signature. vCenter Server manages the STS signing certificates and stores them in the VMware Directory Service (vmdir). Tokens can have a significant lifetime, and historically might have been signed by any one of multiple keys.

Users present their primary credentials to the STS interface to acquire tokens. The primary credential depends on the type of user.

Table 115: STS Users and Credentials

Type of User	Primary Credentials
Solution user	Valid certificate.
Other users	User name and password available in a vCenter Single Sign-On identity source.

STS authenticates the user based on the primary credentials, and constructs a SAML token that contains user attributes.

By default, the VMware Certificate Authority (VMCA) generates the STS signing certificate. You can refresh the STS signing certificate with a new VMCA certificate. You can also import and replace the default STS signing certificate with a custom or third-party generated STS signing certificate. Do not replace the STS signing certificate unless the security policy of your company requires replacing all certificates.

You can use the vSphere Client to:

- Refresh STS certificates
- Import and replace custom and third-party generated STS certificates
- View STS certificate details, such as the expiration date

You can also use the command line to replace custom and third-party generated STS certificates.

STS Certificate Duration and Expiration

A fresh installation of vSphere 7.0 Update 1 and later creates an STS signing certificate with a duration of 10 years. When an STS signing certificate is close to expiring, an alarm warns you starting at 90 days once per week, and then daily when seven days away.

Note: In certain circumstances, replacing your STS signing certificates can change the duration of the certificates. When performing certificate replacement, pay attention to the issuing and expiration dates.

STS Certificate Auto-Renewal

In vSphere 8.0 and later, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. If the auto-renewal fails, vCenter Single Sign-On creates an error message in the log file. If necessary, you can refresh the STS signing certificate manually.

Note: vCenter Single Sign-On does not perform auto-renewal of custom generated or third-party STS signing certificates.

Refreshing and Importing and Replacing STS Certificates

In vSphere 8.0 and later, refreshing or importing and replacing the STS signing certificates does not require a vCenter Server restart and so avoids any downtime. Also, in a linked configuration, refreshing or importing and replacing the STS signing certificates on a single vCenter Server updates the STS certificates on all the linked vCenter Server systems.

Note: In certain circumstances, a refresh or import and replace of STS signing certificates can require you to manually restart the vCenter Server systems.

Refresh a vCenter Server STS Certificate Using the vSphere Client

You can refresh your vCenter Server STS signing certificates using the vSphere Client. The VMware Certificate Authority (VMCA) issues a new certificate and replaces the current certificate.

For certificate management, you must supply the password of the administrator of the local domain (administrator@vsphere.local by default). If you are renewing certificates, you must also supply the vCenter Single Sign-On credentials for a user with administrator privileges on the vCenter Server system.

When you refresh STS signing certificates, the VMware Certificate Authority (VMCA) issues a new certificate and replaces the current certificate in the VMware Directory Service (vmdir). STS starts using the new certificate to issue new tokens. In an Enhanced Linked Mode configuration, vmdir uploads the new certificate from the issuing vCenter Server system to all linked vCenter Server systems. When you refresh STS signing certificates, you do not need to restart the vCenter Server system, nor any other vCenter Server system that is part of an Enhanced Linked Mode configuration.

If you are using a custom generated or third-party STS signing certificate, the refresh overwrites that certificate with a VMCA-issued certificate. To update custom generated or third-party STS signing certificates, use the import and replace option. See [Import and Replace a vCenter Server STS Certificate Using the vSphere Client](#).

The VMCA-issued STS signing certificate is valid for 10 years and is not an external-facing certificate. Do not replace this certificate unless the security policy of your company requires it.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Under the **STS Signing** tab, select the desired certificate and click **Refresh with vCenter Certificate**.
If you are using a custom generated or third-party STS signing certificate, the refresh action overwrites that certificate with a VMCA-generated certificate.

Note: If you were using third-party certificates for compliance reasons, the refresh might cause your vCenter Server systems to go out of compliance. Also, if you are using a custom generated or third-party STS signing certificate, the Security Token Service no longer uses that custom or third-party certificate for token signing.
6. Click **Refresh**.
The VMCA refreshes the STS signing certificate on this vCenter Server system and on any linked vCenter Server systems.
7. Optional: If the **Force Refresh** button appears, vCenter Single Sign-On has detected a problem. Before clicking **Force Refresh**, consider the following potential results.
 - If all the impacted vCenter Server systems are not running at least vSphere 7.0 Update 3, they do not support the certificate refresh.
 - Selecting **Force Refresh** requires that you restart all vCenter Server systems and can render those systems inoperable until you do so.
 - a) If you are unsure of the impact, click **Cancel** and research your environment.
 - b) If you are sure of the impact, click **Force Refresh** to proceed with the refresh then manually restart your vCenter Server systems.

Import and Replace a vCenter Server STS Certificate Using the vSphere Client

You can import and replace the vCenter Server STS certificate with a custom generated or third-party certificate using the vSphere Client.

For certificate management, you must supply the password of the administrator of the local domain (administrator@vsphere.local by default). You also must supply the vCenter Single Sign-On credentials for a user with administrator privileges on the vCenter Server system.

To import and replace the default STS signing certificate, you must first generate a new certificate. When you import and replace STS signing certificates, the VMware Directory Service (vmdir) uploads the new certificate from the issuing vCenter Server system to all linked vCenter Server systems.

The STS certificate is not an external-facing certificate. Do not replace this certificate unless the security policy of your company requires it.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Under the **STS Signing** tab, select the desired certificate and click **Import and Replace Certificate**.
6. Select the PEM file.
The PEM file includes the signing certificate chain and the private key.
7. Click **Replace**.
The STS signing certificate is replaced on this vCenter Server system and on any linked vCenter Server systems. Unless otherwise indicated, you do not need to restart the vCenter Server systems.

Replace a vCenter Server STS Certificate Using the Command Line

You can replace the vCenter Server STS certificate with a custom generated or third-party certificate using the CLI.

Enable SSH login to vCenter Server. See [Manage vCenter Server Using the vCenter Server Shell](#).

To use a company required certificate or to refresh a certificate that is near expiration, you can replace the existing STS signing certificate. To replace the default STS signing certificate, you must first generate a new certificate.

The STS certificate is not an external-facing certificate. Do not replace this certificate unless the security policy of your company requires it.

CAUTION: You must use the procedures described here. Do not replace the certificate directly in the file system.

1. Log in to the vCenter Server shell as root.
2. Create a certificate.
 - a) Create a top-level directory to hold the new certificate and verify the location of the directory.


```
mkdir newsts
cd newsts
pwd
#resulting output: /root/newsts
```
 - b) Copy the certtool.cfg file into the new directory.


```
cp /usr/lib/vmware-vmca/share/config/certtool.cfg /root/newsts
```
 - c) Using a command-line editor such as Vim, open your copy of the certtool.cfg file and edit it to use the local vCenter Server IP address and hostname. The country is required and has to be two characters, as shown in the following example.

```
#
# Template file for a CSR request
#

# Country is needed and has to be 2 characters
Country = US
```



```
Name = STS
Organization = ExampleInc
OrgUnit = ExampleInc Dev
State = Indiana
Locality = Indianapolis
IPAddress = 10.0.1.32
Email = chen@exampleinc.com
Hostname = homecenter.exampleinc.local
```

d) Generate the key.

```
/usr/lib/vmware-vmca/bin/certool --server localhost --genkey --privkey=/root/newsts/sts.key --pub-
key=/root/newsts/sts.pub
```

e) Generate the certificate.

```
/usr/lib/vmware-vmca/bin/certool --gencert --cert=/root/newsts/newsts.cer --privkey=/root/new-
sts/sts.key --config=/root/newsts/certool.cfg
```

f) Create a PEM file with the certificate chain and private key.

```
cat newsts.cer /var/lib/vmware/vmca/root.cer sts.key > newsts.pem
```

3. Update the STS signing certificate, for example:

```
/opt/vmware/bin/sso-config.sh -set_signing_cert -t vsphere.local /root/newsts/newsts.pem
```

The VMCA refreshes the STS signing certificate on this vCenter Server system and on any linked vCenter Server systems.

View the Active vCenter Server STS Signing Certificate Chain Using the vSphere Client

You can use the vSphere Client to view the active vCenter Server STS signing certificate chain and the certificate information, such as the valid until date.

1. Log in with the vSphere Client to the vCenter Server.
2. Enter the user name and password for a user that has at least Read privileges.
3. Navigate to the Certificate Management UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Certificates**, click **Certificate Management**.
4. If the system prompts you, enter the credentials of your vCenter Server.
5. Under the **STS Signing** tab, select a certificate, then expand the certificate.

Certificate and issue information is displayed, including:

 - Valid until date
 - A green check for a valid certificate, and an orange check warning of an expired certificate

Determine the Expiration Date of an LDAPS SSL Certificate Using the Command Line

When using Active Directory over LDAPS, you can upload an SSL certificate for the LDAP traffic. SSL certificates expire after a predefined lifespan. You can use the `sso-config.sh` command to view the certificate's expiration date so that you know to replace or renew the certificate before it expires.

Enable SSH login to vCenter Server. See [Manage vCenter Server Using the vCenter Server Shell](#).

vCenter Server alerts you when an active LDAP SSL certificate is close to its expiration date.

You see certificate expiration information only if you use Active Directory over LDAP or an OpenLDAP identity source and specify an `ldaps://` URL for the server.

1. Log in as root to the vCenter Server.
2. Run the following command.

```
/opt/vmware/bin/sso-config.sh -get_identity_sources
```

Ignore the SLF4J messages.

3. To determine the expiration date, view the SSL certificate's details and verify the `NotAfter` field.

Managing vCenter Single Sign-On Policies

vCenter Single Sign-On policies enforce the security rules for local accounts and tokens in general. You can view and edit the default vCenter Single Sign-On password policy, lockout policy, and token policy.

Edit the vCenter Single Sign-On Password Policy

The vCenter Single Sign-On password policy determines the password format and password expiration. Password policy applies only to users in the vCenter Single Sign-On domain (`vsphere.local`).

By default, vCenter Single Sign-On built-in user account passwords expire after 90 days. The vSphere Client reminds you when your password is about to expire.

See [Change Your vCenter Single Sign-On Password](#).

Note: The administrator account (`administrator@vsphere.local`) does not get locked out nor does its password expire. Proper security practice is to audit logins from this account and rotate the password regularly.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for `administrator@vsphere.local` or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as `administrator@mydomain`.
3. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
4. Click the **Local Accounts** tab.
5. Click **Edit** for the **Password Policy** row.
6. Edit the password policy.

Option	Description
Description	Password policy description.
Maximum lifetime	Maximum number of days that a password is valid before the user must change it. The maximum number of days you can enter is 999999999. A value of zero (0) means that the password never expires.
Restrict reuse	Number of previous passwords that cannot be reused. For example, if you enter 6, the user cannot reuse any of the last six passwords.
Maximum length	Maximum number of characters that are allowed in the password. The maximum number of characters supported for password is 32. The behaviour is undefined if the minimum and maximum password length is beyond this permissible range.

Option	Description
Minimum length	Minimum number of characters required in the password. The minimum length must be no less than the combined minimum of alphabetic, numeric, and special character requirements. The minimum number of characters supported for password is 8. The behaviour is undefined if the minimum and maximum password length is beyond this permissible range.
Character requirements	Minimum number of different character types that are required in the password. You can specify the number of each type of character, as follows: • Special: & # % • Alphabetic: A b c D • Uppercase: A B C • Lowercase: a b c • Numeric: 1 2 3 • Identical Adjacent: The number must be greater than 0. For example, if you enter 1, the following password is not allowed: p@\$\$word. The minimum number of alphabetic characters must be no less than the combined uppercase and lowercase characters.

Note: The password policy picks up the maximum length value only if the minimum length is greater than 20 characters. The behavior of the password policy is undefined or could result in failure of services when the minimum length value is greater than 20 characters and the maximum length is set to any value. To avoid a potential problem, leave the minimum length set to the default value of 8 characters, or no greater than 20 characters.

7. Click **Save**.

Edit the vCenter Single Sign-On Lockout Policy

If a user attempts to log in with incorrect credentials, a vCenter Single Sign-On lockout policy specifies when the user's vCenter Single Sign-On account is locked. Administrators can edit the lockout policy.

If a user logs in to vsphere.local multiple times with the wrong password, the user is locked out. The lockout policy allows administrators to specify the maximum number of failed login attempts, and set the time interval between failures. The policy also specifies how much time must elapse before the account is automatically unlocked.

Note: The lockout policy applies only to user accounts, not to system accounts such as administrator@vsphere.local.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
4. Click the **Local Accounts** tab.
5. Click **Edit** for the **Lockout Policy** row.
You might need to scroll down to see the **Lockout Policy** row.

6. Edit the parameters.

Option	Description
Description	Optional description of the lockout policy.
Maximum number of failed login attempts	Maximum number of failed login attempts that are allowed before the account is locked.
Time interval between failures	Time period in which failed login attempts must occur to trigger a lockout.
Unlock time	Amount of time that the account remains locked. If you enter 0, the administrator must unlock the account explicitly.

7. Click **Save**.

Edit the vCenter Single Sign-On Token Policy

The vCenter Single Sign-On token policy specifies token properties such as the clock tolerance and renewal count. You can edit the token policy to ensure that the token specification conforms to security standards in your corporation.

- Log in with the vSphere Client to the vCenter Server.
- Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
- Navigate to the Configuration UI.
 - From the **Home** menu, select **Administration**.
 - Under **Single Sign On**, click **Configuration**.
- Click the **Local Accounts** tab.
- Click **Edit** for the **Token Trustworthiness** row.
You might need to scroll down to see the **Token Trustworthiness** row.

6. Edit the token policy configuration parameters.

Option	Description
Clock Tolerance	Time difference, in milliseconds, that vCenter Single Sign-On tolerates between a client clock and the domain controller clock. If the time difference is greater than the specified value, vCenter Single Sign-On declares the token invalid.
Maximum Token Renewal Count	Maximum number of times that a token can be renewed. After the maximum number of renewal attempts, a new security token is required.
Maximum Token Delegation Count	Holder-of-key tokens can be delegated to services in the vSphere environment. A service that uses a delegated token performs the service on behalf of the principal that provided the token. A token request specifies a DelegateTo identity. The DelegateTo value can either be a solution token or a reference to a solution token. This value specifies how many times a single holder-of-key token can be delegated.
Maximum Bearer Token Lifetime	Bearer tokens provide authentication based only on possession of the token. Bearer tokens are intended for short-term, single-operation use. A bearer token does not verify the identity of the user or entity that is sending the request. This value specifies the lifetime value of a bearer token before the token has to be reissued.

Option	Description
Maximum Holder-of-Key Token Lifetime	Holder-of-key tokens provide authentication based on security artifacts that are embedded in the token. Holder-of-key tokens can be used for delegation. A client can obtain a holder-of-key token and delegate that token to another entity. The token contains the claims to identify the originator and the delegate. In the vSphere environment, a vCenter Server system obtains delegated tokens on a user's behalf and uses those tokens to perform operations. This value determines the lifetime of a holder-of-key token before the token is marked invalid.

- Click **Save**.

Edit Password Expiration Notification for Active Directory (Integrated Windows Authentication) Users

The Active Directory password expiration notification is separate from the vCenter Server SSO password expiration. The default password expiration notification for an Active Directory user is 30 days but the actual password expiration depends on your Active Directory system. The vSphere Client controls the expiration notification. You can change the default expiration notification to meet the security standards in your corporation.

- Enable SSH login to vCenter Server. See [Manage vCenter Server Using the vCenter Server Shell](#).

- Log in to the vCenter Server shell as a user with administrator privileges.

The default user with the super administrator role is root.

- Change directory to the location of the `vsphere-client.properties` file.

```
cd /etc/vmware/vsphere-ui
```

- Open the `webclient.properties` file with a text editor.

- Edit the following variable.

```
sso.pending.password.expiration.notification.days = 30
```

- Restart the vSphere Client.

```
service-control --stop vsphere-ui
service-control --start vsphere-ui
```

Managing vCenter Single Sign-On Users and Groups

A vCenter Single Sign-On administrator user can manage users and groups in the `vsphere.local` domain from the vSphere Client.

The vSphere Client presents a view of users and groups in your vSphere domain (`vsphere.local` by default). From this view, you can add, edit, and deactivate users. You can also add groups, and manage group membership.

Add vCenter Single Sign-On Users

Users listed on the **Users** tab in the vSphere Client are internal to vCenter Single Sign-On and belong to the vsphere.local domain. You add users to that domain from one of the vCenter Single Sign-On management interfaces.

You can select other domains and view information about the users in those domains, but you cannot add users to other domains from a vCenter Single Sign-On management interface.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the vCenter Single Sign-On user configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
4. If vsphere.local is not the currently selected domain, select it from the drop-down menu.
You cannot add users to other domains.
5. On the **Users** tab, click **Add**.
6. Enter a user name and password for the new user.
The maximum number of characters allowed for the user name is 300.
You cannot change the user name after you create a user. The password must meet the password policy requirements for the system.
7. Optional: Enter the first name and the last name of the new user.
8. Optional: Enter an email address and description for the user.
9. Click **Add**.

When you add a user, that user initially has no privileges to perform management operations.

Add the user to a group in the vsphere.local domain, for example, to the group of users who can administer VMCA (CAAdmins) or to the group of users who can administer vCenter Single Sign-On (Administrators). See [Add Members to a vCenter Single Sign-On Group](#).

Deactivate and Activate vCenter Single Sign-On Users

When a vCenter Single Sign-On user account is deactivated, the user cannot log in to the vCenter Single Sign-On server until an administrator activates the account. You can deactivate and activate accounts from one of the vCenter Single Sign-On management interfaces.

You must be a member of the vCenter Single Sign-On Administrators group to deactivate and activate vCenter Single Sign-On users.

Deactivated user accounts remain available in the vCenter Single Sign-On system, but the user cannot log in or perform operations on the server. Users with administrator privileges can deactivate and activate accounts from the vCenter Server **Users and Groups** page.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.

3. Navigate to the vCenter Single Sign-On user configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
4. Select a user name, click **More**, and click **Disable**.
5. Click **OK**.
6. To activate the user again, click **More**, click **Enable**, and click **OK**.

Delete a vCenter Single Sign-On User

You can delete users that are in the vsphere.local domain from a vCenter Single Sign-On management interface. You cannot delete local operating system users or users in another domain from a vCenter Single Sign-On management interface.

CAUTION: If you delete the administrator user in the vsphere.local domain, you can no longer log in to vCenter Single Sign-On. Reinstall vCenter Server and its components.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the vCenter Single Sign-On user configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
4. Select **Users**, and select the vsphere.local domain from the drop-down menu.
5. In the list of users, select the user that you want to delete.
6. Click **Delete**.
Proceed with caution. You cannot undo this action.
7. Click **Remove**.

Edit a vCenter Single Sign-On User

You can change the password or other details of a vCenter Single Sign-On user from a vCenter Single Sign-On management interface. You cannot rename users in the vsphere.local domain. That means you cannot rename administrator@vsphere.local.

You can create additional users with the same privileges as administrator@vsphere.local.

vCenter Single Sign-On users are stored in the vCenter Single Sign-On vsphere.local domain.

You can review the vCenter Single Sign-On password policies from the vSphere Client. Log in as administrator@vsphere.local and from the **Administration** menu, select **Configuration > Local Accounts > Password Policy**.

See also [Edit the vCenter Single Sign-On Password Policy](#).

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.

3. Navigate to the vCenter Single Sign-On user configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
4. Click **Users**.
5. Select the user and click **Edit**.
6. Edit the user attributes.

You cannot change the user name of the user.

The password must meet the password policy requirements for the system.
7. Click **Save**.

Add a vCenter Single Sign-On Group

The vCenter Single Sign-On **Groups** tab shows groups in the local domain, vsphere.local by default. You add groups if you need a container for group members (principals).

You cannot add groups to other domains, for example, the Active Directory domain, from the vCenter Single Sign-On **Groups** tab.

If you do not add an identity source to vCenter Single Sign-On, creating groups and adding users can help you organize the local domain.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as administrator@*mydomain*.
3. Navigate to the vCenter Single Sign-On user configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
4. Select **Groups** and click **Create Group**.
5. Enter a name and description for the group.

The maximum number of characters allowed for the group name is 300. You cannot change the group name after you create the group.
6. From the **Add Members** drop-down menu, select the identity source that contains the member to add to the group.

If you have configured an external identity provider such as AD FS, the domain of that identity provider is available to select in the **Add Members** drop-down menu.
7. Enter a search term.
8. Select the member.

You can add more than one member.
9. Click **Done**.

See [Add Members to a vCenter Single Sign-On Group](#).

Add Members to a vCenter Single Sign-On Group

Members of a vCenter Single Sign-On group can be users or other groups from one or more identity sources. You can add new members from the vSphere Client.

See the VMware knowledge base article at <https://kb.vmware.com/s/article/2095342> for the background information.

Groups listed on the **Groups** tab in the Web interface are part of the vsphere.local domain. See [Groups in the vCenter Single Sign-On Domain](#).

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the vCenter Single Sign-On user configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
4. Click **Groups** and click the group (for example, Administrators).
5. Click **Edit**.
6. From the **Domain** drop-down menu, select the identity source that contains the member to add to the group.
If you have configured an external identity provider, such as AD FS, the domain of that identity provider is available to select in the **Domain** drop-down menu.
7. Enter a search term.
8. Select the member.
You can add more than one member.
9. For vSphere+ environments if you select **VMware ID** from the **Domain** drop-down menu, then enter the name of the CSP account in the **Username** field.
Note:
Enter the email address of the CSP account in the **Username** field. CSP accounts cannot be searched for in the VMwareID domain.
10. Click **Save**.

Remove Members from a vCenter Single Sign-On Group

You can remove members from a vCenter Single Sign-On group by using the vSphere Client. When you remove a member (user or group) from a group, you do not delete the member from the system.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.

3. Navigate to the vCenter Single Sign-On user configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Users and Groups**.
4. Click **Groups** and select a group.
5. Click **Edit**.
6. In the Current members list, click the user or group that you want to remove.
7. Click **Done**.

The user or group is removed from the group, but is still available in the system.

Change Your vCenter Single Sign-On Password

Users in the local domain, vsphere.local by default, can change their vCenter Single Sign-On passwords from the vSphere Client. Users in other domains change their passwords following the rules for that domain.

The vCenter Single Sign-On lockout policy determines when your password expires. By default, vCenter Single Sign-On passwords expire after 90 days, but administrator passwords such as the password for administrator@vsphere.local do not expire. vCenter Single Sign-On management interfaces show a warning when your password is about to expire.

Note: You can change a password only if it is not expired.

If the password is expired, the administrator of the local domain, administrator@vsphere.local by default, can reset the password by using the `dir-cli password reset` command. Only members of the Administrator group for the vCenter Single Sign-On domain can reset passwords.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as administrator@mydomain.
3. In the upper navigation pane, click your user name to pull down the menu and select **Change Password**.
4. Enter your current password.
5. Enter a new password and confirm it.

The password must conform to the password policy.
6. Click **Confirm**.

As an alternative, you can select **Single Sign On > Users and Groups**, select the user, and click **Edit**.

Other vSphere Authentication Options

In vSphere 7.0 and later, external identity provider federation is the preferred authentication method for vCenter Server. You can still authenticate by using a smart card (UPN-based Common Access Card or CAC), or by using an RSA SecurID token.

Two-Factor Authentication Methods

Government agencies or large enterprises often require two-factor authentication. vSphere supports the following two-factor authentication methods.

External Identity Provider Federation

With external identity provider federation, you can use the authentication mechanisms supported by the external identity provider, including multi-factor authentication.

Smart Card Authentication

Smart card authentication allows access only to users who attach a physical card reader to the computer that they log in to. An example is Common Access Card (CAC) authentication.

The administrator can deploy the PKI so that the smart card certificates are the only client certificates that the CA issues. For such deployments, only smart card certificates are presented to the user. The user selects a certificate, and is prompted for a PIN. Only users who have both the physical card and the PIN that matches the certificate can log in.

RSA SecurID Authentication

For RSA SecurID authentication, your environment must include a correctly configured RSA Authentication Manager. If the vCenter Server is configured to point to the RSA server, and if RSA SecurID Authentication is activated, users can log in with their user name and token.

See the vSphere Blog post, [RSA SecurID setup](#), for details.

Note: vCenter Single Sign-On supports only native SecurID. It does not support RADIUS authentication.

Specifying a vCenter Server Non-default Authentication Method

You can set up a nondefault authentication method from the vSphere Client, or by using the `sso-config` script.

- For smart card authentication, you can perform the vCenter Single Sign-On setup from the vSphere Client or by using `sso-config`. Setup includes activating smart card authentication and configuring certificate revocation policies.
- For RSA SecurID, you use the `sso-config` script to configure RSA Authentication Manager for the domain, and to enable RSA token authentication. You cannot configure RSA SecurID authentication from the vSphere Client. However, if you enable RSA SecurID, that authentication method appears in the vSphere Client.

Combining vCenter Server Authentication Methods

You can activate or deactivate each authentication method separately by using `sso-config`. Leave user name and password authentication enabled initially, while you are testing a two-factor authentication method, and set only one authentication method to enabled after testing.

Smart Card Authentication Login

A smart card is a small plastic card with an embedded integrated circuit chip. Many government agencies and large enterprises use smart cards such as Common Access Card (CAC) to increase the security of their systems and to comply with security regulations. A smart card is used in environments where each machine includes a smart card reader. Smart card hardware drivers that manage the smart card are typically preinstalled.

Note: In vSphere 7.0 Update 2 and later, you can enable FIPS on vCenter Server. See the *vSphere Security* documentation. RSA SecureID and CAC authentication are not supported when FIPS is enabled. Use external identity provider federation for MFA authentication. See [Configuring vCenter Server Identity Provider Federation](#).

Users who log in to a vCenter Server system are prompted to authenticate with a smart card and PIN combination, as follows.

1. When a user inserts the smart card into the smart card reader, the browser reads the certificates on the card.
2. The browser prompts the user to select a certificate, then prompts the user for the PIN for that certificate.
3. vCenter Single Sign-On checks whether the certificate on the smart card is known. If revocation checking is turned on, vCenter Single Sign-On also checks whether the certificate is revoked.

4. If the certificate is known to vCenter Single Sign-On, and is not a revoked certificate, the user is authenticated and can perform tasks for which that the user has permissions.

Note: It usually makes sense to leave user name and password authentication enabled during testing. After testing is complete, deactivate user name and password authentication and activate smart card authentication. Subsequently, the vSphere Client allows only smart card login. Only users with root or administrator privileges on the machine can reactivate user name and password authentication by logging in to the vCenter Server directly.

Configuring and Using Smart Card Authentication

You can set up your environment to require smart card authentication when a user connects to a vCenter Server from the vSphere Client.

Configuring smart card authentication involves the following high-level steps:

1. Configuring the vCenter Server system to request client certificates.
2. Activating the smart card configuration.
You can use either the vSphere Client or the `sso-config` utility to activate the configuration.
3. Customizing certificate revocation checking.
You can use either the vSphere Client or the `sso-config` utility to customize the checking.

Configure vCenter Server to Request Client Certificates

Before you activate smart card authentication, you must configure the vCenter Server to request client certificates.

Copy the certificate authority (CA) certificates to the vCenter Server system to use to create the trusted client CA store. This store must contain the trusted certificates issued by the CA for the client certificate. The client here is the browser from which the smart card process prompts the end user for information.

Note: vCenter Server 7.0 and later supports the HTTP/2 protocol. All modern browsers and applications, including the vSphere Client, connect to vCenter Server using HTTP/2. However, smart card authentication requires use of the HTTP/1.1 protocol. Activating smart card authentication deactivates Application-Layer Protocol Negotiation (ALPN, <https://tools.ietf.org/html/rfc7301>) for HTTP/2, effectively preventing the browser from using HTTP/2. Applications that use only HTTP/2, without relying on ALPN, continue to work.

To complete smart card authentication, clients must be permitted access to port 3128/TCP on the appropriate vCenter Server. Check your perimeter firewalls to ensure that access has been granted.

The connection is redirected to port 3128 during smart card login. Port 3128 only supports pre-configured mutual authentication connections and it is not intended as a direct browser endpoint. It does not return an HSTS header. If your vulnerability scanner reports this behavior, it can be safely ignored.

The configuration uses port 3128, which is set and opened automatically on vCenter Server.

1. Log in to the vCenter Server shell as the root user.
2. Create a trusted client CA store on the vCenter Server using the exact path and PEM name, `/usr/lib/vmware-sso/vmware-sts/conf/clienttrustCA.pem`.

Warning: You must use the exact path and PEM name, `/usr/lib/vmware-sso/vmware-sts/conf/clienttrustCA.pem`.

- a) Change to the `/usr/lib/vmware-sso/` directory.

```
cd /usr/lib/vmware-sso/
```

- b) To create the trusted client CA store, run the `openssl` command, taking as input your trusted signing certificate. For example, the following command creates the `clienttrustCA.pem` file from the `xyzCompanySmartCardSigningCA.cer` trusted signing certificate.

```
openssl x509 -inform PEM -in xyzCompanySmartCardSigningCA.cer > /usr/lib/vmware-sso/vmware-sts/conf/clienttrustCA.pem
```

You can add additional certificates to the trusted client CA store by running the `openssl` command with the ">>" operator to append the additional certificate. For example, the following command appends `xyzCompanySmartCardSigningCA2.cer` to the existing `clienttrustCA.pem` file.

```
openssl x509 -inform PEM -in xyzCompanySmartCardSigningCA2.cer >> /usr/lib/vmware-sso/vmware-sts/conf/clienttrustCA.pem
```

3. To validate that the contents of the `clienttrustCA.pem` file contain the trusted CAs that signed the smart card certificates, run the `keytool` command.

For example:

```
keytool -printcert -file /usr/lib/vmware-sso/vmware-sts/conf/clienttrustCA.pem | grep -i "owner\|sha1\|issuer:\|valid"
```

4. Verify that the CA names match the Smart Card User Certificate Chain.

For example, you can run the following command.

```
sso-config.sh -get_authn_policy -t vsphere.local | grep trusted
```

The root and intermediate certificates must have matching thumbprints, names, valid dates, and so on.

Note: You can also use the vSphere Client (**Administration > Single Sign On > Configuration > Identity Provider > Smart Card Authentication > Smart card authentication settings > Trusted CA certificates > Add**).

5. Restart the STS service.

```
service-control --restart sts
```

Manage Smart Card Authentication Using the vSphere Client

You can activate and deactivate smart card authentication, customize the login banner, and set up the revocation policy from the vSphere Client.

- Verify that an enterprise Public Key Infrastructure (PKI) is set up in your environment, and that certificates meet the following requirements:
 - A User Principal Name (UPN) must correspond to an Active Directory account in the Subject Alternative Name (SAN) extension.
 - The certificate must specify Client Authentication in the Application Policy or Extended Key Usage field or the browser does not show the certificate.

- Add an Active Directory identity source to vCenter Single Sign-On.
- Assign the vCenter Server Administrator role to one or more users in the Active Directory identity source. Those users can then perform management tasks because they can authenticate and they have vCenter Server administrator privileges.
- Ensure that you have set up the reverse proxy and restarted the physical or virtual machine.

If smart card authentication is activated and other authentication methods are deactivated, users are then required to log in using smart card authentication.

If user name and password authentication are deactivated, and if problems occur with smart card authentication, users cannot log in. In that case, a root or administrator user can turn on user name and password authentication from the vCenter Server command line. The following command activates user name and password authentication.

```
sso-config.sh -set_authn_policy -pwdAuthn true -t tenant_name
```

1. Obtain the certificates and copy them to a folder that the `sso-config` utility can see.
 - a) Log in to the vCenter Server console, either directly or by using SSH.
 - b) Activate the shell, as follows.


```
Command> shell
chsh -s "/bin/bash" root
chsh -s "bin/appliancesh" root
```
 - c) Use WinSCP or a similar utility to copy the certificates to the `/usr/lib/vmware-sso/vmware-sts/conf` directory on the vCenter Server.
 - d) Optionally deactivate the shell, as follows.


```
chsh -s "/bin/appliancesh" root
```
2. Log in with the vSphere Client to the vCenter Server.
3. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as administrator@mydomain.
4. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
5. Under the **Identity Provider** tab, click **Smart Card Authentication**, then click **Edit**.
6. Select or deselect authentication methods, and click **Save**.

You cannot activate or deactivate RSA SecurID authentication from this Web interface. However, if RSA SecurID has been activated from the command line, the status appears in the Web interface.

The **Trusted CA certificates** tab appears.
7. Under the **Trusted CA certificates** tab:
 - a) Click **Add**, and click **Browse**.
 - b) Select a trusted CA certificate, and click **Add**.
8. To add additional trusted CA certificates, repeat step 7.

Your environment might require enhanced OCSP configuration.

- If your OCSP response is issued by a different CA than the signing CA of the smart card, provide the OCSP signing CA certificate.
- You can configure one or more local OCSP responders for each vCenter Server site in a multi-site deployment. You can configure these alternative OCSP responders using the CLI. See [Manage Smart Card Authentication Using the CLI](#).

Manage Smart Card Authentication Using the CLI

You can use the `sso-config` utility to manage smart card authentication from the command line. The utility supports all smart card configuration tasks.

- Verify that an enterprise Public Key Infrastructure (PKI) is set up in your environment, and that certificates meet the following requirements:
 - A User Principal Name (UPN) must correspond to an Active Directory account in the Subject Alternative Name (SAN) extension.
 - The certificate must specify Client Authentication in the Application Policy or Extended Key Usage field or the browser does not show the certificate.
- Add an Active Directory identity source to vCenter Single Sign-On.
- Assign the vCenter Server Administrator role to one or more users in the Active Directory identity source. Those users can then perform management tasks because they can authenticate and they have vCenter Server administrator privileges.
- Ensure that you have set up the reverse proxy and restarted the physical or virtual machine.

You can find the `sso-config` script at the following location:

```
/opt/vmware/bin/sso-config.sh
```

Configuration of supported authentication types and revocation settings is stored in VMware Directory Service and replicated across all vCenter Server instances in a vCenter Single Sign-On domain.

If user name and password authentication are deactivated, and if problems occur with smart card authentication, users cannot log in. In that case, a root or administrator user can turn on user name and password authentication from the vCenter Server command line. The following command activates user name and password authentication.

```
sso-config.sh -set_authn_policy -pwdAuthn true -t tenant_name
```

If you use the default tenant, use `vsphere.local` as the tenant name.

If you use OCSP for revocation check, you can rely on the default OCSP specified in the smart card certificate AIA extension. You can also override the default and configure one or more alternative OCSP responders. For example, you can set up OCSP responders that are local to the vCenter Single Sign-On site to process the revocation check request.

Note: If your certificate does not have OCSP defined, use CRL (certificate revocation list) instead.

1. Obtain the certificates and copy them to a folder that the `sso-config` utility can see.
 - a) Log in to the appliance console, either directly or by using SSH.
 - b) Activate the appliance shell, as follows.


```
shell
chsh -s "/bin/bash" root
```
 - c) Use WinSCP or a similar utility to copy the certificates to the `/usr/lib/vmware-sso/vmware-sts/conf` on the vCenter Server.
 - d) Optionally deactivate the shell, as follows.


```
chsh -s "/bin/appliancesh" root
```

2. To activate smart card authentication, run the following command.

```
sso-config.sh -set_authn_policy -certAuthn true -cacerts first_trusted_cert.cer,second_trusted_cert.cer -t tenant
```

For example:

```
sso-config.sh -set_authn_policy -certAuthn true -cacerts MySmartCA1.cer,MySmartCA2.cer -t vsphere.local
```

Separate multiple certificates with commas, but do not put spaces after the comma.

3. To deactivate all other authentication methods, run the following commands.

```
sso-config.sh -set_authn_policy -pwdAuthn false -t vsphere.local
sso-config.sh -set_authn_policy -winAuthn false -t vsphere.local
sso-config.sh -set_authn_policy -securIDAuthn false -t vsphere.local
```

4. Optional: To set a certificate policies allowlist, run the following command.

```
sso-config.sh -set_authn_policy -certPolicies policies
```

To specify multiple policies, separate them with a comma, for example:

```
sso-config.sh -set_authn_policy -certPolicies 2.16.840.1.101.2.1.11.9,2.16.840.1.101.2.1.11.19
```

This allowlist specifies object IDs of policies that are allowed in the certificate's certificate policy extension. An X509 certificate can have a Certificate Policy extension.

5. Optional: Turn on and configure revocation checking using OCSP.

- a) Turn on revocation checking using OCSP.

```
sso-config.sh -set_authn_policy -t tenantName -useOcsf true
```

- b) If the OCSP responder link is not provided by the AIA extension of the certificates, provide the overriding OCSP responder URL and OCSP authority certificate.

The alternative OCSP is configured for each vCenter Single Sign-On site. You can specify more than one alternative OCSP responder for your vCenter Single Sign-On site to allow for failover.

```
sso-config.sh -t tenant -add_alt_ocsp [-siteID yourPSCClusterID] -ocspUrl http://ocsp.xyz.com/ -oc-
spSigningCert yourOcsfSigningCA.cer
```

Note: The configuration is applied to the current vCenter Single Sign-On site by default. Specify the `siteID` parameter only if you configure alternative OCSP for other vCenter Single Sign-On sites.

Consider the following example.

```
.sso-config.sh -t vsphere.local -add_alt_ocsp -ocspUrl http://failover.ocsp.nsn0.rcvs.nit.disa.mil/ -
ocspSigningCert ./DOD_JITC_EMAIL_CA-29__0x01A5__DOD_JITC_ROOT_CA_2.cer
Adding alternative OCSP responder for tenant :vsphere.local
OCSP responder is added successfully!
[
  site:: 78564172-2508-4b3a-b903-23de29a2c342
  [
    OCSP url:: http://ocsp.nsn0.rcvs.nit.disa.mil/
    OCSP signing CA cert: binary value]
  ]
  [
    OCSP url:: http://failover.ocsp.nsn0.rcvs.nit.disa.mil/
    OCSP signing CA cert: binary value]
  ]
]
```

- c) To display the current alternative OCSP responder settings, run this command.

```
sso-config.sh -t tenantName -get_alt_ocsp]
```

- d) To remove the current alternative OCSP responder settings, run this command.

```
sso-config.sh -t tenantName -delete_alt_ocsp [-allSite] [-siteID pscSiteID_for_the_configuration]
```

6. Optional: To list configuration information, run the following command.

```
sso-config.sh -get_authn_policy -t tenantName
```

Set Revocation Policies for Smart Card Authentication

You can customize certificate revocation checking, and you can specify where vCenter Single Sign-On looks for information about revoked certificates.

- Verify that an enterprise Public Key Infrastructure (PKI) is set up in your environment, and that certificates meet the following requirements:
 - A User Principal Name (UPN) must correspond to an Active Directory account in the Subject Alternative Name (SAN) extension.
 - The certificate must specify Client Authentication in the Application Policy or Extended Key Usage field or the browser does not show the certificate.
- Verify that the vCenter Server certificate is trusted by the end user's workstation. Otherwise, the browser does not attempt authentication.
- Add an Active Directory identity source to vCenter Single Sign-On.
- Assign the vCenter Server Administrator role to one or more users in the Active Directory identity source. Those users can then perform management tasks because they can authenticate and they have vCenter Server administrator privileges.

You can customize the behavior by using the vSphere Client or by using the `sso-config` script. The settings that you select depend in part on what the CA supports.

- If revocation checking is deactivated, vCenter Single Sign-On ignores any CRL or OCSP settings. vCenter Single Sign-On does not perform checks on any certificates.
- If revocation checking is activated, the setup depends on the PKI setup.

OCSP only

If the issuing CA supports an OCSP responder, activate **OCSP** and deactivate **CRL as failover for OCSP**.

CRL only

If the issuing CA does not support OCSP, activate **CRL checking** and deactivate **OCSP checking**.

Both OCSP and CRL

If the issuing CA supports both an OCSP responder and a CRL, vCenter Single Sign-On checks the OCSP responder first. If the responder returns an unknown status or is not available, vCenter Single Sign-On checks the CRL. For this case, activate both **OCSP checking** and **CRL checking**, and activate **CRL as failover for OCSP**.

- If revocation checking is activated, advanced users can specify the following additional settings.

OCSP URL

By default, vCenter Single Sign-On checks the location of the OCSP responder that is defined in the certificate being validated. If the Authority Information Access extension is absent from the certificate or if you want to override it, you can explicitly specify a location.

Use CRL from certificate

By default, vCenter Single Sign-On checks the location of the CRL that is defined in the certificate being validated. Deactivate this option if the CRL Distribution Point extension is absent from the certificate or if you want to override the default.

CRL location

Use this property if you deactivate **Use CRL from certificate** and you want to specify a location (file or HTTP URL) where the CRL is located.

You can further limit which certificates vCenter Single Sign-On accepts by adding a certificate policy.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
4. Under the **Identity Provider** tab, click **Smart Card Authentication**.
5. Click **Certificate revocation** and click **Edit** to activate or deactivate revocation checking.
6. If certificate policies are in effect in your environment, you can add a policy in the **Certificate policies** pane.

Set Up RSA SecurID Authentication

You can set up your environment to require that users log in with an RSA SecurID token. SecurID setup is supported only from the command line.

- Verify that your environment has a correctly configured RSA Authentication Manager and that users have RSA tokens. RSA Authentication Manager version 8.0 or later is required.
- Verify that the identity source that RSA Manager uses has been added to vCenter Single Sign-On. See [Add or Edit a vCenter Single Sign-On Identity Source](#).
- Verify that the RSA Authentication Manager system can resolve the vCenter Server host name, and that the vCenter Server system can resolve the RSA Authentication Manager host name.
- Export the `sdconf.rec` file from the RSA Manager by selecting **Access > Authentication Agents > Generate configuration file**. To find `sdconf.rec` file, decompress the resulting `AM_Config.zip` file.
- Copy the `sdconf.rec` file to the vCenter Server node.

See the two vSphere Blog posts about [RSA SecurID setup](#) for details.

Note: RSA Authentication Manager requires that the user ID is a unique identifier that uses 1 to 255 ASCII characters. The characters ampersand (&), percent (%), greater than (>), less than (<), and single quote (') are not allowed.

1. Change to the directory where the `sso-config` script is located.

```
/opt/vmware/bin
```

2. To activate RSA SecurID authentication, run the following command.

```
sso-config.sh -t tenantName -set_authn_policy -securIDAuthn true
```

`tenantName` is the name of the vCenter Single Sign-On domain, `vsphere.local` by default.

3. Optional: To deactivate other authentication methods, run the following command.

```
sso-config.sh -set_authn_policy -pwdAuthn false -winAuthn false -certAuthn false -t vsphere.local
```

4. To configure the environment so that the tenant at the current site uses the RSA site, run the following command.

```
sso-config.sh -set_rsa_site [-t tenantName] [-siteID Location] [-agentName Name] [-sdConfFile Path]
```

For example:

```
sso-config.sh -set_rsa_site -agentName SSO_RSA_AUTHSDK_AGENT -sdConfFile /tmp/sdconf.rec
```

You can specify the following options.

Option	Description
siteID	Optional Platform Services Controller site ID. Platform Services Controller supports one RSA Authentication Manager instance

Option	Description
	or cluster per site. If you do not explicitly specify this option, the RSA configuration is for the current Platform Services Controller site. Use this option only if you are adding a different site.
agentName	Defined in RSA Authentication Manager.
sdConfFile	Copy of the <code>sdconf.rec</code> file that was downloaded from RSA Manager and includes configuration information for the RSA Manager, such as the IP address.

5. Optional: To change the tenant configuration to nondefault values, run the following command.

```
sso-config.sh -set_rsa_config [-t tenantName] [-logLevel Level] [-logFileSize Size] [-maxLogFileCount-
Count] [-connTimeOut Seconds] [-readTimeOut Seconds] [-encAlgList Alg1,Alg2,...]
```

The default is usually appropriate, for example:

```
sso-config.sh -set_rsa_config -t vsphere.local -logLevel DEBUG
```

6. Optional: If your identity source is not using the User Principal Name as the user ID, set up the identity source `userID` attribute. (Supported with Active Directory over LDAP identity sources only.)

The `userID` attribute determines which LDAP attribute is used as the RSA `userID`.

```
sso-config.sh -set_rsa_userid_attr_map [-t tenantName] [-idsName Name] [-ldapAttr AttrName] [-siteID Loca-
tion]
```

For example:

```
sso-config.sh -set_rsa_userid_attr_map -t vsphere.local -idsName ssolabs.com -ldapAttr userPrincipalName
```

7. To display the current settings, run the following command.

```
sso-config.sh -t tenantName -get_rsa_config
```

If user name and password authentication is deactivated and RSA authentication is activated, users must log in with their user name and RSA token. User name and password login is no longer possible.

Note: Use the user name format `userID@domainName` or `userID@domain_upn_suffix`.

Managing the Login Message to the vSphere Client Login Page

You can create a message that is displayed on the vSphere Client login page.

You can set a message, disclaimer, or terms and conditions. Also, you can configure the message to require acknowledgment of the message before login.

Manage the Login Message to the vSphere Client Login Page

You can add a login message to the vSphere Client login page. You can also configure a custom login message and provide a check box for user consent.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for `administrator@vsphere.local` or another member of the vCenter Single Sign-On Administrators group.

If you specified a different domain during installation, log in as `administrator@mydomain`.

3. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
4. Click the **Login Message** tab.
5. Click **Edit** and configure the login message.

Option	Description
Show login message	Toggle on Show login message to enable the login message. You cannot make changes to the login message unless you toggle on this switch.
Login message	Title of the message. By default, when Consent checkbox is toggled on, the login message text is <code>I agree to Terms and Conditions</code> . You must replace <code>Terms and Conditions</code> with your own text. If the Consent checkbox is toggled off, then <code>Login message</code> appears, over which you enter your message.
Consent checkbox	Toggle on Consent checkbox to require that the user clicks a check box before logging in. You can also display a message without a check box.
Details of login message	Message that the user sees when clicking the login message, for example, the text of the terms and conditions. You must enter some details in this text box.

6. Click **Save**.

vCenter Single Sign-On Security Best Practices

Follow vCenter Single Sign-On security best practices to protect your vSphere environment.

The vSphere authentication infrastructure enhances the security of your vSphere environment. To make sure that your infrastructure is not compromised, follow these vCenter Single Sign-On best practices.

Check Password Expiration

The default vCenter Single Sign-On password policy has a password lifetime of 90 days. After 90 days, the password expires and you can no longer log in. Check the expiration and refresh passwords in a timely fashion.

Configure Network Time Protocol

Use Network Time Protocol (NTP) to ensure that all systems use the same relative time source (including the relevant localization offset), and that the relative time source can be correlated to an agreed-upon time standard (such as Coordinated Universal Time—UTC). Synchronized systems are essential for vCenter Single Sign-On certificate validity, and for the validity of other vSphere certificates.

NTP also makes it easier to track an intruder in log files. Incorrect time settings can make it difficult to inspect and correlate log files to detect attacks, and can make auditing inaccurate.

See the *vSphere Security* documentation for instructions on configuring time synchronization using NTP.

Troubleshooting vCenter Server Authentication

The following topics provide a starting point for troubleshooting vCenter Server authentication problems. Search this documentation center and the VMware Knowledge Base system for additional pointers.

Determining the Cause of a Lookup Service Error

vCenter Single Sign-On installation displays an error referring to the vCenter Server or the vSphere Client.

vCenter Server and Web Client installers show the error `Could not contact Lookup Service. Please check VM_ssoreg.log...`

This problem has several causes, including unsynchronized clocks on the host machines, firewall blocking, and services that must be started.

1. Verify that the clocks on the host machines running vCenter Single Sign-On, vCenter Server, and the Web Client are synchronized.

2. View the specific log file found in the error message.

In the message, system temporary folder refers to %TEMP%.

3. Within the log file, search for the following messages.

The log file contains output from all installation attempts. Locate the last message that shows `Initializing registration provider...`

Message	Cause and solution
<code>java.net.ConnectException: Connection timed out: connect</code>	The IP address is incorrect, a firewall is blocking access to vCenter Single Sign-On, or vCenter Single Sign-On is overloaded. Ensure that a firewall is not blocking the vCenter Single Sign-On port (by default 7444). Ensure also that the machine on which vCenter Single Sign-On is installed has adequate free CPU, I/O, and RAM capacity.
<code>java.net.ConnectException: Connection refused: connect</code>	The IP address or FQDN is incorrect and the vCenter Single Sign-On service has not started or has started within the past minute. Verify that vCenter Single Sign-On is working by checking the status of vCenter Single Sign-On vmware-ssso daemon. Restart the service. If restarting does not correct the problem, see the recovery section of the <i>vSphere Troubleshooting Guide</i>.
<code>Unexpected status code: 404. SSO Server failed during initialization</code>	Restart vCenter Single Sign-On. If restarting does not correct the problem, see the Recovery section of the <i>vSphere Troubleshooting Guide</i>.
<code>The error shown in the UI begins with Could not connect to vCenter Single Sign-On</code>	You also see the return code <code>SslHandshakeFailed</code>. This error indicates that the provided IP address or FQDN that resolves to vCenter Single Sign-On host was not the address used when you installed vCenter Single Sign-On. In the <code>VM_ssoreg.log</code>, find the line that contains the following message. host name in certificate did not match: <code><install-configured FQDN or IP> != <A> or or <C></code> where A was the FQDN you entered during the vCenter Single Sign-On installation, and B and C are system-generated allowable alternatives. Correct the configuration to use the FQDN on the right of the <code>!=</code> sign in the log file. In most cases, use the FQDN that you specified during vCenter Single Sign-On installation. If none of the alternatives are possible in your network configuration, recover your vCenter Single Sign-On SSL configuration.

Unable to Log In Using Active Directory Domain Authentication

You log in to a vCenter Server component from the vSphere Client. You use your Active Directory user name and password. Authentication fails.

You add an Active Directory identity source to vCenter Single Sign-On, but users cannot log in to vCenter Server.

Users use their user name and password to log in to the default domain. For all other domains, users must include the domain name (user@domain or DOMAIN\user).

For all vCenter Single Sign-On deployments, you can change the default identity source. After that change, users can log in to the default identity source with user name and password only.

To configure your Integrated Windows Authentication identity source with a child domain within your Active Directory forest, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2070433>. By default, Integrated Windows Authentication uses the root domain of your Active Directory forest.

If changing the default identity source does not resolve the issue, perform the following additional troubleshooting steps.

1. Synchronize the clocks between the vCenter Server and the Active Directory domain controllers.
2. Verify that each domain controller has a pointer record (PTR) in the Active Directory domain DNS service. Verify that the PTR record information for the domain controller matches the DNS name of the controller. When using the vCenter Server, run the following commands to perform the task:

1. To list the domain controllers, run the following command:

```
# dig SRV _ldap._tcp.my-ad.com
```

The relevant addresses are in the answer section, as in the following example:

```
;; ANSWER SECTION:
_ldap._tcp.my-ad.com. (...) my-controller.my-ad.com
...
```

2. For each domain controller, verify forward and reverse resolution by running the following command:

```
# dig my-controller.my-ad.com
```

The relevant addresses are in the answer section, as in the following example:

```
;; ANSWER SECTION:
my-controller.my-ad.com (...) IN Acontroller IP address
...
# dig -x <controller IP address>
```

The relevant addresses are in the answer section, as in the following example:

```
;; ANSWER SECTION:
IP-in-reverse.in-addr.arpa. (...) IN PTR my-controller.my-ad.com
...
```

3. Close all browser sessions connected to the vCenter Server and restart all services.

```
/bin/service-control --restart --all
```

vCenter Server Login Fails Because the User Account Is Locked

When you log in to vCenter Server from the vSphere Client login page, an error indicates that the account is locked.

After several failed attempts, you cannot log in to the vSphere Client using vCenter Single Sign-On. You see the message that your account is locked.

You exceeded the maximum number of failed login attempts.

- If you attempted log in as a user from the system domain (vsphere.local by default), ask your vCenter Single Sign-On administrator to unlock your account. If the lock is set to expire in the lockout policy, you can wait until your account is unlocked. vCenter Single Sign-On administrators can use CLI commands to unlock your account.
- If you log in as a user from an Active Directory or LDAP domain, ask your Active Directory or LDAP administrator to unlock your account.

VMware Directory Service Replication Can Take a Long Time

If your environment includes multiple vCenter Server instances connected through enhanced linked mode, and if one of the vCenter Server instances becomes unavailable, your environment continues to function. When the vCenter Server becomes available again, user data and other information are usually replicated within 30 seconds with partners connected through enhanced linked mode. In certain circumstances, however, replication might take a long time.

In certain situations, for example, when your environment includes multiple vCenter Server instances in different locations, and you make significant changes while one vCenter Server is unavailable, you do not see replication across VMware Directory Service instances right away. For example, you do not see a new user that was added to the available vCenter Server instance in the other instance until replication is complete. Replication might take a long time, depending on your enhanced linked mode topology.

During normal operation, changes to a VMware Directory Service (vmdir) instance in one vCenter Server instance (node) show up in its direct replication partner within about 30 seconds. Depending on the replication topology, changes in one node might have to propagate through intermediate nodes before they arrive at each vmdir instance in each node. Information that is replicated includes user information, certificate information, license information for virtual machines that are created, cloned, or migrated with VMware vMotion, and more.

When the replication link is broken, for example, because of a network outage or because a node becomes unavailable, changes in the federation do not converge. After the unavailable node is restored, each node tries to catch up with all changes. Eventually, all vmdir instances converge to a consistent state but it might take a while to reach that consistent state if many changes occurred while one node was unavailable.

Your environment functions normally while replication happens. Do not attempt to solve the problem unless it persists for over an hour.

Export a vCenter Server Support Bundle

You can export a support bundle that contains the log files for the vCenter Server services from the vSphere Client, or by using an API. After the export, you can explore the logs locally or send the bundle to VMware Support.

Verify that the vCenter Server is successfully deployed and running.

For more information about the API, see *vCenter Server Management Programming Guide*.

1. From a Web browser, connect to the vCenter Server configuration management interface at `https://vcenter_server_ip:5480`.
2. Log in as the root user for the vCenter Server.
3. From the **Actions** menu, select **Create Support Bundle**.
4. Unless browser settings prevent an immediate download, the support bundle is saved to your local machine.

vCenter Server Authentication Services Logs Reference

The vCenter Server authentication services use syslog for logging. You can examine the log files to determine the reasons for failures.

Table 116: vCenter Server Authentication Services Logs

Service	Description
VMware Directory Service	By default, vmdir logging goes to <code>/var/log/messages</code> or <code>/var/log/vmware/vmdird/</code> . For issues at deployment time, <code>/var/log/vmware/vmdir/vmafdvmdirclient.log</code> might also contain useful troubleshooting data.
VMware Single Sign-On	vCenter Single Sign-On logging goes to <code>/var/log/vmware/sso/</code> .
VMware Certificate Authority (VMCA)	VMCA service log is located in <code>/var/log/vmware/vmcad/vmcad-syslog.log</code> .
VMware Endpoint Certificate Store (VECS)	VECS service log is located in <code>/var/log/vmware/vmafdd/vmafdd-syslog.log</code> .
VMware Lookup Service	Lookup service log is located in <code>/var/log/vmware/sso/lookupServer.log</code> .

Managing Host and Cluster Lifecycle

Managing Host and Cluster Lifecycle provides information about configuring and using VMware®vSphere Lifecycle Manager to manage the ESXi hosts and clusters in your environment.

Managing Host and Cluster Lifecycle provides instructions for configuring vSphere Lifecycle Manager, working with the vSphere Lifecycle Manager depot, and using baselines and images to install, update, or upgrade the software and firmware running on your ESXi hosts.

Managing Host and Cluster Lifecycle also provides detailed guidelines about using vSphere Lifecycle Manager recommended images and performing hardware compatibility checks on single hosts or clusters. It also describes how you can configure and use the Update Manager Download Service (UMDS) to download software updates in deployments with no access to the Internet.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

This information is intended for experienced system administrators who are familiar with data center operations and virtual machine technology.

Client Interface

The instructions in this guide reflect the HTML5-based vSphere Client.

Updated Information

This *Managing Host and Cluster Lifecycle* document is updated with each release of the product or when necessary.

This table provides the update history of the *Managing Host and Cluster Lifecycle*.

Revision	Description
24 MAR 2025	Added information about the changes to the online depot URLs in What Is an Online Depot and What Is the Default VMware Online Depot, Configuring the Download Sources, and Installing, Setting Up, and Using Download Service.
27 AUG 2024	Added information about how to set up alarms to monitor drifts between the current and the desired cluster configuration. See Set Up Alarms for Monitoring Configuration Drifts.
26 JUNE 2024	Updated the list of supported Linux-based operating systems for installing the Update Manager Download Service (UMDS). See Supported Linux-Based Operating Systems for Installing UMDS.

What is vSphere Lifecycle Manager

VMware vSphere Lifecycle Manager enables centralized and simplified lifecycle management for VMware ESXi hosts through the use of images and baselines.

Lifecycle Management in vSphere

Lifecycle management refers to the process of installing software, maintaining it through updates and upgrades, and decommissioning it.

In the context of maintaining a vSphere environment, your clusters and hosts in particular, lifecycle management refers to tasks such as installing ESXi and firmware on new hosts, and updating or upgrading the ESXi version and firmware when required.

vSphere Lifecycle Manager Overview

vSphere Lifecycle Manager is a service that runs in vCenter Server and uses the embedded vCenter Server PostgreSQL database. No additional installation is required to start using that feature. Upon deploying the vCenter Server appliance, the vSphere Lifecycle Manager user interface becomes automatically enabled in the HTML5-based vSphere Client.

vSphere Lifecycle Manager encompasses the functionality that Update Manager provides in earlier vSphere releases and enhances it by adding new features and options for ESXi lifecycle management at a cluster level.

In vSphere releases earlier than 7.0, Update Manager provides you with the ability to use baselines and baseline groups for host patching and host upgrade operations. Starting with vSphere 7.0, vSphere Lifecycle Manager introduces the option of using vSphere Lifecycle Manager images as an alternative way to manage the lifecycle of the hosts and clusters in your environment. You can also use vSphere Lifecycle Manager to upgrade the virtual machine hardware and VMware Tools versions of the virtual machines in your environment.

vSphere Lifecycle Manager can work in an environment that has access to the Internet, directly or through a proxy server. It can also work in a secured network without access to the Internet. In such cases, you use the Update Manager Download Service (UMDS) to download updates to the vSphere Lifecycle Manager depot, or you import them manually.

vSphere Lifecycle Manager Operations

The basic vSphere Lifecycle Manager operations are related to maintaining an environment that is up-to-date and ensuring smooth and successful updates and upgrades of the ESXi hosts.

Operation	Description
Compliance Check	An operation of scanning ESXi hosts to determine their level of compliance with a baseline attached to the cluster or with the image that the cluster uses. The compliance check does not alter the object.
Remediation Pre-Check	An operation that you perform before remediation to ensure that the health of a cluster is good and that no issues occur during the remediation process.
Remediation	An operation of applying software updates to the ESXi hosts in a cluster. During remediation, you install software on the hosts. Remediation makes a non-compliant host compliant with the baselines attached to the cluster or with the image for cluster.
Staging	An operation that reduces the time ESXi hosts spend in maintenance mode. When you stage an image or baseline to an ESXi host, vSphere Lifecycle Manager downloads the respective bulletins or components from the vSphere Lifecycle Manager depot to the host without applying them immediately. Staging makes the components, patches, and extensions available locally on the hosts. You can choose to remediate the hosts at a later time, not immediately after staging.

The vSphere Lifecycle Manager Depot

Several components make up vSphere Lifecycle Manager and work together to deliver the vSphere Lifecycle Manager functionality and coordinate the major lifecycle management operations that it provides for. The vSphere Lifecycle Manager depot is an important component in the vSphere Lifecycle Manager architecture, because it contains all software updates that you use to create vSphere Lifecycle Manager baselines and images. You can use vSphere Lifecycle Manager only if the vSphere Lifecycle Manager depot is populated with components, add-ons, base images, and legacy bulletins and patches.

For more information about software updates and how they are distributed, see [Software Packaging Units That vSphere Lifecycle Manager Can Consume](#).

For more information about the vSphere Lifecycle Manager depot, see [The vSphere Lifecycle Manager Depot](#).

Secure Hashing and Signature Verification in vSphere Lifecycle Manager

vCenter Server performs an automatic hash check on all software that vSphere Lifecycle Manager downloads from online depots or from a UMDS-created depot. Similarly, vCenter Server performs an automatic checksum check on all software that you manually import into the vSphere Lifecycle Manager depot. The hash check verifies the sha-256 checksum of the downloaded software to ensure its integrity. During remediation, before vSphere Lifecycle Manager installs any software on a host, the ESXi host checks the signature of the installable units to verify that they are not corrupted or altered during the download.

When you import an ISO image into the vSphere Lifecycle Manager depot, vCenter Server performs an MD5 hash check on the ISO image to validate its MD5 checksum. During remediation, before the ISO image is installed, the ESXi host verifies the signature inside the image.

If an ESXi host is configured with UEFI Secure Boot, the ESXi host performs full signature verification of each package that is installed on the host every time the host boots. For more information, see the *vSphere Security* documentation.

vSphere Lifecycle Manager Scalability

For information about the scalability that vSphere Lifecycle Manager supports, visit the VMware Configuration Maximums Matrix at <https://configmax.vmware.com/>.

The vSphere Lifecycle Manager User Interface in the vSphere Client

After you deploy the vCenter Server appliance, vSphere Lifecycle Manager becomes immediately visible in the vSphere Client.

The vSphere Lifecycle Manager user interface has two main views, which for convenience this guide calls the home view and the compliance view.

vSphere Lifecycle Manager Home View

The vSphere Lifecycle Manager home view is where you configure and administer the vSphere Lifecycle Manager instance that runs on your vCenter Server system. You go to the vSphere Lifecycle Manager home view to configure how vSphere Lifecycle Manager baselines and images work.

You do not need any special privilege to access the vSphere Lifecycle Manager home view.

To access the vSphere Lifecycle Manager home view in the vSphere Client, select **Menu > Lifecycle Manager**.

In the vSphere Lifecycle Manager home view, you specify the vSphere Lifecycle Manager instance that you want to administer by selecting a vCenter Server system from the drop-down menu at the top of the **Lifecycle Manager** pane.

In the **Lifecycle Manager** pane, you have the following top-level tabs: **Image Depot**, **Updates**, **Imported ISOs**, **Baselines**, and **Settings**.

You use the **Image Depot** tab when you work vSphere Lifecycle Manager images. You use the **Updates**, **Imported ISOs**, and **Baselines** tabs when you work with vSphere Lifecycle Manager baselines. For more information about the **Image Depot**, **Updates**, and **Imported ISOs** tabs, see [Browsing the vSphere Lifecycle Manager Depot](#).

The **Settings** tab is where you configure all vSphere Lifecycle Manager remediation settings and download sources. You use the **Settings** tab with both vSphere Lifecycle Manager images and baselines. For more information about how to configure the vSphere Lifecycle Manager settings, see [vSphere Lifecycle Manager Remediation Settings](#).

In the vSphere Lifecycle Manager home view, you can perform the following tasks:

- Browse the vSphere Lifecycle Manager depot.
- Trigger the synchronization of updates with the configured online depots.
- Trigger the synchronization of hardware compatibility data.
- Import offline depots manually.
- Import ISO images to use for the creation of upgrade baselines.
- Create and manage baselines and baseline groups.
- Configure the default vSphere Lifecycle Manager download source.
- Add a URL to an online depot to the list of download sources.
- Allow or disallow downloading from a download source.
- Configure host remediation settings.
- Configure virtual machine rollback settings.

vSphere Lifecycle Manager Compliance View

The vSphere Lifecycle Manager compliance view is where you perform the major vSphere Lifecycle Manager operations - checking the compliance of ESXi hosts against a baseline or an image, staging, remediation pre-checks, remediation, and so on.

You go to the vSphere Lifecycle Manager compliance view to actually use the vSphere Lifecycle Manager baselines and images on your clusters and hosts.

To access the vSphere Lifecycle Manager compliance view in the vSphere Client, you must have the **View Compliance Status** privilege.

Generally, the vSphere Lifecycle Manager compliance view is on the **Updates** tab for a selected object.

Depending on the selected object and whether you use baselines or images to manage the object, you access the vSphere Lifecycle Manager compliance view in two different ways.

- To access the vSphere Lifecycle Manager compliance view for a host or a cluster that you manage with baselines, go to the **Updates** tab for the object and select **Baselines**.

In the **Baselines** pane of the vSphere Lifecycle Manager compliance view, you can perform the following tasks:

- Check the compliance status of ESXi hosts and clusters against baselines or baseline groups.
- Attach and detach baselines and baseline groups to hosts and clusters.
- Generate a remediation pre-check report that lists recommended actions to ensure successful remediation.
- Stage patches or extensions to hosts.
- Check the compliance status of ESXi hosts against an image.
- Remediate hosts against baselines and baseline groups.
- Remediate hosts that are part of a vSAN cluster against system-managed baselines.

- To access the vSphere Lifecycle Manager compliance view for a cluster or a standalone host that you manage with a single image, go to the **Updates** tab for the object and select **Image**.

In the **Image** pane of the vSphere Lifecycle Manager compliance view, you can perform the following tasks:

- Export, import, and edit the image that the cluster or the standalone host use.
- Upgrade the firmware of the ESXi hosts in the cluster or the standalone host.
- Check for and view recommended images for the cluster or the standalone host.
- Check the hardware compatibility for a selected ESXi version against vSAN HCL.
- Check the compliance status of the ESXi hosts in the cluster or the standalone host against the image.
- Run a remediation pre-check to ensure successful remediation.
- Remediate the ESXi hosts against the image that the cluster or the standalone host use.

On the **Updates** tab, you can also perform other tasks.

- Under **Hosts**, select **Hardware Compatibility** to check the hardware compatibility of a host against the VMware Compatibility Guide.
- Under **Hosts**, select **VMware Tools** or **VM Hardware** to check the status of virtual machines and upgrade the VMware Tools version or the virtual hardware version of the virtual machines.

vSphere Lifecycle Manager and vCenter Server Single Sign-On Domain

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, you can configure the settings for each vSphere Lifecycle Manager instance. The modifications to the settings are applied only to the vSphere Lifecycle Manager instance that you specify and are not propagated to the other instances in the group. Similarly, you can use vSphere Lifecycle Manager to perform compliance and status checks, and remediate only those inventory objects that are managed by the vCenter Server system where the respective vSphere Lifecycle Manager instance runs.

Using vSphere Lifecycle Manager with VMware vSphere® Distributed Services Engine™

vSphere Distributed Services Engine provides you with the ability to offload some of the network operations from your server CPU to a data processing unit, or the so-called DPU device on the server. You can use vSphere Lifecycle Manager images and all vSphere Lifecycle Manager operations to manage the life cycle of the software and firmware that runs on the DPU devices of your servers.

Using DPU devices for network acceleration frees up the CPU capacity for business-critical workloads. Besides accelerating networking performance, using DPU devices provides security and compression acceleration.

What Is a DPU Device?

A DPU device is a high-performance network interface card, with added embedded CPU cores, memory, and a hypervisor running on the device independently from the ESXi hypervisor installed on the server. DPU devices are like limited-resource servers with multiple general-purpose compute cores.

The ESXi hypervisor that runs on the DPU device is a fully functional hypervisor, but it can only run on ARM CPU architectures. You don't use ESXi on the DPU to run and provision virtual machines and workloads. The hypervisor that runs on DPU devices is a trimmed version of the ESXi hypervisor that runs on your servers. This trimmed ESXi version is optimized for I/O activities, such as packet offloads, external management, and so on.

From vSphere perspective, the DPU device is a pre-configured device that you can start using without any further configuration or customizations.

You cannot manage DPU devices separately from managing your ESXi hosts. All lifecycle operations that you use to manage the lifecycle of the software and firmware on your hosts are used to also manage the software and firmware lifecycle of DPU devices in a DPU-based environment.

The only DPU devices that vSphere 8.0 supports are NVIDIA BlueField and Pensando Distributed Services Card (Pensando DSC).

How Do I Start Using a Server with a DPU Device?

To be able to use the advantages that vSphere Distributed Services Engine offers, your server must include a supported DPU device. You cannot purchase a DPU device separately and insert it into a legacy non-DPU server.

To start utilizing the DPU device on a server to offload network services, you must perform a fresh ESXi install on the server. You can perform the installation through interactive or scripted mechanisms. During installation, ESXi is installed both on the server and on the DPU device. In vSphere 8.0, the ESXi image contains VIBs for both the ESXi version to be installed on a host and the ESXi version to be installed on the DPU device on that host.

Managing the Lifecycle of the ESXi and Firmware on a DPU Device

You use vSphere Lifecycle Manager to upgrade the ESXi version and firmware on a DPU device. The VIBs that are needed to upgrade the DPU device are included in the base ESXi images that VMware distributes and that you have in the vSphere Lifecycle Manager depot. You don't need to take care of the software and firmware on a DPU device separately. For example, during remediation vSphere Lifecycle Manager updates ESXi both on the host and on the DPU device. In vSphere 8.0 Update 3, vSphere Lifecycle Manager operations happen in parallel on both DPU devices for hosts with dual DPU setup.

Almost all vSphere Lifecycle Manager operations work for DPU-based environments. In vSphere 8.0, only the following vSphere Lifecycle Manager operations don't work for clusters with DPU-backed hosts:

- Recommendation generation
- Hardware compatibility checks

Solutions Interoperability with vSphere Distributed Services Engine

NSX is the only solution that is allowed to push VIBs to ESXi on the server, on the DPU, or both. Currently, no other solutions can work together with the vSphere Distributed Services Engine.

System Requirements for Using vSphere Distributed Services Engine

- ESXi 8.0 or later
- vCenter Server 8.0 or later

Limitations to Using vSphere Lifecycle Manager for a DPU-Backed Cluster

- The vSphere Distributed Services Engine works for clusters and standalone hosts that you manage with a single vSphere Lifecycle Manager image. You cannot add a host with a DPU device to a cluster that you manage with baselines.
- A host can have one or two DPU devices. If a host has two DPU devices, they must run the same firmware and hardware revisions.
- Each host in the DPU-based cluster must have a DPU device.
- All hosts must have DPU devices from the same vendor and of the same model. The DPU generation might vary.

Software Packaging Units That vSphere Lifecycle Manager Can Consume

VMware and other software vendors create and package their software updates into specific units that vSphere Lifecycle Manager can consume. The smallest packaging construct is the component. Components can be further grouped into larger software units, such as add-ons and base images.

vSphere Lifecycle Manager uses software from VMware, original equipment manufacturers (OEMs), and third-party software providers.

- OEMs are VMware partners, for example, Dell, HPE, VMware Cloud on AWS.
- Third-party software providers are providers of I/O filters, device drivers, CIM modules, and so on.

What Is a vSphere Installation Bundle (VIB)?

The VIB is the basic building block for the creation of installation packages for ESXi hosts. A VIB is a software package that contains metadata and a binary payload, which represents the actual piece of software to be installed on an ESXi host. The VIB does not represent an entire feature, but just a single module of the feature. So, the VIB is the smallest installable software unit that VMware and other software vendors ship.

Starting with vSphere 8.0, a VIB can install software on either the ESXi running on the host or the ESXi running on a DPU device on the host. Alternatively, a single VIB can carry updates for both the ESXi versions on the host and on the DPU device.

Table 117: Related Terminology

Term	Definition
VIB Metadata	An XML file (<code>descriptor.xml</code>) that describes the contents of the VIB. It also contains dependency information, textual descriptions, system requirements, and information about bulletins.
Standalone VIB	A VIB that is not included in a bulletin or component.

vSphere Lifecycle Manager does not consume and work with individual VIBs. VIBs must be further packaged into a higher-level construct, such as a base image, an add-on, and so on.

What Is a Bulletin?

The bulletin is a grouping of one or more VIBs. Bulletins are defined within the metadata of the VIB. You use bulletins, and not individual VIBs, to create vSphere Lifecycle Manager baselines, which you attach to inventory objects and use to update and upgrade ESXi hosts.

Table 118: Related Terminology

Term	Definition
Patch	A small software update that provides bug fixes or enhancements to the current version of the software. A patch groups one or more VIBs together to address a particular issue or to provide enhancements to the current version of the software.
Roll-up Bulletin	A collection of patches that are grouped to facilitate downloads and deployment.
Extension	A bulletin that defines a group of VIBs for adding an optional component to an ESXi host. An extension is usually provided by a third party. The third-party provider is also responsible for providing patches and updates for the extension.

What Are Software Components?

The component is the basic packaging construct for VIBs. VMware, OEMs, and third-party software providers deliver software in the form of components. The component is a bulletin with additional metadata specifying the name and the version of the component. Unlike the bulletin, the component is a logical grouping of VIBs that provides you with a complete and visible feature upon installation. Components can contain VIBs for the ESXi version on the server, ESXi version the DPU device of the server, if the server is DPU-backed, or for both of them.

VMware and OEMs do not deliver components independently. VMware bundles components together into fully functional and bootable ESXi base images. OEMs bundle components together into vendor add-ons. Third-party software vendors create and ship software, for example drivers or adapters, as independent components.

What Is a Base Image?

The base image is an ESXi image that VMware provides with every release of ESXi. The base image is a collection of components that is complete and can boot up a server. Base images have a user-readable name and a unique version that is updated with every major or minor release of ESXi.

The version of a base image corresponds to an ESXi release and uses the following naming format:

Release Type	Base Image Naming Example
General Availability release	ESXi 8.0
Update release	ESXi 8.0 U1, ESXi 8.0 U2, and so on
Patch release	ESXi 8.0 a, ESXi 8.0 b, and so on
Security patch release	ESXi 8.0 sa, ESXi 8.0 sb, and so on
Patch release after an Update release	ESXi 8.0 U1 a, ESXi 8.0 U2 sa, and so on

In vSphere 8.0, base images can contain components and VIBs applicable to the ESXi version running on a DPU device. So, a single base image can contain software updates for both ESXi on the server and ESXi on the DPU device.

Base images are hosted and available in the VMware online depot. Additionally, you can download an ESXi installer ISO file and an offline bundle (ZIP file) that contains the ESXi version from customerconnect.vmware.com.

What Is a Vendor Add-On?

OEMs pack one or multiple components into a unit called an add-on. You use vendor add-ons to customize an ESXi image with OEM content and drivers.

The add-on is a collection of components that does not represent a complete, bootable image. You cannot use vendor add-ons on their own. To customize an ESXi release, you must add the vendor add-on to an ESXi base image. The combination of a vendor add-on and an ESXi base image is practically identical to the custom image that OEMs provide.

When combined with an ESXi base image, the add-on can add, update, or remove components that are part of the ESXi base image. In the vSphere Client, for each add-on available in the depot, you can view the list of components that it adds to an ESXi base image. Similarly, you can find information about the components that it removes from a base image.

In addition to custom ISO images and offline bundles, OEMs release ZIP files that contain only the vendor add-on, that is, the delta between the custom image and the ESXi base image. OEMs can release such add-on ZIP files at their discretion. The introduction of the concept of add-ons decouples the release cycle of OEMs from the release cycle of VMware. As a result, you can update vendor add-ons independently of updating the ESXi version of your hosts. Also, the vendor add-on decouples the OEM customization from the VMware stock image. As a result, you can combine software components more freely.

In vSphere 8.0, vendor add-ons support the ESXi software that runs on DPU devices.

Bulletins and Components in the vSphere Lifecycle Manager Depot

vSphere Lifecycle Manager can consume both bulletins and components.

If you use baselines and baseline groups to manage hosts and clusters, vSphere Lifecycle Manager reads and lists the software updates that are available in the vSphere Lifecycle Manager depot as bulletins. You can find the list of available bulletins on the **Updates** tab in the vSphere Lifecycle Manager home view.

If you use vSphere Lifecycle Manager images to manage hosts and clusters, you can only work with components and the related notions of add-ons and base images. You can find the list of the components, add-ons, and ESXi base images on the **Image Depot** tab in the vSphere Lifecycle Manager home view.

vSphere Lifecycle Manager Baselines and Images

vSphere Lifecycle Manager enables you to manage ESXi hosts and clusters with images or baselines. vSphere Lifecycle Manager baselines and vSphere Lifecycle Manager images are different in their essence, the way they work, and the features they support.

You use vSphere Lifecycle Manager baselines and baseline groups to perform the following tasks.

- Upgrade and patch ESXi hosts.
- Install and update third-party software on ESXi hosts.

You use vSphere Lifecycle Manager images to perform the following tasks.

- Install a desired ESXi version on all hosts in a cluster or on a standalone host.
- Install and update third-party software on all ESXi hosts in a cluster or on a standalone host.
- Update and upgrade the ESXi version on all hosts in a cluster or on a standalone host.
- Update the firmware of all ESXi hosts in a cluster or of a standalone host.
- Generate recommendations and use a recommended image for your cluster or host.
- Check the hardware compatibility of hosts and clusters against the VMware Compatibility Guide and the vSAN Hardware Compatibility List.

vSphere Lifecycle Manager Images

You use vSphere Lifecycle Manager images to apply software and firmware updates to the ESXi hosts in a cluster. Using a single image to manage all hosts in a cluster ensures cluster-wide host image homogeneity.

You can use various methods and tools to deploy ESXi hosts and maintain their software lifecycle. For example, you can upgrade hosts by using VMware vSphere® ESXi™ Image Builder CLI, esxcli, vSphere Auto Deploy. The different deployment and upgrade choices involve different workflows and require you to use different ESXi image formats. When you use vSphere Lifecycle Manager images, you follow one workflow and use the same ESXi image format for all software lifecycle-related operations: install, upgrade, update, and patching, which significantly simplifies the lifecycle management process.

Understanding vSphere Lifecycle Manager Images

A vSphere Lifecycle Manager image represents a desired software specification to be applied to all hosts in a cluster. When you set up a vSphere Lifecycle Manager image, you can define the full software stack that you want to run on the hosts in a cluster: the ESXi version, additional VMware software, vendor and third-party software, for example firmware and drivers.

A vSphere Lifecycle Manager image can consist of the following four elements:

- **ESXi base image**
The base image contains an image of VMware ESXi Server and additional components, such as drivers and adapters that are necessary to boot a server. The base image is the only mandatory element in a vSphere Lifecycle Manager image. All other elements are optional.
- **Vendor add-on**
The vendor add-on is a collection of software components that OEMs create and distribute. The vendor add-on can contain drivers, patches, and solutions.
- **Firmware and drivers add-on**
The firmware and drivers add-on is a special type of vendor add-on designed to assist in the firmware update process. The firmware and drivers add-on contains firmware for a specific server type and corresponding drivers. To add a firmware and drivers add-on to your image, you must install the hardware support manager plug-in provided by the hardware vendor for the hosts in the respective cluster.
- **Independent components**
The component is the smallest discrete unit in an image. The independent components that you add to an image contain third-party software, for example drivers or adapters.

You can set up a vSphere Lifecycle Manager image for a cluster during the creation of the cluster. Alternatively, for existing clusters that you manage with vSphere Lifecycle Manager baselines, you can switch from using baselines to using images at a later time.

Note:

If you switch to using images, you cannot revert to using baselines for that cluster. You can only move the hosts to a cluster that uses baselines.

The Desired State Model

The concept of images that vSphere Lifecycle Manager introduces is based on the Desired State model for managing ESXi hosts and clusters.

The desired state of an ESXi host represents both the target software and target configuration for the host as opposed to the software and configuration that it currently runs. The Desired State model is the idea of managing hosts and clusters by defining and applying a desired state instead of listing and following steps to change the current state.

vSphere Lifecycle Manager Baselines and Baseline Groups

You use baselines and baseline groups to update and upgrade the ESXi hosts in your environment. To start managing a cluster with baselines and baseline groups, you must skip setting up an image during the creation of the cluster.

Baselines

A baseline is a grouping of multiple bulletins. You can attach a baseline to an ESXi host and check the compliance of the host against the associated baseline.

Baselines can be classified according to different criteria.

- Depending on the type of content, baselines are patch baselines, extension baselines, and upgrade baselines. Patch and extension baselines contain bulletins of the respective kind. Upgrade baselines contain ESXi images.
- Depending on how the update content is selected, baselines are fixed and dynamic.
- Depending on how they are created and managed, baselines are predefined, recommendation, or custom baselines.

Baseline Groups

A baseline group is a collection of non-conflicting baselines. You can attach the entire baseline group to an inventory object to check the compliance status of the object against all the baselines in the group as a whole.

You can combine custom baselines with any of the predefined baselines to create baseline groups.

Host baseline groups can contain a single upgrade baseline, and various patch and extension baselines.

To update or upgrade ESXi hosts by using baselines or baseline groups, you must first attach the baselines or baseline groups to an inventory object.

Although you can attach baselines and baseline groups to individual objects, a more efficient method is to attach them to container objects, such as folders, vApps, clusters, and data centers. Individual vSphere objects inherit baselines attached to the parent container object. Removing an object from a container removes the inherited baselines from the object.

For more information about creating and managing baselines and baseline groups, see [Types of vSphere Lifecycle Manager Baselines and Baseline Groups](#).

What Is the Difference Between vSphere Lifecycle Manager Images and Baselines

A vSphere Lifecycle Manager baseline is a collection of bulletins. A vSphere Lifecycle Manager image is a combination of components, vendor add-on, and an ESXi base image. Some differences exist between the operations that you can perform on clusters that you manage with a vSphere Lifecycle Manager image and the operations on clusters that you manage with vSphere Lifecycle Manager baselines.

Operation	Baselines	Images
Distribution	Bulletins are distributed through online depots and as offline bundles. You can import and use ISO images to create upgrade baselines.	Base image, vendor add-ons, and components are distributed through online depots and as offline bundles. You cannot use ISO images to set up a vSphere Lifecycle Manager image for a cluster.
Validation	Not supported. You do not validate a baseline before applying the updates to the hosts. You can only perform a remediation pre-check.	Supported. You can validate a vSphere Lifecycle Manager image to check if it is applicable to all hosts in the cluster. You can also perform a remediation pre-check.
Import/Export	You can create a custom baseline and attach it to different objects in the same vCenter Server instance. You cannot export baselines and distribute them across vCenter Server instances.	You can export an image and use it to manage other clusters in the same or in a different vCenter Server instance. Images are portable across vCenter Server instances. You can export an image as an ISO or JSON file, but you can only import images that are in a JSON format.
Compliance checks	With baselines, you can check the compliance of an object against a single or against multiple baselines.	With vSphere Lifecycle Manager images, you can check the compliance of the hosts against a single image. To check the compliance against another image, you must first set up the new image.
Staging	Supported. You can stage updates to the hosts before actually installing them.	Supported. You can stage an image to a cluster or a host in the cluster before remediation. Staging reduces ESXi downtime.

Operation	Baselines	Images
Remediation	With vSphere Lifecycle Manager baselines, you can remediate an object against a single baseline or against multiple baselines. So, with a single operation, you can patch and upgrade a host. However, vSphere Lifecycle Manager baselines list the updates to be applied to hosts, but the remediation result is not always predictable, because ESXi image on the hosts might change after remediation.	With vSphere Lifecycle Manager images, you can add, remove, or modify the components in the image that you use for a cluster. When you remediate the hosts against the new image, all modified components are applied to the host. So, you can upgrade and patch a host with a single remediation operation. vSphere Lifecycle Manager images define the precise image to be applied to the hosts after remediation. No deviation from the defined image is possible after remediation. vSphere Lifecycle Manager does not allow solutions to push VIBs to the hosts.
Firmware updates	Not supported.	Supported. With vSphere Lifecycle Manager images, firmware updates are carried out through firmware and drivers add-ons, which you add to the image that you use to manage a cluster. Updating firmware with images requires an OEM-provided hardware support manager plug-in, which integrates with vSphere Lifecycle Manager.
Hardware compatibility checks	Not supported.	Supported You can check the hardware compatibility of the hosts in a cluster against the VMware Compatibility Guide (VCG). You can also check the compatibility of all hosts in a vSAN-enabled cluster against the vSAN Hardware Compatibility List (vSAN HCL).
Software recommendations	Limited support. Software recommendations are only available for vSAN clusters in the form of recommendation baselines.	Supported. Based on the hardware of the hosts in the cluster, you get recommendations about available and applicable ESXi updates or upgrades.
vCenter Server /Datacenter-level operations	With vSphere Lifecycle Manager baselines, you can trigger any of the main operations at the vCenter Server or data center level.	With vSphere Lifecycle Manager images, you cannot operate at a vSphere Lifecycle Manager or data center level.
Virtual machine management	You can upgrade the VMware Tools and virtual hardware versions of the virtual machines in a cluster that you manage with vSphere Lifecycle Manager baselines.	You can upgrade the VMware Tools and virtual hardware versions of the virtual machines in a cluster that you manage with vSphere Lifecycle Manager images.
Update Manager Download Service (UMDS)	Supported.	Supported.
Remote Office/Branch Office (ROBO) support	Not provided. Although no specific optimization exists for ROBO deployments, you can still use baselines and baseline groups with ROBO clusters.	Provided. With vSphere Lifecycle Manager images, you can set up a local depot and use it in ROBO environments. For more information, see Manage Depot Overrides for a Cluster or a Standalone Host .
REST APIs	Not available.	Available.

Operation	Baselines	Images
Support for DPU devices	Not provided. You cannot use baselines for clusters that contain hosts with DPU devices.	Provided To manage a cluster of DPU-backed hosts, you can only use vSphere Lifecycle Manager images. All vSphere Lifecycle Manager operations that work for updating the ESXi version on the host also deliver updates to the ESXi version on the DPU device on the host.

What Are the Requirements for Using vSphere Lifecycle Manager

Depending on whether you want to use baselines or images for software lifecycle management, you must comply with a different set of requirements. To achieve your goals, you must also know the specifics in behavior and limitations that vSphere Lifecycle Manager has.

Table 119: Requirements for Using vSphere Lifecycle Manager

Scenario	Requirements
Using a single image to manage a cluster.	- All ESXi hosts in the cluster must be of version 7.0 and later. - All ESXi hosts in the cluster must be stateful. A stateful install is one in which the host boots from a disk. - All ESXi hosts in the cluster must be from the same vendor and with identical hardware. Different generations and models of servers need different software drivers, which implies that you must set up different vSphere Lifecycle Manager images to manage each generation or model. However, with vSphere Lifecycle Manager, you use one single image for the entire cluster. Also, vSphere Lifecycle Manager does not detect and handle the hardware differences between the hosts in the cluster. You can use a vSphere Lifecycle Manager image to manage a heterogeneous cluster only if the vSphere Lifecycle Manager image for the cluster includes vendor customization, for example a vendor or firmware add-on, that can address and handle the hardware differences across the hosts in the cluster, which is a rare scenario. - If you want to use DPU devices for network acceleration, all hosts in the cluster must have a DPU device from the same vendor and of the same model. You cannot add a host without a DPU device to a cluster of DPU-backed hosts. - The cluster must include only integrated solutions. For example: - VMware vSAN™ - VMware vSphere® High Availability (HA) - vSphere Supervisor - NSX
Using a single image to manage a standalone host.	- The standalone ESXi host must be of version 7.0 and later. - The standalone ESXi host must be stateful. A stateful install is one in which the host boots from a disk.

Scenario	Requirements
	The host must include only integrated solutions. For example, the witness host for a vSAN stretch cluster must be a standalone host. For more information, see the <i>vSAN Planning and Deployment</i> documentation and vSAN Clusters and vSphere Lifecycle Manager.
Using baselines and baseline groups to manage a cluster or a host.	- To use baselines for ESXi host patching operations, vSphere Lifecycle Manager works with ESXi6.7, ESXi7.0, and ESXi8.0. - To use baselines for ESXi host upgrade operations, vSphere Lifecycle Manager works with ESXi6.7, ESXi7.0, and their respective Update releases.
Switching from using baselines to using a single image to manage a cluster or a host.	- The cluster or the standalone host must meet the requirements for using an image. - The cluster or the standalone host must be eligible for the transition. For more information about the Check cluster's eligibility to be managed with a single image and Check host's eligibility to be managed with a single image tasks, see Cluster or Standalone Host Eligibility to Use vSphere Lifecycle Manager Images.
Upgrading virtual machine hardware and VMware Tools	For VMware Tools and virtual machine hardware upgrade operations, vSphere Lifecycle Manager works with ESXi6.7, ESXi7.0, and ESXi8.0.

Required Privileges for Using vSphere Lifecycle Manager and vSphere Configuration Profiles

To configure vSphere Lifecycle Manager settings and to use successfully vSphere Lifecycle Manager baselines and images, you must have the proper privileges. Similarly, you need the proper privileges to use vSphere Configuration Profiles.

You can assign vSphere Lifecycle Manager and vSphere Configuration Profiles privileges to different roles in the vSphere Client.

vSphere Lifecycle Manager Privileges For Using Images

When you use vSphere Lifecycle Manager images, you need a different set of privileges for each task.

Table 120: VMware vSphere vSphere Lifecycle Manager Privileges For Using Images

Task	Privilege in the vSphere Client	Privilege in the API
Set Up Image	VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Read VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Write VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read	- VcIntegrity.lifecycleSoftwareSpecification.Read - VcIntegrity.lifecycleSoftwareSpecification.Write - VcIntegrity.lifecycleSettings.Read

Task	Privilege in the vSphere Client	Privilege in the API
Import Image	- VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Write - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read - VMware vSphere Lifecycle Manager > Upload File > Upload upgrade images and offline bundles	- VcIntegrity.lifecycleSoftwareSpecification.Read - VcIntegrity.lifecycleSoftwareSpecification.Write - VcIntegrity.lifecycleSettings.Read - VcIntegrity.FileUpload.com.vmware.vcIntegrity.ImportFile
Export Image	VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Read	VcIntegrity.lifecycleSoftwareSpecification.Read
Edit Image	- VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Write - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read	- VcIntegrity.lifecycleSoftwareSpecification.Read - VcIntegrity.lifecycleSoftwareSpecification.Write - VcIntegrity.lifecycleSettings.Read
Work with Recommendations	- VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Write - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read	- VcIntegrity.lifecycleSoftwareSpecification.Read - VcIntegrity.lifecycleSoftwareSpecification.Write - VcIntegrity.lifecycleSettings.Read
Work with Depot	- VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Write	- VcIntegrity.lifecycleSoftwareSpecification.Read - VcIntegrity.lifecycleSoftwareSpecification.Write
Manage Depot Overrides	- VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Write	- VcIntegrity.lifecycleSoftwareSpecification.Read - VcIntegrity.lifecycleSoftwareSpecification.Write
Check Compliance	VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Read	VcIntegrity.lifecycleSoftwareSpecification.Read
Run Remediation Pre-Check	- VMware vSphere Lifecycle Manager > Lifecycle Manager: General Privileges > Read - VMware vSphere Lifecycle Manager > ESXi Health Perspectives > Read	- VcIntegrity.lifecycleGeneral.Read - VcIntegrity.lifecycleHealth.Read

Task	Privilege in the vSphere Client	Privilege in the API
	VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Remediation Privileges > Read	VcIntegrity.lifecycleSoftwareRemediation.Read
Remediate Against an Image	- VMware vSphere Lifecycle Manager > Lifecycle Manager: General Privileges > Read - VMware vSphere Lifecycle Manager > ESXi Health Perspectives > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Remediation Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Remediation Privileges > Write - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read	- VcIntegrity.lifecycleGeneral.Read - VcIntegrity.lifecycleHealth.Read - VcIntegrity.lifecycleSoftwareRemediation.Read - VcIntegrity.lifecycleSoftwareRemediation.Write - VcIntegrity.lifecycleSettings.Read
Edit Remediation Settings	- VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Write	- VcIntegrity.lifecycleSettings.Read - VcIntegrity.lifecycleSettings.Write
Update Firmware	- VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Read - VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Privileges > Write - VMware vSphere Lifecycle Manager > Lifecycle Manager: Settings Privileges > Read	- VcIntegrity.lifecycleSoftwareSpecification.Read - VcIntegrity.lifecycleSoftwareSpecification.Write - VcIntegrity.lifecycleSettings.Write
View Hardware Compatibility	VMware vSphere Lifecycle Manager > Lifecycle Manager: Hardware Compatibility Privileges > Access Hardware Compatibility	VcIntegrity.HardwareCompatibility.Read
Stage an Image	VMware vSphere Lifecycle Manager > Lifecycle Manager: Image Remediation Privileges > Write	VcIntegrity.lifecycleSoftwareRemediation.Write

For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation.

vSphere Lifecycle Manager Privileges For Using Baselines

Each of the vSphere Lifecycle Manager privileges that you need to use baselines and baseline groups covers a distinct functionality.

Table 121: VMware vSphere Lifecycle Manager Privileges For Using Baselines

Task	Privilege in the vSphere Client	Privilege in the API	Description
Configure	Configure > Configure Service	VcIntegrity.General.com.vmware.vcIntegrity.Configure	Configure the vSphere Lifecycle Manager service and the scheduled patch download task.
Manage Baseline	Manage Baselines > Attach Baseline	VcIntegrity.Baseline.com.vmware.vcIntegrity.AssignBaselines	Attach baselines and baseline groups to objects in the vSphere inventory.
	Manage Baselines > Manage Baselines	VcIntegrity.Baseline.com.vmware.vcIntegrity.ManageBaselines	Create, edit, or delete baselines and baseline groups.
Manage Patches and Upgrades	Manage Patches and Upgrades > Remediate to Apply Patches, Extensions, and Upgrades	VcIntegrity.Updates.com.vmware.vcIntegrity.Remediate	Remediate virtual machines and hosts to apply patches, extensions, or upgrades. In addition, this privilege allows you to view the compliance status of objects.
	Manage Patches and Upgrades > Scan for Applicable Patches, Extensions, and Upgrades	VcIntegrity.Updates.com.vmware.vcIntegrity.Scan	Scan virtual machines and hosts to search for applicable patches, extensions, or upgrades.
	Manage Patches and Upgrades > Stage Patches and Extensions	VcIntegrity.Updates.com.vmware.vcIntegrity.Stage	Stage patches or extensions to hosts. In addition, this privilege allows you to view the compliance status of hosts.
	Manage Patches and Upgrades > View Compliance Status	VcIntegrity.Updates.com.vmware.vcIntegrity.ViewStatus	View baseline compliance information for an object in the vSphere inventory.
Upload File	Upload File > Upload upgrade images and offline bundles	VcIntegrity.FileUpload.com.vmware.vcIntegrity.ImportFile	Upload upgrade images and offline patch bundles.

For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation.

Required Privileges for Using vSphere Configuration Profiles

When you use vSphere Configuration Profiles, you need a different set of privileges for each task.

Table 122: Required Privileges For Using vSphere Configuration Profiles

Task	Privilege in the vSphere Client	Privilege in the API	Description
Create Cluster with vSphere Configuration Profiles enabled	Host > Inventory > Create Cluster	Host.Inventory.CreateCluster	Create a cluster along with its initial configuration.

Task	Privilege in the vSphere Client	Privilege in the API	Description
View Host Settings in the Cluster Configuration	VMware vSphere Lifecycle Manager > Desired Configuration Management Privileges > Read-only access to desired configuration management platform	VcIntegrity.ClusterConfiguration.View	View the draft configuration document, current configuration settings, compliance reports, and pre-check reports.
Check Cluster Compliance			
View Compliance Results			
Run Remediation Pre-Check			
View Remediation Pre-Check Results			
View Draft Pre-Check Results			
Remediate a Cluster Against the Desired Cluster Configuration	VMware vSphere Lifecycle Manager > Desired Configuration Management Privileges > Remediate cluster to the desired configuration.	VcIntegrity.ClusterConfiguration.Remediate	Remediate a cluster against the draft configuration or transition a cluster to the desired configuration.
Transition to vSphere Configuration Profiles			
Export a Configuration or Configuration Schema	VMware vSphere Lifecycle Manager > Desired Configuration Management Privileges > Export desired cluster configuration	VcIntegrity.ClusterConfiguration.Export	Export the desired cluster configuration document.
Extract a Configuration Document from a Reference Host	VMware vSphere Lifecycle Manager > Desired Configuration Management Privileges > Modify desired cluster configuration.	VcIntegrity.ClusterConfiguration.Modify	Modify the desired cluster configuration by creating a draft configuration, extracting a configuration document from a reference host, or importing a configuration.
Create a Draft Configuration for a Cluster			
Import a Configuration			

The vSphere Lifecycle Manager Depot

The vSphere Lifecycle Manager depot is the source of software updates for vSphere Lifecycle Manager. Conceptually, the vSphere Lifecycle Manager depot represents all software available for consumption to vSphere Lifecycle Manager.

The vSphere Lifecycle Manager depot is a local depot on the vCenter Server machine. This local depot contains software updates downloaded from the online depots that you configure vSphere Lifecycle Manager to use. You can also manually import updates into the vSphere Lifecycle Manager depot. You can work with vSphere Lifecycle Manager only if the vSphere Lifecycle Manager depot contains software packages, for example, ESXi base images, vendor add-ons, third-party components, and legacy patches and updates.

How Do Software Updates Go to the vSphere Lifecycle Manager Depot

You control how the vSphere Lifecycle Manager depot gets populated with software. You can configure vSphere Lifecycle Manager to download updates from online depots. In environments without access to the Internet, you can use the Update Manager Download Service (UMDS) and configure vSphere Lifecycle Manager to download software from a UMDS-created shared repository. Alternatively, you can use offline depots to import updates into the vSphere Lifecycle Manager depot manually.

What Is an Online Depot and What Is the Default VMware Online Depot

VMware, OEMs, and third-party software providers create and ship their software updates in a well-defined `xml` file structure, which is called a software depot. Depending on the distribution format of the updates, software depots can be online and offline.


An online depot is the hosted version of the software updates that VMware, OEMs, and third-party software providers ship. You enable vSphere Lifecycle Manager to access an online depot by providing a URL to that depot. vSphere Lifecycle Manager is preconfigured to download updates from one online depot, the default VMware online depot.

The VMware Online Depot

The default online depot that VMware provides hosts ESXi base images, vendor add-ons, ESXi-compatible I/O device drivers certified by VMware, and async VMware Tools releases. All software that you need to install, update, or customize the ESXi version of your hosts is available in the official VMware online depot. You don't need to work with separate online depots to access OEM and third-party software updates.

Note:

The default online depot URLs are now unique for each customer. To update the default URLs of the online depot, see the following KB article: <https://knowledge.broadcom.com/external/article/390121>.



Note: Firmware updates are not hosted in the VMware depot. To perform firmware updates, you must install the hardware support manager plug-in that your hardware vendor provides. The plug-in gives you access to depots that contain the necessary firmware and related drivers updates.

You can configure vSphere Lifecycle Manager to also access and use other online depots besides the VMware online depot, for example third-party depots that contain additional components such as CIM modules. However, working with additional third-party depots and independent components is rarely necessary. In most cases, the vendor add-ons that are available in the official VMware depot provide full OEM customization for ESXi.

vSphere Lifecycle Manager downloads to the local vSphere Lifecycle Manager depot the content from all the online depots that you configure it to use.

Synchronization

Synchronization is the process during which vSphere Lifecycle Manager downloads the contents of the online depots that you configure it to use to the local vSphere Lifecycle Manager depot. During synchronization, only software metadata is downloaded. The actual payloads are downloaded when they are needed, for example during staging or remediation. When you deploy vCenter Server, vSphere Lifecycle Manager synchronizes with the official VMware online depot automatically. After the initial synchronization, you can schedule a download task to run at regular intervals or you can initiate a download task manually.

What Is an Offline Depot

An offline depot, also called an offline bundle, is a ZIP file that you download from the Internet or copy from a media drive and you save on a local or a shared network drive. You then can import the offline bundle to the vSphere Lifecycle Manager depot.

All software updates hosted in the official VMware online depot are also available as offline bundles, which you can download from the [Broadcom Support Portal Downloads](#) site. You can also download offline bundles from the VMware website or from the websites of third-party vendors.

In addition to distributing an `offline.zip` file, or offline bundle, and a custom ISO image, OEMs distribute an `Add-on.zip` file that contains the delta between the OEM custom image and the respective base image that VMware provides. For more information about OEM add-ons, see [Software Packaging Units That vSphere Lifecycle Manager Can Consume](#).

Import

Import is the operation through which you upload the contents of an offline bundle into the vSphere Lifecycle Manager depot. During an import operation, both the software metadata and the actual payloads are downloaded into the vSphere Lifecycle Manager depot. In addition to importing offline bundles, you can also import ISO images to the vSphere Lifecycle Manager depot. You can then use the ISO images to create upgrade baselines. You cannot use an ISO image for clusters that you manage with a vSphere Lifecycle Manager image.

Software Deliverables and Their Corresponding Distribution Formats

Software Vendor	Software Deliverable	Software Distribution Format
VMware	Base images	- The default VMware online depot - Offline bundle - ISO image
OEMs	Add-ons	- The default VMware online depot - Offline bundle - Add-on ZIP file - ISO image

Software Vendor	Software Deliverable	Software Distribution Format
Third-party software providers	Components	For device drivers that are certified by VMware: • The default VMware online depot • Offline bundle For other third-party software that is verified and certified by OEMs, for example I/O filters, CIM module: • An online depot • Offline bundle

vSphere Lifecycle Manager Download Sources

The vSphere Lifecycle Manager download source is the place from which vSphere Lifecycle Manager downloads software updates.

If you configure vSphere Lifecycle Manager to use the Internet as its download source, then the download sources are practically all online depots that you use with vSphere Lifecycle Manager.

Note:

The URLs of the official VMware online depot are now unique for each customer. To update the download sources URLs, see the following KB articles: <https://knowledge.broadcom.com/external/article/390121>.

You can also configure vSphere Lifecycle Manager to use a UMDS shared repository as its download source. Using a UMDS repository is appropriate in vCenter Server deployments without access to the Internet. When you configure vSphere Lifecycle Manager to use a UMDS repository, the synchronization of update metadata is not triggered immediately. The metadata is downloaded according to the configured download schedule or when you initiate the download. When you use a UMDS repository as a download source for vSphere Lifecycle Manager, only the metadata of the updates is downloaded and stored in the vSphere Lifecycle Manager depot. The actual payload is downloaded during staging or remediation.

Note:

The default online depot URLs are now unique for each customer. To update the default download URLs in the UMDS configuration files, see the following KB article: <https://knowledge.broadcom.com/external/article/390123>.

Depot Overrides

In ROBO scenarios, you can configure vSphere Lifecycle Manager to use a local depot with updates for a particular cluster instead of the depots that all clusters in that vCenter Server instance use by default.

For more information, see [Manage Depot Overrides for a Cluster or a Standalone Host](#).

Browsing the vSphere Lifecycle Manager Depot

You can use the vSphere Client to view and browse the contents of the vSphere Lifecycle Manager depot.

You can view the vSphere Lifecycle Manager depot in the vSphere Lifecycle Manager home view. The contents of the vSphere Lifecycle Manager depot are displayed on three different tabs: **Image Depot**, **Updates**, and **Imported ISOs**.

Image Depot

On the **Image Depot** tab, you can view all VMware base images, vendor add-ons, and components that are available in the vSphere Lifecycle Manager depot.

You can use the **ESXi Versions**, **Vendor Addons**, and **Component** links at the top of the pane for easier navigation through the lists.

The **ESXi Versions** list contains all base images available in the depot together with information about the version, release date, and category for each image. When you select an image from the list, an information panel appears on the right. The panel displays a list of all components that the base image applies to a host upon remediation.

The **Vendor Addons** list contains all vendor add-ons available in the depot together with information about the version, release date, and category for each add-on. When you select an add-on from the list, an information panel appears on the right. The panel displays information about the components that the add-on applies to the host and the components that the add-on removes from a host upon remediation.

The **Component** list contains all components that are available in the depot together with information about the version, release date, and category for each component. When you select a component from the list, an information panel appears on the right. The panel displays information about the VIBs that the component contains.

You can filter the **Component** list so that it displays only independent components or all components available in the vSphere Lifecycle Manager depot. Independent components are components that are not part of a vendor add-on.

You use the ESXi images, vendor add-ons, and components visible on the **Image Depot** tab to set up images that you can use to manage hosts in clusters collectively.

Updates

On the **Updates** tab, you can see all components available in the vSphere Lifecycle Manager depot as bulletins. You can use the **Filter by Baseline** drop-down menu to view only the bulletins that are part of a particular baseline.

When you select a bulletin from the list, additional information appears below the bulletins list. In the bottom pane, you see information about the baselines that include the selected bulletin.

You use the bulletins visible on the **Updates** tab to create baselines and baseline groups.

Because the official VMware depot hosts certified partner content in addition to VMware content, the **Updates** tab displays a broader set of OEM bulletins, for example vendor add-ons and VMware-certified device drivers. Some of these bulletins might have dependencies that must be pulled into the baselines that you create, so that the remediation against those baselines is successful. As a best practice, always consult the KB article for an individual bulletin to find information about its deployment specifics and required dependencies before including the bulletin in your baselines.

Starting with vSphere 7.0, some changes are also introduced in the way VMware content is packaged. As a result, you might see additional bulletins on the **Updates** tab at patch and update releases. Those bulletins are usually of the Enhancement or BugFix category. When you include those bulletins in a baseline, you might need to also include a base ESXi bulletins in that baseline. As a best practice, to ensure successful application of patches and updates, always include the appropriate rollup bulletin into your baselines. You can use the **Show only rollup updates** toggle switch that is on the **Updates** tab to filter the list of bulletins.

Imported ISOs

On the **Imported ISOs** tab, you can see the ISO images that you import and make available to vSphere Lifecycle Manager.

You use the ISO images visible on the **Imported ISOs** tab to create upgrade baselines. You cannot use an ISO image for clusters configured to use a single vSphere Lifecycle Manager image.

Note: ISO images are not distributed through any online or offline depot, they are a separate software distribution format. As a result, they cannot become available in the vSphere Lifecycle Manager depot through synchronization or the regular import operation that you perform to import offline bundles (ZIP files) to the depot. To make an ISO image available to vSphere Lifecycle Manager, you must trigger the **Import ISO** operation. For more information, see [Import an ISO Image to the vSphere Lifecycle Manager Depot](#)

Working With the vSphere Lifecycle Manager Depot

The vSphere Lifecycle Manager synchronizes to its download source regularly and automatically. But you can change the contents of the vSphere Lifecycle Manager depot by manually importing or deleting updates or ISO images or by triggering synchronization at a convenient time.

Import Updates to the vSphere Lifecycle Manager Depot

You can use an offline bundle in ZIP format and import updates to the vSphere Lifecycle Manager depot manually. When you import offline bundles, you add both the update metadata and actual payload to the vSphere Lifecycle Manager depot.

- Verify that the updates that you import are in ZIP format.
- Required privileges: **VMware vSphere Lifecycle Manager > Upload File > Upload File**.

You use the import option to populate the vSphere Lifecycle Manager depot with updates from an offline bundle. Apart from the legacy patches and extensions, an offline bundle can also contain an ESXi base image, a vendor add-on, or third-party software, for example, asynchronous drivers specific to the OEM hardware requirements. For more information about base images, vendor add-ons, and components, see [Software Packaging Units That vSphere Lifecycle Manager Can Consume](#).

If you want to use vSphere Lifecycle Manager baselines, you can import offline bundles that contain patches and extensions for hosts that run ESXi6.7 and later. In that case, you can use the contents of the offline bundle only for host patching operations. If you import an OEM offline bundle that contains an ESXi image of a version earlier than 7.0, you cannot use the image for upgrade operations. To create upgrade baselines, you need an ISO image. For more information, see [Import an ISO Image to the vSphere Lifecycle Manager Depot](#).

If you want to use vSphere Lifecycle Manager images, you can import offline bundles that contain software for hosts that run ESXi 7.0 and later. In that case, you can use the contents of the offline bundle to set up vSphere Lifecycle Manager images, which you can use to upgrade a standalone host or multiple ESXi hosts collectively.

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. Select **Actions > Import Updates** at the top of the vSphere Lifecycle Manager home view.

The **Import Updates** dialog box opens.
3. Enter a URL or browse to an offline bundle in ZIP format on your local machine.

If the upload fails, check whether the structure of the ZIP file is correct and whether the vSphere Lifecycle Manager network settings are set up correctly.
4. Click **Import**.

The Import updates task appears in the **Recent Tasks** pane.

You imported updates to the vSphere Lifecycle Manager depot. vSphere Lifecycle Manager automatically generates new image recommendations for the clusters or hosts that already have generated recommended images. However, if the imported updates are solution components only, vSphere Lifecycle Manager does not generate new recommendations automatically.

You can view the imported patches and extension on the **Updates** tab in the vSphere Lifecycle Manager home view.

You can view the imported ESXi images, vendor add-ons, and additional components on the **Image Depot** tab in the vSphere Lifecycle Manager home view.

Import an ISO Image to the vSphere Lifecycle Manager Depot

You import ESXi images in ISO format to the vSphere Lifecycle Manager local depot, so that you can create upgrade baselines, which you use for host upgrade operations.

Required privileges: **VMware vSphere Lifecycle Manager > Upload File**

You can use ESXi.iso images to upgrade ESXi6.7.x hosts and ESXi7.0.x hosts to ESXi8.0.

With vSphere Lifecycle Manager 8.0, you cannot perform ESXi upgrades to version 7.0 or 6.7.

ISO images can only be used with vSphere Lifecycle Manager baselines. You cannot use an ISO image to upgrade the hosts in a cluster that uses a single image.

To upgrade hosts, use the ESXi installer image distributed by VMware with the name format `VMware-VMvisor-Installer-7.0.0-build_number.x86_64.iso` or a custom image created by using vSphere ESXi Image Builder. You can also use ISO images created and distributed by OEMs.

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Imported ISOs** tab, click **Import ISO**
3. In the **Import ISO** dialog box, select an image.
 - Click the **Browse** button to import an ESXi image from your local system.
 - Enter an URL address to import an ESXi image that is not on your local system.

Local images are imported immediately, whereas importing images from a URL takes some time.
4. Click **Import**.

The ISO image that you uploaded appears in the list of images. You can view information about the ESXi image, such as product, version, and build details, vendor, acceptance level, and creation date.

Create a host upgrade baseline.

Delete an ISO Image from the vSphere Lifecycle Manager Depot

If you do not need an ESXi image, you can delete it from the vSphere Lifecycle Manager depot.

- Verify that the ISO image that you want to delete is not part of any baseline. You cannot delete images that are included in a baseline.

- Delete any baseline that contains the ISO image that you want to delete.

Unlike components and bulletins, which you cannot delete from the vSphere Lifecycle Manager depot, the ISO images that you import in the depot can be deleted when you no longer need them .

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Imported ISOs** tab, select an image from the list and click **Delete**.

Note: If you try to delete an ESXi image that is used in a baseline, the operation fails with an error message.
3. Click **Yes** to confirm the deletion.

The ISO image is deleted and no longer available.

Synchronize the vSphere Lifecycle Manager Depot

Instead of waiting for the predefined download task to run as scheduled, you can update your local vSphere Lifecycle Manager depot immediately.

At regular configurable intervals, vSphere Lifecycle Manager downloads updates from the configured download sources. The download sources can be online depots or a UMDS-created shared repository.

Regardless of the download schedule, you can initiate synchronization between the vSphere Lifecycle Manager depot and the configured download sources. Similar to scheduled synchronization, when you initiate synchronization manually, vSphere Lifecycle Manager downloads software from all online depots that you configured it to use. For more information about configuring the vSphere Lifecycle Manager download sources, see [Configuring the vSphere Lifecycle Manager Download Sources](#).

During synchronization, vSphere Lifecycle Manager downloads only the update metadata, the actual payloads are downloaded during staging or remediation.

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. Select **Actions > Sync Updates** at the top of the vSphere Lifecycle Manager home view.

The Sync updates task appears in the **Recent Tasks** pane.

You downloaded updates to the vSphere Lifecycle Manager depot. vSphere Lifecycle Manager automatically generates new image recommendations for the clusters that already have generated recommended images. However, if the updates are related to downloading solution components only, vSphere Lifecycle Manager does not generate new recommendations automatically.

You can view the downloaded patches and extension on the **Updates** tab in the vSphere Lifecycle Manager home view.

You can view the downloaded ESXi images, vendor add-ons, and components on the **Image Depot** tab in the vSphere Lifecycle Manager home view.

Configuring the vSphere Lifecycle Manager Download Sources

You can configure vSphere Lifecycle Manager to download software updates for ESXi hosts either from the Internet or from a shared repository of UMDS data.

vSphere Lifecycle Manager downloads only the metadata and not the actual binary payload of the updates. Downloading the metadata saves disk space and network bandwidth. The availability of regularly updated metadata in the vSphere Lifecycle Manager depot lets you perform compliance checks on hosts at any time.

Whatever the download source, vSphere Lifecycle Manager downloads the following types of information:

- Metadata about all ESXi 6.x updates regardless of whether you have hosts of such versions in your environment.
- Metadata about all ESXi 7.x updates regardless of whether you have hosts of such versions in your environment.
- Patch recalls for ESXi 6.x hosts.

vSphere Lifecycle Manager supports the recall of patches for hosts that are running ESXi6.7 or later. A patch is recalled when it has problems or potential issues. After you scan the hosts in your environment, vSphere Lifecycle Manager alerts you if the recalled patch has been installed on any host. Recalled patches cannot be installed on hosts with vSphere Lifecycle Manager. vSphere Lifecycle Manager deletes all the recalled patches from the vSphere Lifecycle Manager depot. After a patch that fixes the problem is released, vSphere Lifecycle Manager downloads the new patch to its depot. If you have already installed the problematic patch, vSphere Lifecycle Manager notifies you that a fix is available and prompts you to apply the new patch.

Downloading host patches from the VMware website is a secure process.

- Download URLs for updates and patches are unique for each customer. For more information about the online depot URLs and how to edit the vSphere Lifecycle Manager configuration, see the following KB article: [390121](#).
- Patches are cryptographically signed with the VMware private keys. Before you try to install a patch on a host, the host verifies the signature. This signature enforces the end-to-end protection of the patch itself and can also address any concerns about downloading the patch.
- vSphere Lifecycle Manager downloads the patch metadata and patch binaries over SSL connections. vSphere Lifecycle Manager verifies both the validity of the SSL certificates and the common name in the certificates. The common name in the certificates must match the names of the servers from which vSphere Lifecycle Manager downloads the patches. vSphere Lifecycle Manager downloads the patch metadata and binaries only after successful verification of the SSL certificates.

Download Sources

If your deployment system is connected to the Internet, you can use the default settings and links for downloading updates to the vSphere Lifecycle Manager depot. You can also add URL addresses to download third-party software, for example drivers.

If your deployment system is not connected to the Internet, you can use a shared repository after downloading the upgrades, patches, and extensions by using Update Manager Download Service (UMDS).

For more information about UMDS, see [Installing, Setting Up, and Using the Update Manager Download Service](#).

The default configuration is for the vSphere Lifecycle Manager to download information directly from the Internet. However, you can change the download source at any time. Changing the download source from a shared repository to the Internet and the reverse is a change in the vSphere Lifecycle Manager configuration. The two options are mutually exclusive. You cannot download updates from the Internet and a shared repository at the same time.

By default, vSphere Lifecycle Manager is configured to use the official VMware online depot as a download source. When you deploy vCenter Server, synchronization to the official VMware depot is triggered automatically. When you change the default download source, synchronization to the new download source is not triggered automatically. The synchronization task runs as per its schedule. To download new data, you must run the **VMware vSphere Lifecycle Manager Update Download** task or trigger synchronization manually.

The **VMware vSphere Lifecycle Manager Update Download** task is a scheduled task that runs at regular intervals. You can change the schedule, and you can also trigger the **VMware vSphere Lifecycle Manager Update Download** task independently of its schedule.

If the **VMware vSphere Lifecycle Manager Update Download** task is running when you apply the new configuration settings, the task continues to use the old settings until it finishes. The next time the download task starts, vSphere Lifecycle Manager uses the new settings.

Using a Proxy Server

You cannot configure vSphere Lifecycle Manager to use a proxy server on its own. vSphere Lifecycle Manager uses the proxy settings of the vCenter Server instance where it runs.

Configure vSphere Lifecycle Manager to Use the Internet as a Download Source

If your deployment system is connected to the Internet, you can configure vSphere Lifecycle Manager to directly download ESXi images, vendor add-ons, and other components from the configured online depots to the local vSphere Lifecycle Manager depot.

Required privileges: **VMware vSphere Lifecycle Manager > Configure**

The Internet is the default download source for vSphere Lifecycle Manager. Downloading takes place at configurable regular intervals. To initiate downloading of updates regardless of the download schedule, see [Synchronize the vSphere Lifecycle Manager Depot](#).

Note:

The default download source URLs are now unique for each customer. To update the default download Internet sources, see the following KB article: <https://knowledge.broadcom.com/external/article/390121>.

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Settings** tab, select **Administration > Patch Setup**.
3. Click the **Change Download Source** button.

The **Change Download Source Type** dialog box opens.
4. Select the **Download patches directly from the Internet** option and click **Save**.

Configure vSphere Lifecycle Manager to Use a Shared Repository as a Download Source

You can configure vSphere Lifecycle Manager to use a shared repository as a source for downloading ESXi images, vendor add-ons, and additional components.

- Create a shared repository by using UMDS and host the repository on a Web server or a local disk. For detailed information about exporting the upgrades, update binaries, and update metadata in [Export the Downloaded Data](#).

- Verify that UMDS is of version compatible with the version of vSphere Lifecycle Manager that you are using. For more information about compatibility, see [Installing, Setting Up, and Using the Update Manager Download Service](#).
- Required privileges: **VMware vSphere Lifecycle Manager > Configure**.

You cannot use folders on a network drive as a shared repository. vSphere Lifecycle Manager does not download updates from folders on a network share in the Microsoft Windows Uniform Naming Convention form (such as `\Computer_Name_or_Computer_IP\Shared`), or on a mapped network drive (for example, `Z:\`).

The downloading of updates takes place at configurable regular intervals. To initiate downloading of updates regardless of the download schedule, see [Synchronize the vSphere Lifecycle Manager Depot](#).

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.
The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.

2. On the **Settings** tab, select **Administration > Patch Setup**.

3. Click the **Change Download Source** button.

The **Change Download Source Type** dialog box opens.

4. Select the **Download patches from a UMDS shared repository** option and enter a path or URL address to the shared repository.

For example, `C:\repository_path\`, `https://repository_path/` or `http://repository_path/`.

In these examples, *repository_path* is the path to the folder with the exported downloaded upgrades, patches, extensions, and notifications. In an environment where vSphere Lifecycle Manager does not have direct access to the Internet, but is connected to a physical machine that has access to the Internet, the folder can be on a Web server.

You can specify an HTTP or HTTPS address, or a location on the disk where vSphere Lifecycle Manager runs. HTTPS addresses are supported without any authentication.

5. Click **Save**.

vCenter Server validates the URL. You can use the path to the shared repository only when the validation is successful. If the validation fails, vSphere Lifecycle Manager reports a reason for the failure.

Important: If the updates in the folder that you specify are downloaded with a UMDS version that is not compatible with the vCenter Server version that you use, the validation fails and you receive an error message.

The shared repository is used as the main source for downloading software updates. Downloading from the repository is enabled by default.

Using a Folder or a Web Server as a Shared Repository

You can use a folder or a Web server as a shared repository.

- When you use a folder as a shared repository, *repository_path* is the path to the top-level directory that stores the patches and notifications exported from UMDS.
For example, use UMDS to export the patches and notifications to the `F:\` drive, which is a drive mapped to a plugged-in USB device on the physical machine where UMDS is installed. Then, plug in the USB device to the physical machine where vSphere Lifecycle Manager runs. The device is mapped as `E:\` and the folder to configure as a shared repository for vSphere Lifecycle Manager is `E:\`.
- When you use a Web server as a shared repository, *repository_path* is the path to the top-level directory on the Web server that stores the patches exported from UMDS.
For example, export the patches and notifications from UMDS to `C:\docroot\exportdata`. If the folder is configured on a Web server and is accessible from other physical machines at the URL

`https://umds_host_name/exportdata`, the URL to configure as a shared repository in vSphere Lifecycle Manager is `https://umds_host_name/exportdata`.

Run the VMware vSphere vSphere Lifecycle Manager Update Download Task

If you change the download source, you must run the VMware vSphere vSphere Lifecycle Manager Update Download task to download any new updates.

1. In the vSphere Client, navigate to a vCenter Server instance.
2. On the **Configure** tab, select **Scheduled Tasks**.
3. In the **Scheduled Tasks** pane, select the **VMware vSphere Lifecycle Manager Update Download** task and click **Run**.

You can see the running task listed in the **Recent Tasks** pane.

Add a New Download Source

You can add a new download source for vSphere Lifecycle Manager by providing a URL to the online depot or the UMDS repository that you want to use. Additionally, you can edit or delete a download source. You can also temporarily stop vSphere Lifecycle Manager from downloading updates from a particular download source.

Add a New Download Source

If you use the Internet as a download source for updates, you can add URL addresses to third-party online depots. vSphere Lifecycle Manager downloads software updates from all the online depots that you configured it to use. Update metadata are downloaded from the online depots to the local vSphere Lifecycle Manager depot.

Required privileges: **VMware vSphere Lifecycle Manager > Configure**

The default download source for vSphere Lifecycle Manager is the official VMware depot.

Note:

The default download source URLs are now unique for each customer. To update the default download Internet sources, see the following KB article: <https://knowledge.broadcom.com/external/article/390121>.

Starting with vSphere 7.0, the official VMware online depot also hosts vendor add-ons and VMware-certified device drivers. Unlike previous releases, all software that you need to install, update, or customize the ESXi version of your hosts is available in the official VMware online depot.

Downloading updates takes place at configurable regular intervals. To initiate the downloading of updates regardless of the download schedule, see [Synchronize the vSphere Lifecycle Manager Depot](#).

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Settings** tab, select **Administration > Patch Setup**.
3. In the **Patches downloaded from the Internet** pane, click **New**.

The **New Download Source** dialog box opens.

4. Enter a URL address to a new download source.

vSphere Lifecycle Manager supports both HTTP and HTTPS URL addresses. Use HTTPS URL addresses to download data securely. The URL addresses that you add must be complete and contain an `index.xml` file, which lists the vendor and the vendor index.

Note: The proxy settings that vSphere Lifecycle Manager uses are also applicable to third-party URL addresses.

5. Optional: Enter a short description for the download source.
6. Click **Save**.

The new location is added to the list of download sources and downloading from it is enabled by default.

Modify a Download Source

You can edit or delete a download source from the list of vSphere Lifecycle Manager download sources. You can also allow or disallow vSphere Lifecycle Manager to download updates from a particular download source.

Required privileges: **VMware vSphere Lifecycle Manager > Configure**

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Settings** tab, select **Administration > Patch Setup**.
3. In the **Patches downloaded from the Internet** pane, select a URL address from the list of download sources and select your task.
 - Click **Edit** to edit the source URL or the description for the selected download source.
 - Click **Enable** or **Disable** to allow or disallow downloading from the selected download source.
 - Click **Delete** to delete the selected download source.

Note: You cannot edit or delete the default VMware download source for ESXi updates. You can only allow or disallow vSphere Lifecycle Manager to use it for downloading update metadata from it.

Configure the vSphere Lifecycle Manager Automatic Download Task

Downloading host updates and related metadata is a predefined automatic process that you can modify. The automatic download task is enabled by default and starts immediately after you deploy vCenter Server. After the initial download, the task runs according to its schedule.

- Verify that the machine on which vSphere Lifecycle Manager runs has access to the Internet.
- Required privileges: **VMware vSphere Lifecycle Manager > Configure**.

The default schedule settings ensure frequent checks, but you can change the schedule if your environment requires you to adjust the frequency of the checks.

If you need the latest software updates, you might want to reduce the time interval between the checks for updates. By contrast, if you are not concerned about the latest updates, if you want to reduce the network traffic, or if you cannot access the update servers, you might want to increase the time interval between the checks for updates.

The automatic download of update metadata is enabled by default and the default task name is VMware vSphere vSphere Lifecycle Manager Update Download. You can change the configuration of the task.

Note:

The default download source URLs are now unique for each customer. To update the default download Internet sources, see the following KB article: <https://knowledge.broadcom.com/external/article/390121>.

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.
 The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Settings** tab, select **Administration > Patch Downloads**.
3. In the **Automatic Download Settings** pane, click the **Edit** button.
 The **Edit Settings for Automatic Patch Downloads** dialog box opens.
4. Select the **Download patches** check box and configure the schedule and settings for the download.
 To receive notification emails after the download task finishes, you must configure mail settings for the vSphere Client. For more information, see the *vCenter Server and Host Management* documentation.
5. Click **Save** to save your changes and close the dialog box.

The automatic download task runs according to the schedule you configured.

Installing, Setting Up, and Using the Update Manager Download Service

VMware vSphere Update Manager Download Service (UMDS) is an optional module of vSphere Lifecycle Manager. UMDS downloads software metadata, software binaries, and notifications that might not otherwise be available to vSphere Lifecycle Manager.

For security reasons and deployment restrictions, vSphere, including vSphere Lifecycle Manager, might be installed in a secured network that is disconnected from other local networks and the Internet. vSphere Lifecycle Manager requires access to patch information to function properly. If you are using such an environment, you can install UMDS on a computer that has Internet access. UMDS downloads upgrades, patch binaries, and patch metadata, and then you can export the downloads to a portable media drive so that they become accessible to vSphere Lifecycle Manager.

Note:

The URLs of the official VMware online depot are now unique for each customer. To update the download sources URLs, see the following KB articles: <https://knowledge.broadcom.com/external/article/390121> and <https://knowledge.broadcom.com/external/article/390123>.

If the server on which vCenter Server is installed has no Internet access, but is connected to a server that has Internet access, you can automate the export process and transfer files from UMDS to the vSphere Lifecycle Manager depot by using a Web server on the machine on which UMDS is installed.

UMDS is available for installation only on Linux-based operating systems. Installing UMDS on a Windows machine is no longer supported. Administrator access is not a requirement for downloading patches with UMDS that runs on Linux. The machine on which you install UMDS must have Internet access.

UMDS 8.0 supports patch recalls and notifications. A patch is recalled if the released patch has problems or potential issues. After you download patch data and notifications with UMDS, and export the downloads so that they become available to vSphere Lifecycle Manager, vSphere Lifecycle Manager deletes the recalled patches.

You can use UMDS to download updates that are packaged and distributed components, but it also works with legacy bulletins.

Compatibility Between UMDS and vSphere Lifecycle Manager

UMDS must be of the same version as vSphere Lifecycle Manager. For example, vSphere Lifecycle Manager 8.0 is compatible and can work only with UMDS 8.0. If you are using vSphere Lifecycle Manager of an Update release version, UMDS must be of the same Update release version.

Installing UMDS

In vSphere 8.0 release, the UMDS 8.0 is bundled with the vCenter Server appliance 8.0. You can use the UMDS bundle from the vCenter Server appliance to install UMDS 8.0 on a separate Linux-based system.

UMDS is a 64-bit application and requires a 64-bit Linux-based system.

Note: You cannot upgrade UMDS that runs on a Linux-based operating system. You can uninstall the current version of UMDS, perform a fresh installation of UMDS according to all system requirements, and use the existing patch store from the UMDS that you uninstalled.

Supported Linux-Based Operating Systems for Installing UMDS

The Update Manager Download Service (UMDS) can run on a limited number of Linux-based operating systems.

- VMware Photon™ 5.0
- Ubuntu 14.0.4
- Ubuntu 18.04
- Ubuntu 18.04 LTS
- Ubuntu 20.04 LTS
- Red Hat Enterprise Linux 7.4
- Red Hat Enterprise Linux 7.5
- Red Hat Enterprise Linux 7.7
- Red Hat Enterprise Linux 8.1

Note: When you use Red Hat Enterprise Linux 8.1, you must install the `libnsl` package version 2.28 or later on the system where UMDS is deployed. If the package is not present on the system, UMDS operations might fail with the following error:

```
Error while loading shared libraries: libnsl.so.1: cannot open shared object file: No such file or directory.
```

- Red Hat Enterprise Linux 8.3
- Red Hat Enterprise Linux 8.5
- Red Hat Enterprise Linux 8.6
- Red Hat Enterprise Linux 8.10
- Red Hat Enterprise Linux 9.0
- Red Hat Enterprise Linux 9.2

Install UMDS on a Linux OS

If the vCenter Server appliance 8.0 where vSphere Lifecycle Manager runs does not have access to the Internet, you can install UMDS on a Linux-based operating system to download update binaries and metadata.

- Verify you have administrative privileges on the Linux machine where you install the UMDS.

- Mount the ISO file of the vCenter Server appliance 8.0 to the Linux machine.
1. In the Linux machine, open the Command Shell.
 2. From the vCenter Server ISO that you mounted to the Linux machine, copy the `VMware-UMDS-8.0.1-build_number.tar.gz` file to the Linux machine.
 3. Unarchive the `VMware-UMDS-8.0.1-build_number.tar.gz` file by running `tar -xvzf VMware-UMDS-8.0.1-build_number.tar.gz` and navigate to the newly extracted directory `/vmware-umds-distrib`.
For example, if you unarchived the `VMware-UMDS-8.0.1-build_number.tar.gz` file, to a directory you created with the name `umds`, your navigation path is `/umds/vmware-umds-distrib`.
 4. Run the file UMDS installation script.
The script has the following file name: `vmware-install.pl`.
 5. Read and accept the EULA.
 6. Select a directory where to install the UMDS.
The default UMDS installation location is `usr/local/vmware-umds`.
 7. Enter the UMDS proxy settings.
You can also change proxy configuration after you install UMDS by using the following command:

```
vmware-umds -S --proxy <proxyAddress:port>
```
 8. Select a directory where to store the patches.
Important: The patch store directory must be different from the UMDS installation directory. The default patch store location is `/var/lib/vmware-umds`.

UMDS is installed.

Uninstall UMDS from a Linux OS

To use the latest UMDS version on your Linux-based system, you must first uninstall the current version of UMDS. No direct upgrade path is available to a later version of UMDS.

- Verify you have administrative privileges on the Linux machine where UMDS runs.
1. In the Linux machine, open the Command Shell.
 2. Navigate to the UMDS installation directory, and locate the file `vmware-uninstall-umds.pl`.
The default UMDS installation location is `usr/local/vmware-umds`.
 3. Run the `./vmware-uninstall-umds.pl` command.
 4. To confirm that you want to uninstall UMDS from the system, enter **Yes**.
The UMDS uninstallation procedure starts.
 5. Optional: Remove PostgreSQL Database from you Linux machine.
For information about uninstalling PostgreSQL Database, go to the official PostgreSQL documentation.

UMDS is uninstalled from the Linux system.

You can upgrade your Linux OS, and install a later compatible version of UMDS.

Setting Up UMDS

You can set up UMDS to download patches for ESXi hosts. You can also set up UMDS to download ESXi6.7, ESXi7.0, and ESXi8.0 patch binaries, patch metadata, and notifications from third-party portals.

Administrator access is not a requirement for downloading patches if UMDS runs on Linux.

You can set up UMDS to download ESXi6.7, ESXi7.0, and ESXi8.0 patches from the official VMware depot and third-party portals.

The best practice is to use a job scheduler, for example cron job, to create a job that periodically triggers UMDS to download the upgrades, patches, and notifications.

Set Up the Data to Download with UMDS

By default, UMDS downloads patch binaries, patch metadata, and notifications for hosts. You can specify which patch binaries and patch metadata to download with UMDS.

1. Log in to the machine where UMDS is installed, and open a **Command Prompt** window.
2. Navigate to the directory where UMDS is installed.
The default location in 64-bit Linux is `/usr/local/vmware-umds`.
3. Specify the updates to download.

- To set up a download of all ESXi host updates, run the following command:

```
vmware-umds -S --enable-host
```

- To deactivate the download of host updates, run the following command:

```
vmware-umds -S --disable-host
```

Download the selected data.

Change the UMDS Patch Repository Location

UMDS downloads upgrades, patch binaries, patch metadata, and notifications to a folder that you can specify during the UMDS installation.

The default folder to which UMDS downloads patch binaries and patch metadata on a Linux machine is `/var/lib/vmware-umds`.

You can change the folder in which UMDS downloads data after you install UMDS.

If you have already downloaded host updates, copy all the files and folders from the old location to the new patch store location. The folder in which UMDS downloads patch binaries and patch metadata must be located on the machine on which UMDS is installed.

1. Log in as an administrator to the machine where UMDS is installed, and open a **Command Prompt** window.
2. Navigate to the directory where UMDS is installed.
The default location in 64-bit Linux is `/usr/local/vmware-umds`.

3. Change the patch repository directory by running the command:

```
vmware-umds -S --patch-store your_new_patchstore_folder
```

In this example, *your_new_patchstore_folder* is the path to the new folder in which you want to download the patch binaries and patch metadata.

You successfully changed the directory in which UMDS stores patch data.

Download data using UMDS.

Configure URL Addresses for Hosts

You can configure UMDS to connect to the websites of third-party vendors to download ESXi6.7, ESXi7.0, and ESXi8.0 host patches and notifications.

1. Log in to the machine where UMDS runs, and open a **Command Prompt** window.
2. Navigate to the directory where UMDS is installed.
The default location in 64-bit Linux is `/usr/local/vmware-umds`.
3. Configure UMDS to download data from the new URL address.
 - To add a new URL address for downloading patches and notifications for ESXi6.7, ESXi7.0, or ESXi8.0 hosts, run the following command:

```
vmware-umds -S --add-url https://host_URL/index.xml --url-type HOST
```

4. Optional: Remove a URL address, so that UMDS does not download data from it anymore.
Downloaded data is retained and can be exported. Use the following command:

```
vmware-umds -S --remove-url https://URL_to_remove/index.xml
```

You configured UMDS to download host patches and notifications from specific URL addresses.

Download the patches and notifications by using UMDS.

Using UMDS

You use UMDS in air-gap scenarios, where vSphere Lifecycle Manager has no access to the Internet. You first use UMDS to download software and notifications and then you export the data, so that it becomes available to vSphere Lifecycle Manager.

To use UMDS, the machine on which you install it must have Internet access. After you download the upgrades, patch binaries, patch metadata, you can export the data to a Web server or a portable media drive, such as CD or USB flash drive. Then, you must set up vSphere Lifecycle Manager to use the folder on the Web server or the media drive (mounted as a local disk) as a download source.

Download the Specified Data Using UMDS

After you set up UMDS, you can download upgrades, patches and notifications to the machine on which UMDS is installed.

Administrator level access is not a requirement for downloading data with UMDS that runs on Linux.

Note:

The URLs of the official VMware online depot are now unique for each customer. To update the download sources URLs, see the following KB articles: <https://knowledge.broadcom.com/external/article/390121> and <https://knowledge.broadcom.com/external/article/390123>.

1. Log in to the machine where UMDS is installed, and open a **Command Prompt** window.
2. Navigate to the directory where UMDS is installed.
The default location in 64-bit Linux is `/usr/local/vmware-umds`.
3. Download the selected updates.

```
vmware-umds -D
```

This command downloads all the upgrades, patches and notifications from the configured sources for the first time. Subsequently, it downloads all new patches and notifications released after the previous UMDS download.

4. Optional: If you have already downloaded upgrades, patches, and notifications and want to download them again, you can include the start and end times to restrict the data to download.

The command to re-download patches and notifications deletes the existing data from the patch store (if present) and re-downloads it.

To re-download the upgrades, patches and notifications that were downloaded in November 2010, for example, run the following command:

```
vmware-umds -R --start-time 2010-11-01T00:00:00 --end-time 2010-11-30T23:59:59
```

The data previously downloaded for the specified period is deleted and downloaded again.

Export the downloaded upgrades, patches, and notifications.

Export the Downloaded Data

You can export downloaded upgrades, patches, and notifications to a specific location that serves as a shared repository for vSphere Lifecycle Manager. You can configure vSphere Lifecycle Manager to use the shared repository as a patch download source. The shared repository can also be hosted on a Web server.

If you installed UMDS with an existing download directory, verify that you perform at least one download by using UMDS 8.0 before you export updates.

Administrator level access is not a requirement for exporting the downloaded data with UMDS that runs on Linux.

1. Log in to the machine where UMDS is installed and open a **Command Prompt** window.
2. Navigate to the directory where UMDS is installed.

The default location in 64-bit Linux is `/usr/local/vmware-umds`.

3. Specify the export parameters and export the data.

```
vmware-umds -E --export-store repository_path
```

In the command, you must specify the full path of the export directory.

If you are working in a deployment in which the vCenter Server is installed on a machine connected to the machine on which UMDS is installed, `repository_path` can be the path to the folder on the Web server that serves as a shared repository.

If the vCenter Server is installed on a machine in an isolated and secure environment, `repository_path` can be the path to a portable media drive. Export the downloads to the portable media drive to physically transfer the patches to the machine on which vCenter Server is installed and vSphere Lifecycle Manager runs.

The data you downloaded by using UMDS is exported to the path you specify. Make sure that all files are exported. You can periodically export from UMDS and populate the shared repository so that vSphere Lifecycle Manager can use the new patch binaries and patch metadata.

4. Optional: You can export the ESXi patches that you downloaded during a specified time window. For example, to export the patches downloaded in November 2010, run the following command:

```
vmware-umds -E --export-store repository_path --start-time 2010-11-01T00:00:00 --end-time 2010-11-30T23:59:59
```

Configure vSphere Lifecycle Manager to use a shared repository as a patch download source. For more information, see [Configure vSphere Lifecycle Manager to Use a Shared Repository as a Download Source](#).

vSphere Lifecycle Manager Remediation Settings

Whether you manage the ESXi hosts in your environment with baselines or with images, you can configure the behavior of vSphere Lifecycle Manager during remediation.

The vSphere Lifecycle Manager remediation settings for hosts and clusters that use baselines differ from the remediation settings for hosts and clusters that you manage with a single vSphere Lifecycle Manager image. For example, allowing the installation of software on PXE booted hosts and the removal of media devices before maintenance mode are settings that you can configure only for hosts and clusters that use baselines. Virtual machine migration settings, maintenance mode settings, and Quick Boot are examples of remediation settings that you can configure for both hosts and clusters that use baselines or images.

You can modify the default vSphere Lifecycle Manager settings only if you have the appropriate privileges. The permission must be assigned to the vCenter Server instance where vSphere Lifecycle Manager runs. For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation. For a list of the vSphere Lifecycle Manager privileges and their descriptions, see [Required Privileges for Using vSphere Lifecycle Manager and vSphere Configuration Profiles](#).

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, you can configure the remediation settings for each vSphere Lifecycle Manager instance. The configuration properties that you modify are applied only to the vSphere Lifecycle Manager instance that you specify, and are not propagated to the other instances in the domain.

How Do Cluster Settings Affect Remediation?

When you remediate ESXi hosts that are in a cluster, certain cluster settings might cause remediation failure. You must configure the cluster settings in such a way as to ensure successful remediation.

Distributed Resource Scheduler (DRS)

Updates might require a host to enter maintenance mode during remediation. Virtual machines cannot run when a host is in maintenance mode. To ensure availability, you can activate DRS for the cluster and you can configure it for vSphere vMotion. In this case, before the host is put in maintenance mode, vCenter Server migrates the virtual machines to another ESXi host within the cluster.

To help ensure vSphere vMotion compatibility between the hosts in the cluster, you can enable Enhanced vMotion Compatibility (EVC). EVC ensures that all hosts in the cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. EVC prevents migration failures due to incompatible CPUs. You can enable EVC only in a cluster where the host CPUs meet the compatibility requirements. For more information about EVC and the requirements that the hosts in an EVC cluster must meet, see the *vCenter Server and Host Management* documentation.

Distributed Power Management (DPM)

If a host has no running virtual machines, DPM might put the host in standby mode, which might interrupt a vSphere Lifecycle Manager operation. So, to make sure that all vSphere Lifecycle Manager operations finish successfully, you must deactivate DPM during these operations.

For successful remediation, you must configure vSphere Lifecycle Manager to deactivate DPM. After the remediation task finishes, vSphere Lifecycle Manager restores DPM. If DPM has already put a host in standby mode, vSphere Lifecycle Manager powers on the host before compliance checks, remediation, and staging. After the

respective task finishes, vSphere Lifecycle Manager turns on DPM and lets DPM put the host into standby mode, if needed. vSphere Lifecycle Manager does not remediate powered off hosts.

If a host is put in standby mode and DPM is manually deactivated for a reason, vSphere Lifecycle Manager does not remediate or power on the host.

HA Admission Control

Within a cluster, you must deactivate HA admission control temporarily to let vSphere vMotion proceed. This action prevents downtime for the machines on the hosts that you remediate. You can configure vSphere Lifecycle Manager to deactivate HA admission control during remediation. After the remediation of the entire cluster is complete, vSphere Lifecycle Manager restores the HA admission control settings. vSphere Lifecycle Manager deactivates HA admission control before remediation, but not before compliance checks. Additionally, for clusters that you manage with baselines, vSphere Lifecycle Manager deactivates HA admission control before staging.

Note:

Deactivating HA admission control before you remediate a two-node cluster that uses a single vSphere Lifecycle Manager image causes the cluster to practically lose all its high availability guarantees. The reason is that when one of the two hosts enters maintenance mode, vCenter Server cannot failover virtual machines to that host and HA failovers are never successful. For more information about HA admission control, see the *vSphere Availability* documentation.

Fault Tolerance (FT)

If FT is turned on for any of the virtual machines on a host within a cluster, you must temporarily turn off FT before performing any vSphere Lifecycle Manager operation on the cluster. If FT is turned on for any of the virtual machines on a host, vSphere Lifecycle Manager does not remediate that host. You must remediate all hosts in a cluster with the same updates, so that FT can be reactivated after remediation. A primary virtual machine and a secondary virtual machine cannot reside on hosts of different ESXi versions and patch levels.

Configuring vSphere Lifecycle Manager for Fast Upgrades

During cluster or host remediation, you can preserve the state of the virtual machines in the host memory and restore them from memory after the remediation finishes. Suspending virtual machines to memory and using the Quick Boot functionality significantly reduces the time for remediation, minimizes system boot time, and reduces the downtime of system and services.

During remediation with vSphere Lifecycle Manager, migrating virtual machines from the host that is under remediation to another host takes a considerable amount of time. After remediation, vSphere Lifecycle Manager migrates back the virtual machines to the remediated host. However, you can configure vSphere Lifecycle Manager to suspend virtual machines to memory instead of migrating them, powering them off, or suspending them to disk.

You can use the suspend to memory functionality only for patching operations, for example, when you remediate a cluster or a standalone host to apply to it a hot patch, express patch, and so on. You cannot use the suspend to memory option for upgrade operations, for example when you upgrade your ESXi hosts from version 7.0 to version 8.0.

Suspend Virtual Machines to Memory

Suspend to memory is an option that you can use only for clusters or hosts that you manage with vSphere Lifecycle Manager images. The functionality works together with the Quick Boot setting to optimize the remediation process and minimize virtual machine downtime.

You enable vSphere Lifecycle Manager to suspend virtual machines to memory when you configure the vSphere Lifecycle Manager host remediation settings. During remediation pre-check and remediation, vSphere Lifecycle Manager verifies that the suspend to memory option is indeed applicable to the host or cluster under remediation. If for some reason suspend to memory is inapplicable, vSphere Lifecycle Manager reports an error and prevents remediation from proceeding.

During a suspend to memory operation, virtual machines remain in a suspended state for some time. So, suspending virtual machines to memory might impact the workloads running on those virtual machines. The impact is similar to the impact that the suspend to disk operation might have on virtual machines and workloads.

CAUTION: As a best practice, always take snapshots of the virtual machines with critical workloads before you start remediation when the suspend to memory option is activated.

vSphere Lifecycle Manager might not suspend to memory all virtual machines on the host even if you activated the feature for the entire cluster or the standalone host. In some cases, vSphere Lifecycle Manager is still able to proceed with the remediation of the host, even if some virtual machines cannot be suspended to memory.

- vSphere ESX Agent Manager (EAM) virtual machines
vSphere Lifecycle Manager powers off the EAM virtual machines after all other virtual machines are suspended. Similarly, vSphere Lifecycle Manager powers on the EAM virtual machines before any other virtual machines are resumed from memory. None of the suspended virtual machines is resumed until the EAM virtual machines are powered on.
- vSphere Cluster Services virtual machines
vSphere Lifecycle Manager first migrates to another host the vSphere Cluster Services virtual machines, and then suspends to memory the rest of the virtual machines on the host.

Similarly, vSphere Lifecycle Manager does not suspend to memory the management virtual appliances for some VMware products and solutions. However, if a virtual machine for any of the following products or solutions runs on a host, the suspend to memory pre-check fails and vSphere Lifecycle Manager does not proceed with the remediation of the respective host:

- vCenter Server
- vSAN witness virtual machine
- vSphere Supervisor
- NSX
- VMware HCX
- vSphere Replication
- Site Recovery Manager
- VMware Aria products

Note: Third-party virtual machines do get suspended during remediation, if the Suspend to memory option is activated.

Quick Boot

Quick Boot is a setting that you can use with clusters or standalone hosts that you manage with vSphere Lifecycle Manager images and vSphere Lifecycle Manager baselines. Using Quick Boot optimizes the host patching and upgrade operations. Quick Boot lets vSphere Lifecycle Manager reduce the remediation time for hosts that undergo patch and upgrade operations. Patch and upgrade operations do not affect the hardware of a host. If the Quick Boot feature is activated, vSphere Lifecycle Manager skips the hardware reboot (the BIOS or UEFI firmware reboot). As a result, the time an ESXi host spends in maintenance mode is reduced and the risk of failures during remediation is minimized.

To configure vSphere Lifecycle Manager to suspend virtual machines to the host memory, you must activate Quick Boot. However, you can activate Quick Boot even if you decide not to use the Suspend to memory option.

Requirements for Using Suspend to Memory

Several factors might hinder the applicability of the suspended to memory option. If for some reason suspend to memory is inapplicable, vSphere Lifecycle Manager reports an error and prevents remediation from proceeding. Suspend to memory works under the following conditions:

- The host supports the suspend to memory functionality.
- Quick Boot is activated for the cluster or the standalone host and the host under remediation supports Quick Boot.
- The remediation does not involve host upgrades or firmware upgrade.
- The host and the virtual machines meet certain requirements.

Host Requirements	Virtual Machine Requirements
• The host has enough free memory. • The host has sufficient free low memory. • The host has enough free memory per NUMA node to start after a reboot. • The host has enough reservation available • The host does not use swapped or compressed pages of virtual machines.	• The virtual machines do not have any passthrough devices. • The virtual machines do not have latency sensitivity set to high. • The virtual machines are not fault tolerant. • The virtual machines are not encrypted. • The virtual machines do not use persistent memory. • The virtual machines do not have virtual SGX or SEV devices. • The virtual machines do not have the suspend feature deactivated. • The virtual machines are not frozen source virtual machines during an Instant Clone operation.

Suspend to Memory and vSphere High Availability (HA)

When you configure vSphere Lifecycle Manager to suspend virtual machines to memory during remediation, vSphere HA provides protection for the suspended virtual machines in cases of failure at the virtual machine or host level. By modifying the vSphere HA advanced options, you can set a timeout value for suspended to memory virtual machines. If a suspended to memory virtual machine is not responsive for the specified time, vSphere HA powers on the virtual machine on the original host or on another host.

- If you deactivate or reconfigure vSphere HA for the cluster during remediation, vSphere HA can no longer protect the suspended virtual machines. Before you change the vSphere HA configuration, make sure that no hosts in the cluster are in maintenance mode and the suspended virtual machines are powered on.
- If you modify the `das.failoverDelayForSuspendToMemoryVmsSecs` advanced option for vSphere HA after you configure vSphere Lifecycle Manager to use the suspend to memory option, the newly specified timeout value might not apply to the virtual machines. If you need to modify the default value of the `das.failoverDelayForSuspendToMemoryVmsSecs` option, ensure that you modify it before you start remediation to ensure that the new value is in effect.
- If the suspend to memory operation fails, vSphere HA determines the most appropriate failover host after the specified timeout value expires. The failover host might be the original host or another one.
- You must synchronize the server time for all ESXi hosts in the cluster. If the hosts are not synchronized, vSphere HA might not respect the specified timeout period and initiate failover earlier or later.

For more information about using and configuring vSphere HA, see the *vSphere Availability* documentation.

Configuring vSphere Lifecycle Manager for Live Patches

When you enable the Live Patch setting, you can apply security patches and urgent bug fixes to all hosts in a cluster that you manage with a vSphere Lifecycle Manager image with minimal disruption to your workloads.

Starting with vSphere 8.0 Update 3, you can use the Live Patch setting to deploy patches with resolved security and functional issues to all ESXi hosts in a cluster managed with a single image. During remediation, vSphere Lifecycle Manager does not place the hosts in the cluster into maintenance mode, hosts are not rebooted and there is no need to migrate the virtual machines running on the hosts in the cluster.

To configure vSphere Lifecycle Manager to perform live patch operations on the hosts in a cluster managed with an image, you must activate Live Patch. See [Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image](#).

Requirements for Using Live Patch

You use vSphere Lifecycle Manager to install live patches on hosts that are part of a cluster managed with a single image, only in case the following conditions are met:

- vCenter Server must be of version 8.0 Update 3 and later.
- All ESXi hosts in the cluster managed with a vSphere Lifecycle Manager image must be of version 8.0 Update 3 and later.
- You use the Live Patch functionality only to upgrade hosts in a cluster that you manage with a vSphere Lifecycle Manager image.
- All hosts in the cluster managed with a single image must support the Live Patch functionality. If you activate Live Patch for a vCenter Server instance or a cluster and you try to install a live patch on the cluster with hosts that require maintenance mode, remediation is blocked. You can view details about the compliance of the cluster against the selected live patch image for further information.
- vSphere Lifecycle Manager can install a live patch only to a host that is eligible to be updated with a specific live patch release of the VMware base image. vSphere Lifecycle Manager image depot provides information about all the possible versions of the ESXi host image which you can update with a specific Live Patch release. For more information about how to browse the depot, see [Browsing the vSphere Lifecycle Manager Depot](#).
- vSphere DRS must be enabled on the cluster managed with a single image before you try to install a live patch on all hosts.
- You must deactivate the parallel remediation setting before you try to install a live patch on the hosts in the cluster. If you have both settings enabled, the remediation pre-check reports this as an issue for a successful remediation and provides you with possible solutions. Live patches can only be installed on the hosts in a cluster managed with a single image in sequence.

How to Configure the vSphere Lifecycle Manager Remediation Settings

In the vSphere Client, you configure separately the vSphere Lifecycle Manager remediation settings for the hosts and clusters that you manage with a single image and with baselines.

The global remediation settings are valid for all hosts and clusters that you manage with images or baselines respectively. For clusters or standalone hosts that you manage with a single image, you can override the global remediation settings and use specific remediation settings for a particular cluster or host.

Hosts that are in a vSAN cluster can enter maintenance mode only one at a time. This behavior is a peculiarity of the vSAN cluster. For more information about the vSphere Lifecycle Manager behavior during the remediation of hosts in a vSAN cluster, see [Remediation Specifics of vSAN Clusters](#).

Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image

You can configure how ESXi hosts and VMs behave before and during the remediation of a cluster or a standalone host that you manage with a single image.

- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).
- To use Quick Boot, verify that the ESXi host is compatible with the feature. For more information, see [Quick Boot](#).

For information about automatically triggered hardware compatibility checks, which is a functionality that is also applicable only to vSAN clusters, see [When Does the Hardware Compatibility Check Task Run](#).

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.
 The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Settings** tab, select **Cluster lifecycle > Images**.
3. Edit the cluster or standalone host remediation settings.
 - In the **Images** pane, click the **Edit** button.
 The **Edit Cluster Remediation Settings** dialog box opens.
 - In the **Standalone Host Images** pane, click the **Edit** button.
 The **Edit Standalone Host Remediation Settings** dialog box opens.
4. Configure the remediation settings for clusters and standalone hosts that use images, and click **Save**.

Setting	Description
Quick Boot	Quick Boot reduces the host reboot time during remediation. By default, Quick Boot is turned on. You can configure this setting for a cluster or a standalone host.
Live Patch	Live Patch allows you to perform low-disruption upgrades of the ESXi hosts in a cluster managed with an image. When you turn on Live Patch, you can install the latest security patches and bug fixes without placing the hosts into maintenance mode or evacuating virtual machines from the hosts. See Configuring vSphere Lifecycle Manager for Live Patches . By default, Live Patch is turned off. You can configure this setting only for a cluster that you manage with a single image.
VM power state	The VM power state option lets you control the behavior of the virtual machines that run on the ESXi host. You can select from the following options. • Do not change power state • Suspend to disk • Suspend to memory To select the Suspend to memory option, you must turn on Quick Boot. Otherwise, the Suspend to memory option is dimmed.

Setting	Description
	Together with Quick Boot, the Suspend to memory option provides faster host upgrades. vSphere Lifecycle Manager suspends to the host memory and not to the disk the powered on virtual machines on the host. After the Quick Boot, the suspended virtual machines are resumed from memory. • Power off The default selection is Do not change power state. You can configure these options for a cluster or a standalone host.
VM migration	You can configure vSphere Lifecycle Manager to migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster. By default, the Migrate powered off and suspended VMs to other hosts in the cluster check box is selected. If you have virtual machines with passthrough devices on a host in a DRS cluster, make sure that you automate the migration of such virtual machines. For more information, see the VMware knowledge base article at https://kb.vmware.com/s/article/88271. You can configure this setting only for a cluster that you manage with a single image.
Retry policy	You can configure how vSphere Lifecycle Manager behaves if a host fails to enter maintenance mode or if staging or remediation for that host fails. You can configure vSphere Lifecycle Manager to wait for a specified retry delay period and to retry to put the host into maintenance mode as many times as you indicate in the Number of retries text box. You can configure this setting for a cluster or a standalone host.
HA admission control	Admission control is a policy that vSphere HA uses to ensure failover capacity within a cluster. If vSphere HA admission control is enabled during remediation, vMotion might be unable to migrate the virtual machines within the cluster. Deactivating admission control allows a virtual machine to be powered on even if it causes insufficient failover capacity. When this happens, no warnings are presented, and the cluster does not turn red. If a cluster has insufficient failover capacity, vSphere HA can still perform failovers, and uses the VM Restart Priority setting to determine which virtual machines to power on first. Note: Deactivating HA admission control before you remediate a two-node cluster causes the cluster to practically lose all its high availability guarantees. • If you select the Disable HA admission control on the cluster option, vSphere Lifecycle Manager remediates the hosts in the cluster and re-enables HA admission control after remediation is complete. • If you deselect the Disable HA admission control on the cluster option, vSphere Lifecycle Manager skips

Setting	Description
	remediating the clusters on which HA admission control is enabled. By default, the Disable HA admission control on the cluster check box is deselected. You can configure this setting only for a cluster that you manage with a single image.
DPM	VMware Distributed Power Management (DPM) monitors the resources consumed by the running virtual machines in the cluster. If sufficient excess capacity exists, VMware DPM recommends moving virtual machines to other hosts in the cluster and placing the original host into standby mode to conserve power. If the capacity is insufficient, VMware DPM might recommend returning standby hosts to a powered-on state. - If you select the Disable DPM on the cluster option, vSphere Lifecycle Manager remediates the hosts in the cluster and re-enables DPM after remediation is complete. - If you deselect the Disable DPM on the cluster option, vSphere Lifecycle Manager skips remediating the clusters on which DPM is enabled. By default, the Disable DPM on the cluster check box is selected. You can configure this setting only for a cluster that you manage with a single image.
Hardware compatibility issues	vSphere Lifecycle Manager performs a hardware compatibility check as part of the remediation pre-check and the remediation tasks for vSAN clusters. You can configure vSphere Lifecycle Manager to prevent remediation when hardware compatibility issues exist for the cluster. - If you select the Prevent remediation if hardware compatibility issues are found option, vSphere Lifecycle Manager reports hardware compatibility issues as an error, which prevents remediation. - If you deselect the Prevent remediation if hardware compatibility issues are found option, vSphere Lifecycle Manager reports hardware compatibility issues as a warning, which does not prevent remediation. If the cluster is not vSAN-enabled, vSphere Lifecycle Manager does not perform a hardware compatibility check as part of the remediation pre-check or the remediation tasks. You can configure this setting only for a cluster that you manage with a single image.
Parallel remediation	Enabling parallel remediation allows you to remediate multiple ESXi hosts simultaneously. By selecting the Parallel remediation check box, you enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. You can also specify the maximum number of concurrent remediations manually. Note: If the hosts have NSX virtual distributed switches that are ready to be migrated to vSphere Distributed Switches,

Setting	Description
	you must manually set the maximum number of parallel remediations to no more than 4. In cases when host switch migration is needed, if more than 4 hosts are remediated in parallel, the remediation might fail, because the host switch migration takes more time than the time vSphere Lifecycle Manager needs to complete the parallel remediation. You can configure this setting only for a cluster that you manage with a single image.

These settings become the default remediation settings for clusters or standalone hosts that you manage with vSphere Lifecycle Manager images. For any cluster or host, you can change the default settings when you configure individual remediation tasks.

Override the Global vSphere Lifecycle Manager Remediation Settings for a Cluster or Host That You Manage with a Single Image

You can customize the remediation settings for a particular cluster or a standalone host while the global remediation settings remain intact and applicable to all other clusters or hosts that you manage with vSphere Lifecycle Manager images and that are in the same vCenter Server instance.

- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).
- To enable Quick Boot, verify that the ESXi host is compatible with the feature. For more information, see [Quick Boot](#).

When you override the global remediation settings for a single cluster or host, the overrides are used during the remediation of that specific cluster or host. For all other clusters or hosts, the global remediation settings apply.

1. In the vSphere Client, navigate to a cluster or host that you manage with a single image.
2. On the **Updates** tab, select **Hosts > Image**.
3. In the **Image Compliance** card, click the horizontal ellipsis icon and select **Edit remediation settings**.
4. In the **Edit Remediation Settings** dialog box, configure the remediation settings of the target cluster or host and click **Save**.

Setting	Description
Quick Boot	Quick Boot reduces the host reboot time during remediation. By default, Quick Boot is enabled. You can configure this setting for a cluster or a standalone host.
Live Patch	Live Patch allows you to perform low-disruption upgrades of the ESXi hosts in a cluster managed with an image. When you turn on Live Patch, you can install the latest security patches and bug fixes without placing the hosts into maintenance mode or evacuating virtual machines from the hosts. See Configuring vSphere Lifecycle Manager for Live Patches. By default, Live Patch is turned off. You can configure this setting only for a cluster that you manage with a single image.
VM power state	The VM power state option lets you control the behavior of the virtual machines that run on the ESXi host. You can select from the following options. • Do not change power state

Setting	Description
	• Suspend to disk • Suspend to memory To select the Suspend to memory option, you must enable Quick Boot. Otherwise, the Suspend to memory option is dimmed. Together with Quick Boot, the Suspend to memory option provides faster host upgrades. vSphere Lifecycle Manager suspends to the host memory and not to the disk the powered on virtual machines on the host. After the Quick Boot, the suspended virtual machines are resumed from memory. • Power off The default selection is Do not change power state. You can configure these options for a cluster or a standalone host.
VM migration	You can configure vSphere Lifecycle Manager to migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster. By default, the Migrate powered off and suspended VMs to other hosts in the cluster check box is selected. If you have virtual machines with passthrough devices on a host in a DRS cluster, make sure that you automate the migration of such virtual machines. For more information, see the VMware knowledge base article at https://kb.vmware.com/s/article/88271. You can configure this setting only for a cluster that you manage with a single image.
Retry policy	You can configure how vSphere Lifecycle Manager behaves if staging or remediation fails, or if a host fails to enter maintenance mode. You can configure vSphere Lifecycle Manager to wait for a specified retry delay period and to retry to put the host into maintenance mode as many times as you indicate in the Number of retries text box. You can configure this setting for a cluster or a standalone host.
HA admission control	Admission control is a policy that vSphere HA uses to ensure failover capacity within a cluster. If vSphere HA admission control is enabled during remediation, vMotion might be unable to migrate the virtual machines within the cluster. Deactivating admission control allows a virtual machine to be powered on even if it causes insufficient failover capacity. When this happens, no warnings are presented, and the cluster does not turn red. If a cluster has insufficient failover capacity, vSphere HA can still perform failovers, and uses the VM Restart Priority setting to determine which virtual machines to power on first. Note: Deactivating HA admission control before you remediate a two-node cluster causes the cluster to practically lose all its high availability guarantees.

Setting	Description
	- If you select the Disable HA admission control on the cluster option, vSphere Lifecycle Manager remediates the hosts in the cluster and re-enables HA admission control after remediation is complete. - If you deselect the Disable HA admission control on the cluster option, vSphere Lifecycle Manager skips remediating the clusters on which HA admission control is enabled. By default, the Disable HA admission control on the cluster check box is deselected. You can configure this setting only for a cluster that you manage with a single image.
DPM	VMware Distributed Power Management (DPM) monitors the resources consumed by the running virtual machines in the cluster. If sufficient excess capacity exists, VMware DPM recommends moving virtual machines to other hosts in the cluster and placing the original host into standby mode to conserve power. If the capacity is insufficient, VMware DPM might recommend returning standby hosts to a powered-on state. - If you select the Disable DPM on the cluster option, vSphere Lifecycle Manager remediates the hosts in the cluster and re-enables DPM after remediation is complete. - If you deselect the Disable DPM on the cluster option, vSphere Lifecycle Manager skips remediating the clusters on which DPM is enabled. By default, the Disable DPM on the cluster check box is selected. You can configure this setting only for a cluster that you manage with a single image.
Hardware compatibility issues	vSphere Lifecycle Manager performs a hardware compatibility check as part of the remediation pre-check and the remediation tasks for vSAN clusters. You can configure vSphere Lifecycle Manager to prevent remediation when hardware compatibility issues exist for the cluster. - If you select the Prevent remediation if hardware compatibility issues are found option, vSphere Lifecycle Manager reports hardware compatibility issues as an error, which prevents remediation. - If you deselect the Prevent remediation if hardware compatibility issues are found option, vSphere Lifecycle Manager reports hardware compatibility issues as a warning, which does not prevent remediation. If the cluster is not vSAN-enabled, vSphere Lifecycle Manager does not perform a hardware compatibility check as part of the remediation pre-check or the remediation tasks. You can configure this setting only for a cluster that you manage with a single image.
Parallel remediation	Enabling parallel remediation allows you to remediate multiple ESXi hosts simultaneously. By selecting the Parallel

Setting	Description
	remediation check box, you enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. You can also specify the maximum number of concurrent remediations manually. Note: If the hosts have NSX virtual distributed switches that are ready to be migrated to vSphere Distributed Switches, you must manually set the maximum number of parallel remediations to no more than 4. In cases when host switch migration is needed, if more than 4 hosts are remediated in parallel, the remediation might fail, because the host switch migration takes more time than the time vSphere Lifecycle Manager needs to complete the parallel remediation. You can configure this setting only for a cluster that you manage with a single image.

These settings become the remediation settings for the selected cluster. vSphere Lifecycle Manager uses those settings for that cluster for all future remediation tasks. The global remediation settings remain unchanged and are applied to all other clusters.

In the **Image Compliance** card, vSphere Lifecycle Manager displays a message that the global remediation settings are overridden. Also, an option to reset the values appears in the card.

Configure the vSphere Lifecycle Manager Remediation Settings for Hosts and Clusters That You Manage with Baselines

You can configure how vSphere Lifecycle Manager behaves before and during remediation against a baseline or a baseline group. The remediation settings help ensure that vSphere Lifecycle Manager puts ESXi hosts in maintenance mode before remediation.

- Required privileges: **VMware vSphere Lifecycle Manager > Configure**
- To enable Quick Boot, verify that the ESXi host is compatible with the feature. For more information, see [Quick Boot](#).

When you use vSphere Lifecycle Manager baselines, you can configure vSphere Lifecycle Manager to let other software initiate the remediation of PXE booted ESXi hosts. The remediation installs software modules on the hosts, but typically those host updates are lost after a reboot. To retain updates on stateless hosts after a reboot, use a PXE boot image that contains the updates. You can update the PXE boot image before applying the updates with vSphere Lifecycle Manager, so that the updates are not lost because of a reboot. vSphere Lifecycle Manager itself does not reboot the hosts, because it does not install updates requiring a reboot on PXE booted ESXi hosts.

- Navigate to the vSphere Lifecycle Manager home view.
 - In the vSphere Client, select **Menu > Lifecycle Manager**.
 - Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
- On the **Settings** tab, select **Cluster Lifecycle > Baselines**.
- Click the **Edit** button.

The **Edit Settings for Host Remediation** dialog box opens.

4. Configure the baselines remediation settings and click **Save**.

Option	Description
VM power State	You can configure vSphere Lifecycle Manager to power off or suspend all running virtual machines before host remediation. Alternatively, you can configure vSphere Lifecycle Manager not to change the power state of the virtual machines.
Maintenance mode failures	You can configure how vSphere Lifecycle Manager behaves if a host fails to enter maintenance mode before remediation. If you enable vSphere Lifecycle Manager to retry to put the host in maintenance mode, you must specify the number of retries and the retry delay. vSphere Lifecycle Manager waits for as much time as you configure for the Retry Delay option, and retries to put the host in maintenance mode as many times as you indicate in the Number of retries text box.
PXE booted hosts	You can allow the installation of software for solutions on the PXE booted ESXi hosts in the vSphere inventory that you manage with vSphere Lifecycle Manager baselines.
VM migration	If you enable virtual machine migration by selecting the respective option, vSphere Lifecycle Manager migrates the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster.
Removable media devices	vSphere Lifecycle Manager does not remediate hosts on which virtual machines have connected CD/DVD or floppy drives. All removable media drives that are connected to the virtual machines on a host might prevent the host from entering maintenance mode and interrupt remediation. So, you can disconnect all removable media devices to ensure that the respective host enters maintenance mode successfully. After remediation, vSphere Lifecycle Manager reconnects the removable media devices if they are still available.
Quick Boot	Quick Boot reduces the host reboot time during remediation. By default, Quick Boot is deactivated.
Parallel remediation	Enabling parallel remediation allows you to remediate multiple ESXi hosts simultaneously. By selecting the Parallel remediation option, you enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. Alternatively, you can specify the maximum number of concurrent remediations manually. Note: If the hosts have NSX virtual distributed switches that are ready to be migrated to vSphere Distributed Switches, you must manually set the maximum number of parallel remediations to no more than 4. In cases when host switch migration is needed, if more than 4 hosts are remediated in parallel, the remediation might fail, because the host switch

Option	Description
	migration takes more time than the time vSphere Lifecycle Manager needs to complete the parallel remediation.

These settings become the default failure response settings with vSphere Lifecycle Manager baselines. You can specify different settings when you configure individual remediation tasks.

Configure vSphere Lifecycle Manager Remediation Settings for vSAN Clusters that You Manage with vSphere Lifecycle Manager Baselines

You can configure what type of baseline to include in the recommendation baseline group that vSphere Lifecycle Manager generates for a vSAN cluster.

- Verify that you manage the vSAN cluster with vSphere Lifecycle Manager baselines and not a single vSphere Lifecycle Manager image.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

A recommendation baseline group can contain upgrades or only host patches and updates. By default, vSphere Lifecycle Manager is set to generate recommendation baselines that contain upgrades, not only patches and updates. However, you can change the default configuration at any time. For any vSAN cluster that you manage with baselines, you can also configure vSphere Lifecycle Manager to generate no recommendation baseline group at all. In such cases, you can still manually create baselines and perform host upgrades.

1. In the vSphere Client, navigate to a vSAN cluster that you manage with baselines.
2. On the **Updates** tab, select **Hosts** > **Cluster Settings**.
3. In the **Remediation Settings for this Cluster** pane, click the **Edit** button.
The **Edit Cluster Remediation Settings** dialog box appears.
4. Select what type of baseline to include in the recommendation baseline group that vSphere Lifecycle Manager generates for the selected vSAN cluster.
 - To include upgrade baselines into the recommendation baseline group for that cluster, select the **Include upgrades to new ESXi versions** radio button.
The **Include upgrades to new ESXi versions** options is the default selection for any newly created vSAN cluster.
 - To include only patches and updates in the recommendation baseline group for that cluster, select the **Include patches and updates for current ESXi version** radio button.
 - To stop the generation of the recommendation baseline group for that cluster, select the **No recommendation** baseline group.

5. Click the **Done** button to save your selection and exit the dialog box.

The option that you select becomes the default configuration for the vSAN cluster.

Creating and Managing vSphere Lifecycle Manager Clusters

A vSphere Lifecycle Manager cluster is a cluster of ESXi hosts that you manage either with baselines or with a single image. You decide whether to manage a cluster with baselines or with a single image during the creation of the cluster.

Creating a vSphere Lifecycle Manager Cluster That Uses a Single Image

When you create a cluster, the option to use a single image for that cluster is selected by default. You must keep it selected unless you want to create a cluster that uses baselines. During cluster creation, you must set up an image yourself or use the image of a host within or outside the current vCenter Server instance. In both cases, you can also enable vSphere Configuration Profiles for the cluster that you create.

Starting with vSphere 8.0, the default selections in the **New Cluster** wizard are based on the presumption that you are creating a cluster that uses a single image.

Creating a vSphere Lifecycle Manager Cluster That Uses Baselines

To create a cluster that uses baselines, during cluster creation, you deselect the option to manage the cluster with a single image. You can switch from using baselines to using images at a later time. For more information about switching from using baselines to using images, see [Convert a Cluster or a Host That Uses Baselines Into a Cluster or a Host That Uses vSphere Lifecycle Manager Images](#).

Adding Hosts to a vSphere Lifecycle Manager Cluster

You can add hosts of any version to a cluster that you manage with baselines.

You can add hosts of ESXi version 7.0 or later to a cluster that you manage with a single image. You can add a host to a cluster and at the same time use the image on that host as an image for the entire cluster.

Removing Hosts from a vSphere Lifecycle Manager Cluster

Removing a host from a cluster is a straightforward procedure. If you remove a host from a cluster that uses a single image, the host retains the software and firmware installed during the last remediation against the image for the cluster.

Note: When you remove a host from a vSAN cluster that you manage with a single image, vSphere Lifecycle Manager invalidates the results from the last hardware compatibility check for the cluster. To obtain valid hardware compatibility information about the cluster, you must re-run a hardware compatibility check. For instructions on how to check the hardware compatibility for a cluster, see [Check the Hardware Compatibility of a Cluster](#).

All cluster-related operations are described in full detail in the *vCenter Server and Host Management* documentation.

For information about using Auto Deploy to deploy and provision ESXi hosts, see the *VMware ESXi Installation and Setup* documentation.

Create a Cluster That Uses a Single Image by Composing a New Image

You can create a vSphere Lifecycle Manager cluster that uses a single image by setting up the desired image during the creation of the cluster. When you choose to set up the image manually instead of importing it from a host, you must have the necessary software available in the vSphere Lifecycle Manager depot.

- Review the requirements for using a single image in [What Are the Requirements for Using vSphere Lifecycle Manager](#).
- Verify that the hosts are of ESXi version 7.0 or later.
- Verify that a data center exists in the vCenter Server inventory.
- Verify that you have an ESXi image available in the vSphere Lifecycle Manager depot.
- If you want to enable vSphere Configuration Profiles on the cluster, verify that the following requirements are met:
 - You have the required license.
 - The hosts in the cluster are of ESXi version 8.0 or later.
 - vCenter Server is version 8.0 or later.

When you set up an image manually, you must specify the ESXi version for the image and, optionally, a vendor add-on. After the cluster creation is complete, you can further customize the image. For example, you can add components to the image. You can also configure a hardware support manager and add a firmware and drivers add-on to the image. For information about editing an image, see [Edit a vSphere Lifecycle Manager Image](#).

If you decide to enable vSphere Configuration Profiles on the cluster, you must set up the host configuration manually after you complete the cluster creation.

1. In the vSphere Client, navigate to the **Hosts and Clusters** inventory.
2. Right-click a data center and select **New Cluster**.
The **New Cluster** wizard opens.
3. On the **Basics** page, enter a name for the cluster and enable vSphere DRS, vSphere HA, or vSAN.
4. Leave the **Manage all hosts in the cluster with a single image** check box and **Compose a new image** radio button selected.
Those options are selected by default.
5. Optional: To enable vSphere Configuration Profiles on the cluster, select the **Manage configuration at a cluster level**.
6. Click **Next**.
7. On the **Image** page, set up the desired image.
 - a) Select an ESXi version.
 - b) Optional: Select a vendor add-on and the version of the add-on.
8. Click **Next**.
9. On the **Review** page, review your selections and the image setup.
10. Click **Finish** to finish the creation of the cluster.

A cluster that uses a single image appears in the vCenter Server inventory. You can view and customize the cluster image on the **Updates** tab for the cluster.

If the cluster has vSphere Configuration Profiles enabled, you can configure the settings of the hosts in the cluster collectively by going to the **Configure** tab for the cluster and clicking **Desired State > Configuration**. See [Edit the Host Settings of a Draft Configuration in vSphere Client](#).

Add hosts to the cluster.

Create a Cluster That Uses a Single Image by Importing an Image from a Host

During cluster creation, instead of composing a new image, you can import the desired software specification from a reference host. If you choose to import an image, vSphere Lifecycle Manager extracts the software specification from the reference host and uses it for the newly created cluster. Importing an image saves you the time and effort of ensuring that you have all necessary components and images available in the vSphere Lifecycle Manager depot. Also, because you use a ready-made image, you do not need to spend time validating the new image.

- Verify that the vCenter Server version is 7.0 Update 2
- Verify that a data center exists in the vCenter Server inventory.
- Verify that the reference host is version ESXi 7.0 Update 2 or later.
- Verify that the ESX-OSDATA volume is backed by a dedicated NVMe, SSD, or HDD device.
- Obtain the user name and password of the root user account for the reference host if it is not in your vCenter Server instance.
- Review the requirements for using a single image in [What Are the Requirements for Using vSphere Lifecycle Manager](#).

- If you want to enable vSphere Configuration Profiles on the cluster, verify that the following requirements are met:
 - You have the required license.
 - The hosts in the cluster are of ESXi version 8.0 or later.
 - vCenter Server is version 8.0 or later.

During image import, along with extracting the software specification from the reference host, vSphere Lifecycle Manager also extracts the software depot associated with the image, and imports the software components to the vSphere Lifecycle Manager depot in the vCenter Server instance where you create the cluster. As a result, in air-gap scenarios, you only need one reference host to obtain the necessary ESXi image and components in the local depot and to create a software specification for your clusters.

You can import an image from an ESXi host that is in the same or a different vCenter Server instance. You can also import an image from an ESXi host that is not managed by vCenter Server. The reference host can also be in a cluster that you manage with baselines. Along with importing the image, you can also choose to move the reference host to the cluster. As a result, the newly created cluster uses the same image as the image on the reference host, which is now part of that cluster. If the reference host is in another vCenter Server instance, you can import the image from that host, but you cannot move it to the cluster.

Note: When you import an image from a host, vSphere Lifecycle Manager retrieves the ESXi version, vendor add-on, and user-added components from the host. vSphere Lifecycle Manager does not extract the components from solutions and firmware updates installed on the reference host. Therefore, the image for the new cluster does not contain solution components or a firmware and drivers add-on. To obtain firmware updates in the depot and add a firmware and drivers add-on to your cluster image, you must configure a hardware support manager for the cluster after the cluster is created. For more information about firmware updates, see [Firmware Updates with vSphere Lifecycle Manager](#).

In vSphere 8.0, you can use the image of a reference host that is in a DPU-based cluster and import it into a non-DPU cluster. Similarly, you can select a non-DPU host as the reference host for a DPU-based cluster.

1. In the vSphere Client, navigate to the **Hosts and Clusters** inventory.
2. Right-click a data center and select **New Cluster**.
The **New Cluster** wizard opens.
3. On the **Basics** page, enter a name for the cluster and enable vSphere DRS, vSphere HA, or vSAN.
4. Leave the **Manage all hosts in the cluster with a single image** check box selected.
This option is selected by default.
5. Choose the import method of creating an image for the cluster.
 - To import an image from a host that is in the same vCenter Server inventory, select the **Import image from an existing host in the vCenter inventory** radio button.
 - To import an image from a host that is in a different vCenter Server instance or a standalone host that is not added to your vCenter Server, select the **Import image from a new host** radio button.
6. Optional: To enable vSphere Configuration Profiles on the cluster, select the **Manage configuration at a cluster level** check box.
7. Click **Next**.
8. Select the host image you want to import and use for the cluster.

Import Option	Steps
Import an image from an existing host in the vCenter Server inventory	1. On the Select Image page, select the reference host which you want to extract the image from and click Next . Note: After you select a reference host, detailed information about the image on the selected host appears at the

Import Option	Steps
	bottom of the page. You can view the ESXi version and all additional components.
Import an image from a new host	1. On the Image page, enter the host details and click the Find Host button. 2. If a Security Alert dialog box appears, click Yes to confirm that you want to connect to the host. 3. To move a standalone host to the cluster, select the Also move selected host to cluster check box and click Next.

9. On the **Review** page, review your selections and verify that the selected reference host and its image are what you need.

10. Click **Finish**.

A cluster that uses a single image appears in the vCenter Server inventory. Depending on your selections, the reference host might be in the newly created cluster. The image for that cluster is identical to the image on the selected reference host. You can view and customize the cluster image on the **Updates** tab for the cluster.

If the cluster has vSphere Configuration Profiles enabled, you can configure the settings of the hosts in the cluster collectively by going to the **Configure** tab for the cluster and clicking **Desired State > Configuration**. See [Edit the Host Settings of a Draft Configuration in vSphere Client](#).

Add other hosts to the cluster.

Add Hosts to a Cluster that Uses a Single Image

When you add hosts to a cluster, you can appoint one of those hosts as a reference host. vSphere Lifecycle Manager extracts and uses the image on the reference host as the new image for the cluster. The option to import a host's image to the cluster facilitates and simplifies the cluster upgrade operation by removing the need for you to manually import components to the vSphere Lifecycle Manager depot, set up and validate a new image, check the compliance of the hosts against the image, and then remediate the cluster against the image to apply the new software specification to all hosts.

- Verify that the vCenter Server version is 7.0 or later.
- Verify that the hosts to add are of version ESXi 7.0 or later.
- Verify that the hosts that you add to the cluster are of the same main and patch version as the rest of the hosts.
- Obtain the user name and password of the root user account for the hosts that are not in your vCenter Server instance.
- Review the requirements for using a single image in [What Are the Requirements for Using vSphere Lifecycle Manager](#).
- To add a host and import its image to the cluster, verify that the following requirements are met.
 - The vCenter Server instance is of version 7.0 Update 2 or later.
 - The ESXi version on the reference host is 7.0 Update 2 or later.
- To add a host to a cluster that has vSphere Configuration Profiles enabled, verify that the following requirements are met.
 - The vCenter Server instance is of version 8.0 or later.
 - The ESXi version of the host 7.0 Update 3 or later.
 - The host doesn't have NSX or vSphere Distributed Switch configured.
 - No host profile is attached to the host.

Along with extracting the software specification from the appointed reference host, vSphere Lifecycle Manager also extracts the software depot associated with the image, and imports the components to the vSphere Lifecycle Manager depot in the vCenter Server instance where the target cluster is.

Note: When you import an image from a host, vSphere Lifecycle Manager retrieves the ESXi version, vendor add-on, and user-added components from the host. vSphere Lifecycle Manager does not extract the components from solutions and firmware updates installed on the reference host. Therefore, the new image for the cluster does not contain solution

components or a firmware and drivers add-on. To obtain firmware updates in the depot and add a firmware and drivers add-on to your cluster image, you must configure a hardware support manager. For more information about firmware updates, see [Firmware Updates with vSphere Lifecycle Manager](#).

Adding a host to a cluster and importing its image to the target cluster changes the compliance state of the other hosts in the cluster. After adding the host and setting its image as the new cluster image, you can run a compliance check. The newly added host is compliant against the new cluster image. The rest of the hosts become non-compliant. To apply the new cluster image to all the hosts in the cluster and make them compliant, you must remediate the cluster.

Note: You cannot downgrade the software that is actually installed on the hosts in the cluster. If the image on the reference host contains software components of lower version, you can still import and use that image for the cluster. However, the hosts in the cluster become incompatible with the new image, and you cannot proceed and remediate the cluster against that image.

1. In the vSphere Client, navigate to the **Hosts and Clusters** inventory.
2. Right-click cluster and select **Add Hosts**.
The **Add Hosts** wizard opens.
3. On the **Add hosts**, specify the hosts that you want to add to the cluster and click **Next**.
 - Add a host that is not in the same vCenter Server inventory.
 - a. Click the **New Hosts** tab.
 - b. Enter the required information about the host in the text boxes.
 - c. To add more new hosts, click the **Add Host** button and enter the required information.
 - d. Select the **Use the same credentials for all hosts** check box.
 - Add a host that is in the same vCenter Server inventory.
 - a. Click the **Existing hosts** tab.
 - b. Select one or multiple hosts from the list.

Note: If you want to add multiple hosts, you do not need to specify only new hosts or only existing hosts. You can specify new hosts and select from the existing hosts at the same time.

4. On the **Host Summary** page, review the information about the hosts and click **Next**.
5. On the **Import Image** page, select the host whose image to use as the image for the cluster.
 - To add the specified hosts to the cluster without changing the current image for that cluster, select the **Don't import an image** radio button.
 - To use any of the specified hosts as a reference host and use its image as the new image for that cluster, select the **Select which host to import the image from** radio button and select a host from the list.

Note: If the image on the host that you select is of version earlier than ESXi 7.0 Update 2, you cannot proceed with importing the image from that host.

When you select a host to import an image from, a card with information about the image appears. you can view the ESXi version, the vendor add-on, and the additional components that the image on the reference host contains.

6. If you selected the **Select which host to import the image from** radio button, click **Show Cluster's Current Image** at the bottom of the page.

A card with information about the current image for the cluster appears. Before you proceed, you can review and compare the current cluster image with the image on the selected host

before you proceed. In this way, you can ensure that you selected the right reference host.

Add hosts

- Add hosts
- Host Summary
- Import Image**
- Ready to Complete

Import Image

Optionally, you can import an image from these hosts to set as the new cluster image.

☐ Don't import an image
☒ Select which host to import the image from

Select which host to import the image from
Manually select the host to import the cluster's new image from:

	Host	Cluster	ESXi
☒	10.41.77.204	cls	7.0.2

Image on selected host:

ESXi 7.0 Update 2 - 17536818

0 additional components

SHOW COMPONENTS

- Click **Next**.
- On the **Ready to Complete** page, review the information about the selected hosts and the new cluster image.
- Click **Finish** to add the hosts to the cluster.

The specified hosts are added to the cluster.

If you chose to import an image from a host during the host addition procedure, the image for the cluster changes. The new software specification for the cluster is identical to the one on the selected and added reference host.

If you chose to import an image from a host during the host addition procedure, remediate the cluster to apply the new image to all the hosts.

Managing Standalone ESXi Hosts with vSphere Lifecycle Manager Images

Starting with vSphere 8.0, you can use an image to manage the lifecycle of any standalone ESXi host that is part of your vCenter Server inventory. In vSphere 8.0 Update 1, you can use both, the vSphere Client and the vSphere Automation API, to manage the lifecycle of a standalone host.

A standalone host is a host that is connected to a vCenter Server instance but is not part of any cluster. In your inventory, standalone hosts are the hosts directly under a data center object or in a host folder.

The permissions and privileges that you need to manage a standalone host with an image are the same as the vSphere Lifecycle Manager permissions and privileges that you need for the respective operations at a cluster level.

No limits exist for the number of standalone hosts that you can manage with vSphere Lifecycle Manager images. The real limit is the number of hosts that vCenter Server can support.

By using the vSphere Client and the vSphere Automation API commands, you can start using images on a standalone host that you already manage with baselines. When you take a host out of a cluster that you manage with vSphere Lifecycle Manager image, the host does not retain the image. The host becomes managed by baselines. The reverse action, moving a standalone host that you manage with an image to a cluster that you manage with a single image is not supported.

In vSphere 8.0 Update 1, you can perform on standalone hosts all vSphere Lifecycle Manager operations that you can perform at a cluster level by using the vSphere Client or the respective API commands.

When you add a standalone host to a data center or a folder, the option to use a vSphere Lifecycle Manager image for the host is selected by default. You must keep it selected unless you want to manage the lifecycle of that standalone host with baselines. During the addition of the standalone host, you can either set up an image manually, use the existing image on the host, or import an image from a reference host in the vCenter Server inventory. For information, see [Set Up a New Image to a Host When Adding the Host to a Data Center or Folder](#), [Import an Image from a Reference Host When Adding a Host to a Data Center or Folder](#), and [Use the Current Image of a Host When Adding the Host to a Data Center or Folder](#).

After the standalone host is added to the data center or folder, you can further customize the image. For more information, see [Edit a vSphere Lifecycle Manager Image](#).

For information about how to convert an existing standalone host to a host that is managed with a vSphere Lifecycle Manager image and all available workflows for managing a standalone host with images, see [How to Use vSphere Lifecycle Manager Images](#).

Use the Current Image of a Host When Adding the Host to a Data Center or Folder

When you add a standalone host to a data center or folder in vCenter Server, you can use the current image on the host to manage its lifecycle.

- Review the requirements for using a single image in [What Are the Requirements for Using vSphere Lifecycle Manager](#).
- Verify that a data center or a folder exists in the inventory.
- Verify that the vCenter Server instance is version 8.0 Update 3 or later.
- Verify that the host is of ESXi version 7.0 or later.
- Obtain the user name and password of the root user account for the host.

- Verify that you have an ESXi image available in the vSphere Lifecycle Manager depot.
- Verify that you have the proper privileges. For more information, see [vSphere Lifecycle Manager Privileges For Using Images](#).

During the addition of a standalone host to a data center or a folder, you can choose to import an image from a host in the inventory and use this image to manage the lifecycle of the host. After the standalone host is added to the data center or folder, you can further customize the image. For example, you can customize the additional components or configure a hardware support manager and add firmware and drivers add-ons to the image. For information about editing an image, see [Edit a vSphere Lifecycle Manager Image](#).

1. In the vSphere Client, navigate to a data center or folder within a data center.
2. Right-click the data center or folder and select **Add Host**.
The **Add Host** wizard opens.
3. On the **Name and location** page, enter the IP address or the name of the host and click **Next**.
4. On the **Connection settings** page, enter administrator credentials and click **Next**.
5. On the **Host summary** page, review the host summary and click **Next**.
6. On the **Host lifecycle** page, leave the **Manage host with an image** check box selected.
This option is selected by default.
7. Select the **Import image from an existing host in the vCenter inventory** radio button and click **Next**.
8. On the **Select image** page, select the reference host which you want to extract the image from and click **Next**.
After you select the reference host, detailed information about the image appears at the bottom of the page. You can view the ESXi version and all additional components.
9. On the **Assign license** page, add a license to the host and click **Next**.
10. Optional: On the **Lockdown mode** page, select a lock-down mode option to deactivate the remote access for the administrator account after vCenter Server takes control of this host and click **Next**.
11. On the **VM location** page, select a location for the virtual machines that reside on the host and click **Next**.
12. On the **Ready to complete** page, review your selections and the image setup and click **Finish**.

A standalone host that uses a vSphere Lifecycle Manager image appears in the vCenter Server inventory.

View and customize the host image on the **Updates** tab for the standalone host.

Import an Image from a Reference Host When Adding a Host to a Data Center or Folder

When you add a standalone host to a data center or folder, you can import an image from another host in vCenter Server to manage the host lifecycle.

- Review the requirements for using a single image in [What Are the Requirements for Using vSphere Lifecycle Manager](#).
- Verify that a data center or a folder exists in the inventory.
- Verify that the vCenter Server instance is version 8.0 Update 3 or later.
- Verify that the reference host is of ESXi version 7.0 Update 2 or later.
- Verify that the ESX-OSDATA volume is backed by a dedicated NVMe, SSD, or HDD device.
- Obtain the user name and password of the root user account for the host.
- Verify that you have an ESXi image available in the vSphere Lifecycle Manager depot.

- Verify that you have the proper privileges. For more information, see [vSphere Lifecycle Manager Privileges For Using Images](#).

When you add a standalone host to a data center or a folder, you can use the image that is currently present on the host. After the standalone host is added to the data center or folder, you can further customize this image. For example, you can add additional components or configure a hardware support manager and add firmware and drivers add-ons to the image. For information about editing an image, see [Edit a vSphere Lifecycle Manager Image](#).

1. In the vSphere Client, navigate to a data center or folder within a data center.
2. Right-click the data center or folder and select **Add Host**.
The **Add Host** wizard opens.
3. On the **Name and location** page, enter the IP address or the name of the host and click **Next**.
4. On the **Connection settings** page, enter administrator credentials and click **Next**.
5. On the **Host summary** page, review the host summary and click **Next**.
6. On the **Host lifecycle** page, leave the **Manage host with an image** check box selected.
This option is selected by default.
7. Select the **Use the current image on host** radio button and click **Next**.
After you select this option, detailed information about the image on the host appears on this page. You can view the ESXi version and all additional components.
8. On the **Assign license** page, add a license to the host and click **Next**.
9. Optional: On the **Lockdown mode** page, select a lock-down mode option to deactivate the remote access for the administrator account after vCenter Server takes control of this host and click **Next**.
10. On the **VM location** page, select a location for the virtual machines that reside on the host and click **Next**.
11. On the **Ready to complete** page, review your selections and the image setup and click **Finish**.

A standalone host that uses a vSphere Lifecycle Manager image appears in the vCenter Server inventory.

View and customize the host image on the **Updates** tab for the standalone host.

Set Up a New Image to a Host When Adding the Host to a Data Center or Folder

You can add a standalone host to a data center or folder and set up a new vSphere Lifecycle Manager image to manage the lifecycle of the host.

- Review the requirements for using a single image in [What Are the Requirements for Using vSphere Lifecycle Manager](#).
- Verify that a data center or a folder exists in the inventory.
- Verify that the host is of ESXi version 7.0 or later.
- Obtain the user name and password of the root user account for the host.
- Verify that you have an ESXi image available in the vSphere Lifecycle Manager depot.
- Verify that you have the proper privileges. For more information, see [vSphere Lifecycle Manager Privileges For Using Images](#).

When you add a standalone host to a data center or a folder, you must specify the ESXi version for the image and, optionally, a vendor add-on. After the standalone host is added to the data center or folder, you can further customize the image. For example, you can add components to the image. You can also configure a hardware support manager and

add a firmware and drivers add-on to the image. For information about editing an image, see [Edit a vSphere Lifecycle Manager Image](#).

1. In the vSphere Client, navigate to a data center or folder within a data center.
2. Right-click the data center or folder and select **Add Host**.
The **Add Host** wizard opens.
3. On the **Name and location** page, enter the IP address or the name of the host and click **Next**.
4. On the **Connection settings** page, enter administrator credentials and click **Next**.
5. On the **Host summary** page, review the host summary and click **Next**.
6. On the **Host lifecycle** page, leave the **Manage host with an image** check box and the **Compose a new image** radio button selected.
These options are selected by default.
7. Click **Next**.
8. On the **Image** page, set up the desired image and click **Next**.
 - a) Select an ESXi version.
 - b) Optional: Select a vendor add-on and the version of the add-on.
9. On the **Assign license** page, add a license to the host and click **Next**.
10. Optional: On the **Lockdown mode** page, select a lock-down mode option to deactivate the remote access for the administrator account after vCenter Server takes control of this host and click **Next**.
11. On the **VM location** page, select a location for the virtual machines that reside on the host and click **Next**.
12. On the **Ready to complete** page, review your selections and the image setup and click **Finish**.

A standalone host that uses a single image appears in the vCenter Server inventory. You can view and customize the host image on the **Updates** tab for the standalone host.

How to Use vSphere Lifecycle Manager Images

Using vSphere Lifecycle Manager images provides a simplified and unified workflow for patching and upgrade of ESXi hosts. You can also use vSphere Lifecycle Manager images for bootstrapping purposes and firmware updates.

An image defines the exact software stack to run on all ESXi hosts in a cluster or on a standalone host.

General Workflow

Using images to apply software and firmware updates to ESXi hosts is a multi-stage process.

1. Software updates must become available in the vSphere Lifecycle Manager depot.
To set up and use an image, you use the software updates that are available in the vSphere Lifecycle Manager depot. The depot contains base images, vendor add-ons, and additional components.
Updates get into the vSphere Lifecycle Manager local depot through synchronization with configurable download sources. By default, vSphere Lifecycle Manager is configured to synchronize with the official VMware depot. You can also import updates into the depot manually.
You can see the contents of the vSphere Lifecycle Manager depot in the vSphere Lifecycle Manager home view.
For more information, see [The vSphere Lifecycle Manager Depot](#).

2. Start using vSphere Lifecycle Manager images.

vSphere Lifecycle Manager provides you with the option to start using images with the very creation of a cluster. If you do not set up an image during the creation of a cluster, you can switch from using vSphere Lifecycle Manager baselines to using vSphere Lifecycle Manager images at a later time.

Even when you save the image, no software is installed on the ESXi hosts during image setup.

For more information, see [Creating and Managing vSphere Lifecycle Manager Clusters](#), [Managing Standalone ESXi Hosts with vSphere Lifecycle Manager Images](#), and [Convert a Cluster or a Host That Uses Baselines Into a Cluster or a Host That Uses vSphere Lifecycle Manager Images](#).

3. Check the compliance of the ESXi hosts in the cluster or the standalone host against the image specification.

The compliance check compares the current image on the ESXi hosts in the cluster or the standalone host against the desired image that you specified during the setup process, and defines compatibility status of the hosts.

For more information, see [Check the Compliance of a Cluster Against a vSphere Lifecycle Manager Image](#) and [View Host Compliance Information](#).

4. Review the compliance statuses of the hosts in the cluster or the standalone host.

5. You can run a remediation pre-check on an ESXi host to ensure software and hardware compatibility with the image.

Running a remediation pre-check is optional. The remediation pre-check ensures that all requirements for successful remediation are met. For more information, see [Run a Remediation Pre-Check for a Cluster, a Host Within a Cluster, or a Standalone Host](#).

6. Remediate the non-compliant ESXi hosts in the cluster or the non-compliant standalone host.

Remediation is the process through which the software specification defined by the vSphere Lifecycle Manager image that you use for a cluster or a host is actually applied to the hosts in the cluster or to the standalone host.

For more information about remediating hosts against an image, see [Remediating a Cluster or a Standalone Host Against a Single Image](#).

Limitations

- When you set up and save an image for a cluster, the image is not applied to the hosts in the cluster unless you remediate the hosts. The mere action of changing the management method does not alter the hosts in the cluster.
- After you set up an image for the cluster and remediate the hosts in the cluster against the image, standalone VIBs are deleted from the hosts.
- After you set up an image for the cluster and remediate the hosts in the cluster against the image, non-integrated solution agents are deleted from the hosts.

Convert a Cluster or a Host That Uses Baselines Into a Cluster or a Host That Uses vSphere Lifecycle Manager Images

You can manage a cluster or a standalone host with either baselines or images. You cannot use both at the same time for a single cluster or host. Even if you do not set up an image for the cluster during cluster creation or for the host when adding it to the data center or folder, you can at any time switch from using baselines to using images for that cluster or host.

To switch a cluster to vSphere Lifecycle Manager images, you must set up a new image or import an existing one. In vSphere 8.0 Update 1, to switch a standalone host to a vSphere Lifecycle Manager images, you can either set up a new image or import an existing one. Before you proceed with setting up or importing an image, vCenter Server checks and reports if the cluster or host is eligible for using images. For more information about cluster or host eligibility, see [Cluster or Standalone Host Eligibility to Use vSphere Lifecycle Manager Images](#).

System Requirements

To switch to using images, the cluster or host must meet multiple requirements.

- All ESXi hosts in the cluster or the standalone host must be of version 7.0 and later.

- All ESXi hosts in the cluster or the standalone host must be stateful.
A stateful install is one in which the host boots from a disk.
- No host in the cluster or standalone host can contain any unknown components.

If a host is of version earlier than ESXi 7.0, you must first use an upgrade baseline to upgrade the host and then you can successfully switch to using images. For more information about using baselines for host patching and upgrade operations, see [Using vSphere Lifecycle Manager Baselines and Baseline Groups](#).

For more information about converting a stateless host into a stateful host, find information about Auto Deploy in the *VMware ESXi Installation and Setup* documentation.

Specifics of the Transitioning Workflow

Several behavioral specifics exist when you switch to using vSphere Lifecycle Manager images.

- If you switch to using images, you cannot revert to using baselines for the cluster. You can move the hosts to a cluster that uses baselines, but you cannot change a cluster that already uses a single image for management purposes.
- If you switch to using images for a standalone host, you cannot revert to using baselines for the host except if you re-add the standalone host to the vCenter Server inventory.
- When you set up and save an image for a cluster or host, the image is not applied to the hosts in the cluster or to the standalone host unless you remediate the hosts. The mere action of changing the management method does not alter the hosts in the cluster or the standalone host.
- After you set up an image for the cluster or the standalone host and remediate the hosts in the cluster or the standalone host against the image, standalone VIBs are deleted from the hosts.
- After you set up an image for the cluster or host and remediate the hosts in the cluster or the standalone host against the image, non-integrated solution agents are deleted from the hosts.
- If you enable a solution that cannot work with vSphere Lifecycle Manager, for example Dell EMC VxRail, on an empty cluster and attempt to switch to using an image for that cluster, the transition operation succeeds. However, the result is an unsupported cluster configuration, because both vSphere Lifecycle Manager and the non-integrated solution are enabled on the cluster.
- If you take a host out of a cluster that you manage with a vSphere Lifecycle Manager image, the host becomes standalone. The host can retain the image of the cluster in case the following criteria are met:
 - The vCenter Server and ESXi host are of version vSphere 8.0 Update 3 and later.
 - The host does not have the NSX component installed.
 - If the host contains the NSX component, NSX is of version 4.2 or later. See [Using vSphere Lifecycle Manager Images for Standalone Hosts with NSX 4.2 and later](#).
 - During the move operation the host stays connected or reachable.
 Otherwise, the image is removed from the host and the host is reverted back to using baselines.
- The reverse action, moving a standalone host that you manage with a single image to a cluster that you manage with an image is supported. After you move the standalone host managed with a vSphere Lifecycle Manager image, you must check the compliance of the hosts in the cluster against the specified cluster image.

Cluster or Standalone Host Eligibility to Use vSphere Lifecycle Manager Images

Switching from baselines to images requires that you set up or import a vSphere Lifecycle Manager image to manage the cluster or host with. As part of the transition, before you set up the image, vCenter Server triggers an automatic task that checks whether the cluster or host is eligible to use vSphere Lifecycle Manager images.

The **Check cluster's eligibility to be managed with a single image** and **Check host's eligibility to be managed with a single image** tasks ensure that the cluster or host is not undergoing remediation against a baseline and check whether all requirements for using vSphere Lifecycle Manager images are met.

Also, the tasks check for standalone VIBs and ensure that no unintegrated solutions are enabled for the cluster or host. You might not be able to switch to using vSphere Lifecycle Manager images if unintegrated solutions are enabled on the cluster or host.

Note:

If you use any third-party products or solutions, you must confirm with your third-party software vendor whether the respective solution works with vSphere Lifecycle Manager.

The tasks return three types of notifications: error, warning, and info.

Errors

The **Check cluster's eligibility to be managed with a single image** and **Check host's eligibility to be managed with a single image** tasks report an error in case one of the following conditions occurs:

- The cluster contains at least one host that is not stateful or the standalone host is not stateful.
- The cluster contains at least one host that is not of a compatible ESXi version or the standalone host is of version different than 7.0 and later.
- The cluster or the standalone host contain VIBs of unintegrated solutions. In that case, you must deactivate the unintegrated solution and retry the transition.

Warnings

The **Check cluster's eligibility to be managed with a single image** and **Check host's eligibility to be managed with a single image** tasks issue a warning if the cluster contains at least one host with or the standalone host contains a standalone VIB or an unknown VIB. Warnings do not block the transition to using vSphere Lifecycle Manager images, but they require special attention or a user action.

For example, you see a warning notification if a host in the cluster contains a standalone VIB, for example a driver, for which a component is available in the vSphere Lifecycle Manager depot. If you want to keep the VIB, you must add the respective component to the vSphere Lifecycle Manager image. Otherwise, the standalone VIB is deleted upon remediation.

You also get a warning if a host in the cluster or the standalone host contain an unknown VIB. Unknown VIBs are standalone VIBs for which no component is available in the vSphere Lifecycle Manager depot. If vSphere Lifecycle Manager detects an unknown VIB, you must import a component that contains the VIB into the vSphere Lifecycle Manager depot and restart the transition. Otherwise, the unknown VIB is deleted upon remediation.

Info

The **Check cluster's eligibility to be managed with a single image** and **Check host's eligibility to be managed with a single image** tasks give an info notification if the cluster contains at least one host with a standalone VIB or the standalone host contains a standalone VIB, but you can still proceed to setting up a vSphere Lifecycle Manager image for the cluster or host without any additional actions.

For example, you see an info notification if the cluster is enabled for an integrated solution, for example vSphere HA or vSAN.

Option 1: Use an Image from a Host in the Cluster

To start managing an existing cluster or host that uses baselines with a single vSphere Lifecycle Manager image, you can select and use one of the images installed on a host in the cluster or the current image on the standalone host.

- Verify that all ESXi hosts in the cluster or the standalone host are of version 7.0 Update 2 and later.

- Verify that your vCenter Server instance is of version 8.0 Update 3 and later.
- Verify that all ESXi hosts in the cluster or the standalone host are stateful. A stateful install is one in which the host boots from a disk.
- Verify that all ESXi hosts in the cluster are from the same hardware vendor.
- Verify that no unintegrated solution is enabled for the cluster or the host.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

If you switch to using images, you cannot revert to using baselines for the cluster. You can move the hosts to another cluster, which uses baselines, but you cannot change the cluster that already uses a single image.

If you switch to using images for a standalone host, you cannot revert to using baselines for the host.

For conceptual information about vSphere Lifecycle Manager images, see [vSphere Lifecycle Manager Images](#).

For information about how to use vSphere Lifecycle Manager images to manage hosts and clusters, see [How to Use vSphere Lifecycle Manager Images](#).

With vSphere 8.0 Update 3, instead of importing or manually setting up an image, you can use an image from a host that is already in the cluster or the image currently running on the standalone host.

When you initiate the transitioning workflow for a cluster that uses baselines, vSphere Lifecycle Manager triggers the **Get Installed Host Images** task after the eligibility checks finish successfully. During this task, vSphere Lifecycle Manager retrieves information about all images currently installed on the hosts in the cluster. Then vSphere Lifecycle Manager calculates which are the most complete, valid, and suitable images that you can use to remediate the cluster. The list of suitable images is generated based on two criteria: the ESXi version and the software running on the hosts in the cluster. Based on these criteria, vSphere Lifecycle Manager provides the following image options to choose from:

- The image that contains the highest ESXi version.
- The image that has the most widely used software on most of the hosts in the cluster.
- The images that have a different software setup than the most widely used setup on the hosts in the cluster.

After you select an image from a host in the cluster, you can make changes to the image or save it without editing.

When you decide to add a host to or remove a host from a cluster that already has some generated list of suitable images, this list is invalidated. vSphere Lifecycle Manager automatically triggers a new **Get Installed Host Images** task. During this task, vSphere Lifecycle Manager retrieves information about the new set of images installed on the hosts in the cluster and computes a new list of suitable images based on the cluster membership changes. You can also regenerate the list of images at any time.

When you start converting a standalone host that uses baselines to a host that is managed with an image, you can view and use the image currently installed on the host. vSphere Lifecycle Manager automatically retrieves information about

the current image on the standalone host. You can select and edit the current image on the standalone host or use it without changes.

1. In the vSphere Client, navigate to a cluster or host that you manage with baselines and baseline groups.
2. On the **Updates** tab, click **Image**.

vSphere Lifecycle Manager starts checking if the cluster or host is eligible for using images. If no problems are reported, vSphere Lifecycle Manager retrieves information about the images currently installed on the hosts in the cluster or the image running on the standalone host.

- If you selected a cluster, a list of the most valid and suitable images for the cluster is displayed in the **Currently running images in the cluster** page section.

Cluster managed with baselines

ACTIONS

Summary

Monitor

Configure

Permissions

Hosts

VMs

Datastores

Networks

Hosts

Baselines

Image

VMware Tools

VM Hardware

⚠

Transition from vSphere Baselines to vSphere Lifecycle Manager Images is not supported in vSphere Foundation.

Manage with a single image

Lifecycle Manager enables you to have all hosts in a cluster inherit the same image, the same image means faster upgrades, improved reliability and easier overall maintenance.

① For the list of host requirements to be managed with a single image, please refer to the [vSphere Lifecycle Manager documentation](#).

Currently running images in the cluster

The following images have been detected running on hosts in this cluster.
Checked on 06/09/2024, 7:23:39 PM [Check again](#)

✓

ESXi 8.0 Update 3 - 23870251

No additional components

📱

1/2 hosts is running this image

✓ This is the highest versioned image in this cluster

1 more images found

[VIEW ALL IMAGES](#)

PROCEED WITH THIS IMAGE

Other ways to set up an image for the cluster

IMPORT IMAGE

SETUP IMAGE MANUALLY

VMware by Broadcom

1625

- If you selected a standalone host, you can view the image card with details for the current host image in the **Get started quick** page section.

The screenshot displays the VMware vSphere 8.0 interface for a standalone host with IP 10.168.171.169. The left sidebar shows the 'Hosts' menu with 'Image' selected. The main content area is titled 'Manage with an image' and includes a description of Lifecycle Manager. A 'Get started quick' section provides instructions on using the image on the host, featuring a card for 'ESXi 8.0 U1c - 22543607' with '0 additional components' and a 'USE IMAGE ON HOST' button. Below this, a 'Use a different image' section offers options to 'IMPORT IMAGE' or 'SETUP IMAGE MANUALLY'.

10.168.171.169 | ACTIONS

Summary Monitor Configure Permissions VMs Resource Pools Datastores Networks

Hosts

- Baselines
- Image**
- Hardware Compatibility
- VMware Tools
- VM Hardware

Manage with an image

Lifecycle Manager enables you to manage a host with an image which means less maintenance.

① For the list of host requirements to be managed with a single image, please refer to the [documentation](#).

Get started quick

The easiest and quickest way to set up an image is to use the image on the host.


ESXi 8.0 U1c - 22543607

0 additional components

USE IMAGE ON HOST

Use a different image

Alternatively you can set up the host with a different image. Import an image or setup manually.

IMPORT IMAGE **SETUP IMAGE MANUALLY**

- If you selected a cluster, select a host image from the provided list of images and click **Proceed with selected image** to confirm your selection.

The **Convert to an Image** pane appears. The elements of the selected host image appear in the **Define Image** card.

4. If you selected a standalone host, click **Use image on host**.
The **Convert to an Image** pane appears. The elements of the current host image appear in the **Define Image** card.
5. If you want to further customize the host image, edit any of the image elements.
6. Optional: To validate the image, click the **Validate** button.
You validate an image to check for missing dependencies and component conflicts.
7. Click **Save**.
Saving the image triggers an automatic compliance check. All hosts in the cluster or the standalone host are checked against the image.
8. In the **Convert to an Image** pane, finish the image setup.
 - a) Click the **Finish image setup** button.
 - b) In the **Finish image setup** dialog box, click **Yes, finish image setup**.

You now manage all hosts in the cluster collectively or the standalone host with a single image. Upon remediation, the image is installed on all hosts in the cluster or on the standalone host.

To apply the image to all hosts in the cluster or the standalone host, remediate the cluster or host against the image.

Option 2: Import an Image from a JSON File

To start managing with a single vSphere Lifecycle Manager image an existing cluster or host that already uses baselines, you can import an image from another cluster or host instead of setting up a new image manually.

- Verify that all ESXi hosts in the cluster or the standalone host are of version 7.0 and later.
- Verify that all ESXi hosts in the cluster or the standalone host are stateful. A stateful install is one in which the host boots from a disk.
- Verify that all ESXi hosts in the cluster are from the same hardware vendor.
- Verify that no unintegrated solution is enabled for the cluster or the host.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

If you switch to using images, you cannot revert to using baselines for the cluster. You can move the hosts to another cluster, which uses baselines, but you cannot change the cluster that already uses a single image.

If you switch to using images for a standalone host, you cannot revert to using baselines for the host.

For conceptual information about vSphere Lifecycle Manager images, see [vSphere Lifecycle Manager Images](#).

For information about how to use vSphere Lifecycle Manager images to manage hosts and clusters, see [How to Use vSphere Lifecycle Manager Images](#).

1. In the vSphere Client, navigate to a cluster or host that you manage with baselines and baseline groups.
2. On the **Updates** tab, click **Image**.
3. Click the **Import Image** button.
The **Import Image** dialog box appears.
4. Select a JSON file to import and click **Next**.
 - Click the **Browse** button and select a JSON file on your local machine.
 - Enter a valid URL to a JSON file on a remote web server.

vSphere Lifecycle Manager starts checking if the cluster or host is eligible for using images. If no problems are reported, the **Convert to an Image** pane appears. The elements of the imported image appear in the **Define Image** card.

5. Optional: Customize the imported image by changing any of its elements.
6. Optional: To validate the image, click the **Validate** button.
You validate an image to check for missing dependencies and component conflicts.
7. Click **Save**.
Saving the image triggers an automatic compliance check. All hosts in the cluster or the standalone host are checked against the image.
8. In the **Convert to an Image** pane, finish the image setup.
 - a) Click the **Finish image setup** button.
 - b) In the **Finish image setup** dialog box, click **Yes, finish image setup**.

You now manage all hosts in the cluster collectively or the standalone host with a single image for the cluster or host. Upon remediation, the image is installed on all hosts in the cluster or on the standalone host. To apply the image to all hosts in the cluster or the standalone host, remediate the cluster or host against the image.

Option 3: Set Up a New Image Manually

You can start using vSphere Lifecycle Manager images for an existing cluster or host that uses baselines by setting up an image manually.

- Verify that all ESXi hosts in the cluster or the standalone host are of version 7.0 and later.
- Verify that all ESXi hosts in the cluster or the standalone host are stateful. A stateful install is one in which the host boots from a disk.
- Verify that all ESXi hosts in the cluster are from the same hardware vendor.
- Verify that no unintegrated solution is enabled for the cluster or the host.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

If you switch to using images, you cannot revert to using baselines for the cluster. You can move the hosts to another cluster, which uses baselines, but you cannot change the cluster that already uses a single image.

If you switch to using images for a standalone host, you cannot revert to using baselines for the host.

For conceptual information about vSphere Lifecycle Manager images, see [vSphere Lifecycle Manager Images](#).

For information about how to use vSphere Lifecycle Manager images to manage hosts and clusters, see [How to Use vSphere Lifecycle Manager Images](#).

1. In the vSphere Client, navigate to a cluster or host that you manage with baselines and baseline groups.
2. On the **Updates** tab, click **Image**.
3. Click the **Set up image manually** button.

vSphere Lifecycle Manager starts checking if the cluster or host is eligible for using images. If no problems are reported, the **Convert to an Image** pane appears.

4. If the **Check cluster's eligibility to be managed with a single image** or **Check host's eligibility to be managed with a single image** tasks report an error or a warning that requires an action, resolve the issue and restart the procedure.
5. From the **ESXi Version** drop-down menu, select an ESXi image.
6. Optional: Add a vendor add-on to the image.
 - a) Click **Select**.
The **Select Vendor Addon** appears.
 - b) Select an add-on from the list.
An information panel appears on the right. The information panel shows information about the components that the add-on adds to the ESXi image and the components that it removes from the image.
 - c) From the **Version** column use the drop-down menu for the selected add-on and select the add-on version.
 - d) Click **Select**.
7. Optional: Add a firmware and drivers add-on to the image.
 - a) Click **Select**.
The **Select Firmware and Drivers Addon** appears.
 - b) Select a hardware support manager from the respective drop-down menu.
A list of firmware and drivers add-ons appears.
 - c) Select an add-on from the list.
An information panel appears on the right. The information panel shows information about the supported ESXi versions and whether the add-on contains a driver or not.
 - d) From the **Version** column use the drop-down menu for the selected add-on and select the add-on version.
 - e) Click **Select**.
8. Optional: Manage the additional components of the image.
 - a) Click **Show details**.
 - b) Click **Add components**.
The **Add Components** dialog box appears.
 - c) Optional: Use the **Show** drop-down menu to sift out the components that are not part of the selected vendor add-on.
 - d) Select one or multiple components from the list.
An information panel appears on the right. The information panel shows information about the component that you selected first.
 - e) From the **Version** column use the drop-down menu for the selected component to select the component version.
CAUTION: If your vCenter Server instance and all ESXi hosts in the cluster are of version 8.0 Update 3 or later, you can select and retain a lower version for an additional component that is part of the ESXi base image. To avoid critical errors on any of the hosts in the cluster, make sure you perform thorough testing of the customized image before remediating the cluster.
 - f) Click **Select**.
The selected components appear in the list of components that the image contains. You can use the **Show** drop-down menu to sift out the additional components. For example, you can list only the customized components.
 - g) Optional: From the **Notes** column, click the trash icon to remove a component.
CAUTION: If your vCenter Server instance and all ESXi hosts in the cluster are of version 8.0 Update 3 or later, you can remove the VMware Host Client and ESXi VM Tools components from the base image. To avoid critical

errors on any of the hosts in the cluster, make sure you perform thorough testing of the customized image before remediating the cluster.

h) Optional: Click **Hide details** to hide the list of components.

9. Optional: To validate the image, click the **Validate** button.

You validate an image to check for missing dependencies and component conflicts.

10. Click **Save**.

Saving the image triggers an automatic compliance check. All hosts in the cluster or the standalone host are checked against the image.

11. In the **Convert to an Image** pane, finish the image setup.

a) Click the **Finish image setup** button.

b) In the **Finish image setup** dialog box, click **Yes, finish image setup**.

You set up an image for the cluster or standalone host. You now manage all hosts in the cluster collectively or the standalone host with a vSphere Lifecycle Manager image for the cluster or host. Upon remediation, the image is installed on all hosts in the cluster or the standalone host.

To apply the image to all hosts in the cluster or the standalone host, remediate the cluster or host against the image.

Working with Images

When you manage a cluster or host with a single image, you can change the image at any time. You can edit the image by changing the software that it includes, for example, you can add or remove components, and you can also change the version of the included components. You can also reuse an image by exporting it and importing it into a different cluster or a standalone host.

Using vSphere Lifecycle Manager images starts with setting up an image for a cluster or a host. After you start managing a cluster or a host with a single image, you can edit the image at any time. You can validate the image before saving it to verify that it includes no conflicting components or missing dependencies.

Setting Up an Image for a Cluster or a Host

You can set up an image during the creation of the cluster or when you add a standalone host to a data center or a folder or later. During the creation of a cluster or the addition of a standalone host to a data center or a folder, you can only define the ESXi version and, optionally, a vendor add-on to be included in the image for the cluster or host. You can later edit the image to include additional components or a firmware add-on. For detailed information about creating a cluster and adding hosts to it, and about adding a standalone host to a data center or folder, see the *vCenter Server and Host Management* documentation.

Note:

When you set up an image, you select an ESXi version and a vendor add-on from the vSphere Lifecycle Manager depot. If no ESXi base images and vendor add-ons are available in the vSphere Lifecycle Manager depot, you must populate the depot with software updates by synchronizing the depot or uploading updates to the depot manually. For detail information about the corresponding procedures, see [Synchronize the vSphere Lifecycle Manager Depot](#) and [Import Updates to the vSphere Lifecycle Manager Depot](#).

If you do not set up an image during the creation of the cluster or the addition of the standalone host to a data center or folder, the cluster or the host use baselines, but you can switch to images at any time. When you set up an image during the transition workflow, you can define the full stack of software to run on the hosts in the cluster or on the standalone host. For more information about switching from baselines to images, see [Convert a Cluster or a Host That Uses Baselines Into a Cluster or a Host That Uses vSphere Lifecycle Manager Images](#).

Whereas switching from baselines to images is possible, the reverse operation is not. If a cluster or a host use a single image, regardless of whether you set up the image during the cluster creation or the host addition to the data center or folder, you cannot switch to using baselines for that cluster or standalone host.

Viewing Details About the Image of a Cluster or Host

You can view the software elements that the image for a cluster or host contains on the **Updates** tab for that cluster or host.

The **Image** pane consists of two cards.

The **Image** card contains information about the image that the cluster or host uses. In that card, you perform all image-related operations. You edit the image, you export the image, you validate your selections, and so on. You can also check and view the recommendations that VMware provides.

The **Image Compliance** card contains compliance information about the hosts in the cluster or the standalone host. In that card, you perform host-related operations. You check the compliance of the hosts in the cluster or the standalone host, you run remediation pre-checks, you remediate the hosts in the cluster or the standalone host, and so on.

In the **Image Compliance** card, you edit remediation settings for that cluster or host and manage depot overrides.

Edit a vSphere Lifecycle Manager Image

For a cluster or host that you manage with a single image, you can edit the image at any time to add, remove, or update an image element. For example, you can edit the image to update the vendor add-on version that it includes, to add or remove a driver, to upgrade the ESXi version in the image, and so on.

Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

Working with Image Drafts

When you edit an image, vSphere Lifecycle Manager saves the working copy of the image as a draft. The draft is an edited but unsaved version of an image. If you edit an image but for some reason you do not save the new image set-up, when you restart editing the image, you can use the saved draft version as a starting point or you can altogether discard the changes that you previously made.

Validation

You can validate an image draft before you save it. Validation checks whether the image is correct and complete. During validation, vSphere Lifecycle Manager checks for missing dependencies and conflicting components. In case of issues, vSphere Lifecycle Manager returns messages with information about the existing issues. You must resolve all issues before you can save the image.

Editing an image and validating a draft before saving it are supported operations for clusters with hosts or standalone hosts that have a DPU device. During validation, vSphere Lifecycle Manager validates both the VIBs that are applicable to the ESXi version on the host and to the ESXi version on a DPU device.

1. In the vSphere Client, navigate to a cluster or host that you manage with a vSphere Lifecycle Manager image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. In the **Image** card, click the **Edit** button.
4. In the **Edit Image** card, modify the image set-up.

Image Element	Possible Modifications
ESXi Version	From the ESXi Version drop-down menu, select a new ESXi base image.

Image Element	Possible Modifications
Vendor Addon	- To add a vendor add-on to the image, click Select and select a vendor add-on. - To change the version of the vendor add-on in the image or to select a new vendor add-on, click the pencil icon and make your changes. - To remove the vendor add-on from the image altogether, click the trash icon.
Firmware and Drivers Addon	- To add a firmware add-on to the image, click Select. In the Select Firmware and Drivers Addon dialog box, specify a hardware support manager and select a firmware add-on to add to the image. - To select a new firmware add-on, click the pencil icon and make your changes. - To remove the firmware add-on element from the image altogether, click the trash icon. Selecting a firmware add-on for a family of vendor servers is possible only if the respective vendor-provided hardware support manager is registered as an extension to the vCenter Server where vSphere Lifecycle Manager runs.
Components	Click Show details and view the list of additional components in the image. - To add components to the image, click Add Components and select the components and their corresponding versions to add to the image. CAUTION: If your vCenter Server instance and all ESXi hosts in the cluster are of version 8.0 Update 3 or later, you can select and retain a lower version for an additional component that is part of the ESXi base image. To avoid critical errors on any of the hosts in the cluster, make sure you perform thorough testing of the customized image before remediating the cluster. - To delete a component from the image, click the trash icon in the table with components. CAUTION: If your vCenter Server instance and all ESXi hosts in the cluster are of version 8.0 Update 3 or later, you can remove the VMware Host Client and ESXi VM Tools components from the base image. To avoid critical errors on any of the hosts in the cluster, make sure you perform thorough testing of the customized image before remediating the cluster. - To delete a manually added component that overrides a component in the selected vendor add-on or base image, click the undo icon in the table with components. This action reverts the override.

5. Optional: To validate the image, click the **Validate** button.

You validate an image to check for missing dependencies and component conflicts.

6. Click **Save**.

The save operation triggers validation. If the image is valid, vSphere Lifecycle Manager saves it and runs a compliance check against the new image. You can view compliance information in the **Image Compliance** card.

If the image is invalid, saving the image fails and vSphere Lifecycle Manager returns an error.

The new image is validated and displayed in the **Image** card. vSphere Lifecycle Manager performs an automatic hardware compatibility check against the new image.

If there are recommended images generated for the cluster or host, those recommendations become invalidated and updated. vSphere Lifecycle Manager automatically generates a new recommendation based on the new image.

Export or Import a vSphere Lifecycle Manager Image

You can reuse an image that you already set up for a cluster by exporting it and importing it to another cluster in the same or different vCenter Server instance. Depending on your goal, the image can be exported, distributed, and consumed in three different formats.

You can export the image of a cluster as an ISO image, ZIP file, or a JSON file.

You can import an image to a cluster that uses vSphere Lifecycle Manager images only in the JSON format.

In vSphere 8.0, the export and import operations work for clusters with hosts that have DPU devices.

ISO Image

Distributing an image created with vSphere Lifecycle Manager in an ISO format is useful when you need the image to perform clean installs of ESXi and for bootstrapping purposes, for example the kickstart workflow.

You cannot use an image exported as an ISO file with another cluster that uses vSphere Lifecycle Manager images.

ZIP File

Distributing an image created with vSphere Lifecycle Manager as an offline bundle is useful when you want to import the components that the image contains into the depot of the target vSphere Lifecycle Manager instance.

Unlike the ISO image, you cannot use a ZIP file to create upgrade baselines. You also cannot use the ZIP file to create a vSphere Lifecycle Manager image for a cluster.

JSON File

Distributing an image created with vSphere Lifecycle Manager as a JSON file is useful when you want to reuse the same image for other clusters that use images for host management.

When you distribute the JSON file to clusters in a different vCenter Server instance, you must make sure that the depot of the target vSphere Lifecycle Manager instance contains all components that the JSON file contains.

The JSON file contains only metadata and not the actual software payloads.

Which Distribution Format to Choose?

To reuse an existing image for a cluster in the same vCenter Server system, you must export the image as a JSON file and then import the JSON file to the target cluster.

However, when you want to use an existing image for a cluster in another vCenter Server instance, exporting the image as a JSON file might not be enough. You might also need to export the image as a ZIP file. At the target location, you must import the JSON file as an image to the target cluster. But you might also need to import the ZIP file to the target vSphere Lifecycle Manager depot to make sure that all components included in the image are available to the target vSphere Lifecycle Manager instance.

To reuse an existing cluster image for a cluster that you manage with baselines, you must export the source image as an ISO image. You must then import the ISO image into the local depot of the target vSphere Lifecycle Manager instance and use the imported ISO file to create an upgrade baseline.

Export an Image

You export an image when you want to use the same image for another cluster or standalone host in the same or in a different vCenter Server instance.

Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

Depending on your goals, you can export an image as a JSON file, as an installable ISO image, or as an offline bundle that contains all software packages included in the image. The export format depends on your needs and goals.

For example, if you intend to use the image for a cluster or host in another vCenter Server, you must export it as a JSON file and as a ZIP file. Afterwards, you must import both the JSON file and the ZIP file to the target vCenter Server system. For information about importing updates to the vSphere Lifecycle Manager depot, see [Import Updates to the vSphere Lifecycle Manager Depot](#).

1. In the vSphere Client, navigate to a cluster or host that you manage with a vSphere Lifecycle Manager image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. Click the horizontal ellipsis icon and select **Export**.
4. In the **Export Image** dialog box, select a file format, and click **Export**.

You can export the image in one file format at a time.

The exported file is saved on your local machine.

Import the image to a target cluster or standalone host in the same or in a different vCenter Server instance.

Import an Image

Instead of setting up a new image manually, you can reuse an existing image by importing it to a cluster or a standalone host. Upon remediation, the imported image is applied to all hosts in the cluster or to the standalone host.

- Verify that the vSphere Lifecycle Manager depot contains all components included in the image that you import.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

You can import an image only if it is in a JSON format. The JSON file contains only the image metadata, but not the actual software payloads. To successfully import an image to a cluster or host and apply the software specification to the hosts in the cluster or to the standalone host, all the components must specified in the image be available in the vSphere Lifecycle Manager depot.

So, if you want to distribute and reuse an image across vCenter Server instances, importing the JSON file might not be enough if the components from the image are not available in the target vSphere Lifecycle Manager depot. In such cases, before you import the JSON file to the target cluster or host, you must first import an offline bundle that contains all components included in the image to the target vSphere Lifecycle Manager depot. If you try to import a JSON file to a cluster or host but the target vSphere Lifecycle Manager depot does not contain the corresponding components, the import operation fails due to validation errors.

For information about importing updates to the vSphere Lifecycle Manager depot, see [Import Updates to the vSphere Lifecycle Manager Depot](#).

1. In the vSphere Client, navigate to a cluster or host that you manage with a vSphere Lifecycle Manager image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. Click the horizontal ellipsis icon and select **Import**.
4. In the **Import Image** dialog box, select a JSON file and click **Next**.
 - Enter a URL address to the JSON file that you want to import.
 - Browse to a JSON on your local machine.

5. Optional: In the **Edit Image** card, modify the image set-up.

Image Element	Possible Modifications
ESXi Version	From the ESXi Version drop-down menu, select a new ESXi base image.
Vendor Addon	- To add a vendor add-on to the image, click Select. - To change the version of the vendor add-on in the image or to select a new vendor add-on, click the pencil icon. - To remove the vendor add-on element from the image altogether, click the trash icon .
Firmware and Drivers Addon	- To add a firmware add-on to the image, click Select. - To select a new firmware add-on, click the pencil icon. - To remove the firmware add-on element from the image altogether, click the trash icon. Selecting a firmware add-on for a family of vendor servers is possible only if the respective vendor-provided hardware support manager is registered as an extension to the vCenter Server where vSphere Lifecycle Manager runs.
Components	Click Show details and view the list of additional components in the image. - To add components to the image, click Add Components and select components to add to the image. - To delete a component from the image, click the trash icon in the table with components. - To delete a manually added component that overrides a component in the selected vendor add-on, click the undo icon in the table with components. This action reverts the override.

6. If the image contains conflicting components or unresolved dependencies, resolve the issues and retry the procedure.

7. Optional: To validate the image, click the **Validate** button.

You validate an image to check for missing dependencies and component conflicts.

8. Click **Save**.

A compliance check task is automatically triggered. You can view compliance information in the **Image Compliance** card.

The imported JSON file is imported and set as your new image for the target cluster or host. At that stage, nothing is installed on the hosts in the cluster or on the standalone host. The installation of software on the hosts happens during remediation.

Remediate the hosts in the cluster or the standalone host against the new image. See [Run a Remediation Pre-Check for a Cluster, a Host Within a Cluster, or a Standalone Host](#) and [Remediate a Cluster Against a Single Image](#).

Check the Compliance Against a Single Image

When you check the compliance of a cluster or host against an image, vSphere Lifecycle Manager compares the software on each host in the cluster or on the standalone host with the software specified in the image. If the image contains a firmware and drivers add-on, the compliance check also calculates the firmware compliance of the hosts with the image. If some of the hosts in the cluster or the standalone host contain DPU devices, vSphere Lifecycle Manager compares the software and firmware on the DPU devices on the hosts with the software and firmware specified in the cluster or host image.

In addition to calculating the compliance state for each host in the cluster or the standalone host, the compliance check gives you information about the impact that the remediate operation will have on a single host, for example if remediation will cause host reboot or if maintenance mode is needed for the host.

A host can have any of the four compliance states: compliant, non-compliant, incompatible, and unknown. If a host has a DPU device, the compliance state of the DPU device is merged into the overall host compliance result.

Compliant

A host is compliant if the image on the host matches the image that you set for the cluster or host.

A DPU-backed host is compliant when the software and firmware on the server and on the DPU device are the same as the software and firmware specified in the image for the cluster or host.

Non-Compliant

A host is non-compliant if the image on the host does not match the image that you set for the cluster or host. A compliant host becomes non-compliant when you set a new image for the cluster or host or manually add or remove components on the host. You remediate non-compliant hosts to make them compliant.

For example, a host is non-compliant in the following cases.

- The ESXi version on the host is earlier than the ESXi version included in the image for the cluster or host.
- The firmware on the host is different from the firmware add-on in the image for the cluster or host.
- The host has a component that is not included in the image for the cluster or host.
- The host contains a standalone VIB.

A DPU-backed host is non-compliant in the following cases:

- ESXi on the server, ESXi on the DPU device, or both have software components of earlier versions than the software and firmware specified in the image for the cluster or host.
- You edit the image for the cluster or host and add a component that contains one or multiple VIBs that are applicable to ESXi on the server, ESXi on the DPU device, or both platforms.
- The remediation of a cluster or a host fails and as a result, the software and firmware components on the host are updated, but the DPU is not updated or the reverse.

Incompatible

A host is incompatible when the image for the cluster or host cannot be applied to the host.

For example, a host is incompatible in the following cases.

- The ESXi version on the host is later than the ESXi version included in the image for the cluster or host.
- The host does not have sufficient resources, for example RAM.
- The hardware of the host is incompatible with the vSphere Lifecycle Manager image for the cluster or host.

A DPU-backed host has the incompatible compliance state when the ESXi version on the DPU is later than the ESXi version included in the image for the cluster or host.

Unknown

The unknown compliance state indicates that there is no compliance information about the host.

For example, the compliance state of a host is unknown in the following cases.

- You add a new host to the cluster. The compliance state of the newly added hosts is unknown until you perform a compliance check operation on the cluster.

- You add a new standalone host to the data center or folder. The compliance state of the standalone host is unknown until you perform a compliance check operation on the host.
- You edit the image for the cluster or host and save the modifications. The compliance state of all hosts in the cluster or the standalone host is unknown until you check the compliance of the cluster or host against the new image.

Compliance information about the hosts in a cluster or the standalone hosts is displayed on the **Updates** tab for that cluster or host, in the **Image Compliance** card. If you select a cluster, the **Image Compliance** card displays a list of all hosts in the cluster that are out of compliance with the image for the cluster. When you select a host from the cluster, the compliance information about the host appears on the right. If you select a standalone host, the **Image Compliance** card displays the compliance information about the host.

If a host has a DPU device, it is listed directly below the host. The compliance information about the DPU device appears on the right and is presented in an identical manner as the overall compliance information about the host. You can see full image comparison or drift comparison.

Image Compliance

Last checked on 08/04/2022, 3:43:02 AM (5 days ago)

⚠️ 1 of 1 hosts is out of compliance with the image

REMEDIATE ALL

RUN PRE-CHECK

STAGE ALL

Hosts

⚠️ 10.161.85.217

↪️ ⚠️ 28b7171f-2bc2-5433-84f9-f41daed0dd7b

2 hosts

28b7171f-2bc2-5433-84f9-f41daed0dd7b

⚠️ DPU is out of compliance with the image

⚠️ The CPU in this host may not be supported in future

ℹ️ The host will be rebooted during remediation.

Software compliance

Image	
Components	DPU
VMware-VIB-Test-Certs	8.0.0
VMware-Virtual-Dpu-Ctl	8.0.0

Check the Compliance of a Cluster Against a vSphere Lifecycle Manager Image

You check the cluster compliance against an image to understand how each of the hosts in the cluster compares to the specified image.

Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

When you initiate the check compliance operation on an object that contains multiple clusters that you manage with a single image, for example a data center or vCenter Server instance, vSphere Lifecycle Manager performs compliance checks on all those clusters.

1. In the vSphere Client, navigate to a cluster that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. In the **Image Compliance** card, click the **Check Compliance** button.

The **Image Compliance** card displays information about the overall number of non-compliant and incompatible hosts in the cluster. You can view detailed compliance information for every host. If the host has a DPU device, you can also view compliance information about the DPU device only.

Remediate the cluster to make the non-compliant hosts compliant. See [Run a Remediation Pre-Check for a Cluster, a Host Within a Cluster, or a Standalone Host](#) and [Remediate a Cluster Against a Single Image](#).

View Host Compliance Information

You can view detailed compliance information for every non-compliant host in a cluster or for every standalone host that you manage with a single image. You can also view compliance information about the DPU device on the host, if the host has a DPU device. As a result, you can easily find what causes the host to become out of compliance with the cluster or host image.

- Run a compliance check.
- Verify that no host is added to the cluster after your last compliance check.

Detailed compliance information is displayed only for hosts that are out of compliance with the image in the cluster or the standalone host. vSphere Lifecycle Manager displays no compliance details for compliant hosts.

For hosts that have the incompatible compliance state, vSphere Lifecycle Manager shows in a signpost information about what causes the compatibility issues.

1. In the vSphere Client, navigate to a cluster or host that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. View compliance information for the hosts in a cluster or for a standalone host.

Option	Action
View the compliance information for a host in the cluster.	1. In the Image Compliance card, select a host from the Hosts list. Note: Only non-compliant hosts are listed. 2. To view the full comparison between the image on the host and the image for the cluster, select Full image comparison from the drop-down menu for the Software compliance table. 3. To view only the image elements that make the host non-compliant with the image for the cluster, select Only drift comparison from the drop-down menu for the Software compliance table.
View the compliance information for a standalone host.	You can view the compliance information for a standalone host in the Image Compliance card. • To view the full comparison between the software that runs on the host and the image for the host, select Full image comparison from the drop-down menu for the Software compliance table.

Option	Action
	To view only the image elements that make the host non-compliant with the image, select Only drift comparison from the drop-down menu for the Software compliance table.

4. If you want to view compliance details about a DPU device on a host, select the DPU device listed for the host.

DPU devices are listed directly under the host they belong to.

- To view the full comparison between the image on the DPU device and the image for the cluster or the standalone host, select **Full image comparison** from the drop-down menu for the **Software compliance** table.
- To view only the image elements that make the DPU device non-compliant with the image for the cluster or the standalone host, select **Only drift comparison** from the drop-down menu for the **Software compliance** table.

An information panel appears on the right. In the **Software compliance** table, you can see what software runs on the selected host and what is the software specification in the image for the cluster or the standalone host.

Staging vSphere Lifecycle Manager Images to ESXi Hosts in a Cluster or Standalone Hosts

Staging is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. Staging reduces the time that ESXi hosts spend in maintenance mode.

In vSphere 8.0, staging exists as a separate operation that you can initiate. Staging is also automatically triggered when you remediate a cluster, a host in the cluster, or a standalone host. By default, vSphere Lifecycle Manager performs the staging operation onto all hosts in the cluster in parallel.

If you manage a cluster with a single image, you can stage the image to all hosts in the cluster or to a single host from that cluster. Staging is also supported for clusters with hosts and standalone hosts that have DPU devices. During staging, vSphere Lifecycle Manager downloads to the hosts all components that are applicable to the DPU device.

After staging is successful, a staged icon (📁) appears for any host to which the image is staged. The icon shows that both the software and firmware components from the image are successfully staged to the respective host. If for any reason vSphere Lifecycle Manager can't stage the firmware components to the host, staging does not fail, but you don't see the staged icon for that host.

Staging as Part of Remediation

Starting with vSphere 8.0, the remediation process happens in two steps, staging and then the actual remediation. vSphere Lifecycle Manager puts the ESXi hosts in maintenance mode only after staging is complete. The result is reduced ESXi downtime and disruption in the virtual workloads running on the virtual machines that reside on the respective hosts.

If for some reason staging is successful but the subsequent remediation process fails, the compliance status of the cluster or the standalone host changes from non-compliant to staged.

What Happens With the Staged Content?

The staged content is downloaded to scratch partition if scratch partition has enough space for the staged image. As a result, the staged content is persistent across power-cycles and reboots. If you delete the staged content, vSphere Lifecycle Manager re-stages the components needed for applying the image to the hosts in the cluster or the standalone host as part of the remediation process.

During staging, vSphere Lifecycle Manager validates the checksum and signature of the VIBs and downloads them to the appointed folder only after the validation is successfully completed. When the VIBs become available locally on the hosts, only a root user can modify them.

Staging Firmware Components

If an image contains a drivers and firmware add-on, during staging, the firmware VIBs are staged to the host or hosts along with the other software components defined in the image. If the firmware vendor does not support firmware staging, you get a notification in the **Image Compliance** card. You won't see the staged icon for that host even if staging the rest of the software is successful. In such a case, vSphere Lifecycle Manager installs the firmware components on the hosts during remediation.

Stage an Image to a Cluster

Staging an image to a cluster before remediating the cluster reduces the time that ESXi hosts spend in maintenance mode. After successful staging, you can immediately proceed with remediation, or you can remediate the cluster at a later time.

- Verify that the hosts in the cluster are non-compliant with the image.
- Verify that the hosts in the cluster are connected to vCenter Server.
- Verify that you have the required privileges to perform the task.

When you stage an image to a cluster, vSphere Lifecycle Manager downloads all software and firmware components defined in the image from the vSphere Lifecycle Manager depot to all hosts in the cluster. Before you can stage an image to a cluster or a host, you must run a compliance check for the cluster and ensure that the host or hosts to which you stage the image are non-compliant with that image. Staging content to compliant hosts doesn't alter the hosts in any way. Staging the image on hosts whose compliance status is unknown triggers an automatic compliance check.

1. In the vSphere Client, navigate to a cluster that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. In the **Image Compliance** card, click the **Stage All** button.
4. Optional: Click **Skip Remaining Hosts** if you want to stop the staging process while it is ongoing.
The image is not staged to the skipped hosts.

After staging finishes successfully, an icon (📁) appears for each host for the cluster. The icon shows that all software and firmware from the image are successfully staged to the hosts in the cluster. If for any reason vSphere Lifecycle Manager can't stage the firmware components to the host, staging is still successful, but you don't see the staged icon for the hosts.

Remediate the cluster to install the staged software and firmware to the hosts in the cluster.

Stage an Image to a Single Host Within a Cluster or a Standalone Host

Staging is an operation that you can perform for a single host within the cluster or for a standalone host. Learn how to stage an image to a single host from a cluster or to a standalone host.

- Verify that the hosts in the cluster or the standalone host are non-compliant with the image.
- Verify that the hosts in the cluster or the standalone host are connected to vCenter Server.
- Verify that you have the required privileges to perform the task.

1. In the vSphere Client, navigate to a cluster or host that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. Stage a single host.

Option	Action
Stage a single host in a cluster.	1. In the Image Compliance card, select a host from the Hosts list. A card with information about the host appears on the right.

Option	Action
	2. Click Actions > Stage > .
Stage a standalone host.	1. In the Image Compliance card, click the Stage button.

After staging finishes successfully, an icon () appears for the host. The icon shows that all software and firmware components from the image are successfully staged to the host. If for any reason vSphere Lifecycle Manager can't stage the firmware components to the host, staging does not fail, but you don't see the staged icon for the host.

Remediate the host to install the staged software and firmware components to the host.

Run a Remediation Pre-Check for a Cluster, a Host Within a Cluster, or a Standalone Host

To ensure that the health of a cluster or standalone host is optimal and that no problems occur during remediation, you can perform a remediation pre-check.

Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

The remediation pre-check operation includes a series of checks for the cluster and for each host in the cluster, and for the standalone host. These checks include extensive health checks to determine whether the cluster or host is in a stable state and to ensure successful remediation. Also, the remediation pre-check triggers a compliance check for the cluster or for the standalone host. As a result, after the remediation pre-check, you can view compliance information for each host and whether host reboot or maintenance mode are necessary for successful remediation.

For vSAN clusters, the remediation pre-check operation includes a hardware compatibility check. Depending on how you configure the vSphere Lifecycle Manager remediation settings, vSphere Lifecycle Manager might prevent remediation if hardware compatibility issues exist. For information about configuring the global vSphere Lifecycle Manager remediation settings, see [Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image](#) . For information about configuring the remediation settings for a particular cluster, see [Override the Global vSphere Lifecycle Manager Remediation Settings for a Cluster or Host That You Manage with a Single Image](#).

1. In the vSphere Client, navigate to a cluster or host that you manage with a single image.
2. On the **Updates** tab, select **Hosts > Image**.
3. In the **Image Compliance** card, click the **Run pre-check** button.

You can observe the progress of the pre-check task. When the task finishes, vSphere Lifecycle Manager displays information about the issues found during the pre-check.

If vSphere Lifecycle Manager reports no issues, remediate the cluster or the standalone host. See [Remediate a Cluster Against a Single Image](#) or [Remediate a Single Host Within a Cluster or a Standalone Host Against an Image](#).

If issues are reported, resolve the issues before you remediate the cluster or the standalone host.

Run a Remediation Pre-Check for a Single Host Within a Cluster or a Standalone Host

Instead of generating a remediation pre-check report for the entire cluster, you can run a remediation pre-check task for a single host in the cluster. You can also run a remediation pre-check for a standalone host in your inventory.

Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

The remediation pre-check task ensures that the host can be successfully remediated.

1. In the vSphere Client, navigate to a cluster or a standalone host that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. Run a remediation pre-check for the standalone host or a single host within the cluster.

Option	Action
Run a pre-check for a single host within the cluster.	1. In the Image Compliance card, select a host from the list of hosts. An information panel appears on the right. 2. Select Actions > Run pre-check. The Running pre-check card appears above the list of hosts. You can observe the progress of the task.
Run a pre-check for a standalone host.	1. In the Image Compliance card, click the Run pre-check button. The Running pre-check dialog appears inside the Image Compliance card. You can observe the progress of the task.

When the pre-check finishes, vSphere Lifecycle Manager displays information about the issues found during the pre-check.

If vSphere Lifecycle Manager reports no issues, you can remediate the host. See [Remediate a Single Host Within a Cluster or a Standalone Host Against an Image](#).

If vSphere Lifecycle Manager reports any issues, resolve the issues before you remediate the host.

Remediating a Cluster or a Standalone Host Against a Single Image

When you set up or import a vSphere Lifecycle Manager image to use with a cluster or host, the software specified in the image is not immediately installed on the hosts in that cluster or on the standalone host. To apply the software specification from the image to the hosts, you must remediate the cluster or the standalone host against that image.

To initiate remediation of a cluster or host, you must have the required privileges. For a list of all vSphere Lifecycle Manager privileges and their descriptions, see [vSphere Lifecycle Manager Privileges For Using Images](#). For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation.

During remediation, the image that you set up for the cluster is installed on all ESXi hosts in the cluster.

When you remediate a cluster that contains a single ESXi host or that has vSphere DRS deactivated or in manual mode, the remediation process cannot put that host into maintenance mode. So, to proceed with the remediation, you must power off the virtual machines that are running on the host, move them to another host, or select a user policy that allows the remediation process to power off the virtual machines. You can also set a user policy to power on the virtual machines after the host is remediated.

For vSAN clusters, the remediation operation includes a hardware compatibility check. Depending on how you configure the vSphere Lifecycle Manager remediation settings, vSphere Lifecycle Manager might not proceed with the remediation task if hardware compatibility issues exist. For information about configuring the global vSphere Lifecycle Manager remediation settings, see [Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image](#). For information about configuring the remediation settings for a particular cluster or standalone cluster, see [Override the Global vSphere Lifecycle Manager Remediation Settings for a Cluster or Host That You Manage with a Single Image](#).

Maintenance Mode

If the update requires it, hosts are put into maintenance mode before remediation. Virtual machines cannot run when a host is in maintenance mode. To ensure a consistent user experience, vCenter Server migrates the virtual machines to other hosts within the cluster before a host is put into maintenance mode. vCenter Server can migrate the virtual machines if the cluster is configured for vMotion and if DRS and VMware Enhanced vMotion Compatibility (EVC) are enabled. EVC guarantees that the CPUs of the hosts are compatible, but it is not a prerequisite for vMotion.

You can configure vSphere Lifecycle Manager to deactivate HA admission control for the cluster before remediation. However, deactivating HA admission control before you remediate a two-node cluster that uses a single vSphere Lifecycle Manager image causes the cluster to practically lose all its high availability guarantees. The reason is that when one of the two hosts enters maintenance mode, vCenter Server cannot failover virtual machines to that host and HA failovers are never successful. For more information about HA admission control, see the *vSphere Availability* documentation.

Parallel Remediation

During the remediation of a cluster against a vSphere Lifecycle Manager image, the ESXi hosts in the cluster are remediated sequentially by default. So, if the remediation for a single host in the cluster fails, the remediation of the entire cluster stops. However, you can configure vSphere Lifecycle Manager to remediate in parallel the hosts within a cluster that uses images. Parallel remediation reduces the overall remediation time and optimizes the maintenance window for the cluster. You can remediate in parallel only ESXi hosts that are already in maintenance mode. During parallel remediation, hosts do not enter maintenance mode automatically. Similarly, after remediation finishes, the hosts do not exit maintenance mode automatically. To remediate hosts in parallel, you must manually enter and exit maintenance mode. If you activate parallel remediation, vSphere Lifecycle Manager does not remediate the ESXi hosts that are not in maintenance mode.

When you configure vSphere Lifecycle Manager to remediate hosts in parallel, you can set the maximum number of hosts to be remediated in a single remediation task. Alternatively, you can let vSphere Lifecycle Manager calculate the optimal number of hosts to remediate in parallel.

When you remediate hosts in parallel, if the remediation of a single host fails, the remediation task for the entire cluster does not stop and the rest of the hosts are remediated successfully. After remediation finishes, vSphere Lifecycle Manager reports an error for the respective host.

Parallel remediation is deactivated by default, but you can activate it during remediation or in the vSphere Lifecycle Manager general remediation settings. Parallel remediation and all other remediation settings apply to remediation tasks that you start in vCenter Server. Solutions such as NSX, for example, might have separate parallel remediation settings.

You cannot remediate in parallel the witness host and the hosts in the associated vSAN cluster. Parallel remediation is also not possible for clusters that are enabled for vSphere Supervisor or NSX.

Remediating a DPU-Backed Cluster or Standalone Host

During remediation, vSphere Lifecycle Manager applies VIBs to the ESXi version on the DPU, if the host has a DPU device and the image for the cluster or host contains components with VIBs applicable to the DPU.

Remediate a Cluster Against a Single Image

By remediating a cluster against an image, you apply the software specified in the image to all the hosts in the cluster. So, by remediating a cluster, you make the non-compliant hosts compliant with the image that you set for the cluster.

Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

During remediation, the hosts in the cluster are remediated in sequence by default. You can configure vSphere Lifecycle Manager to remediate hosts in parallel.

The hosts that have the incompatible compliance state are not remediated.

If a vCenter HA failover is initiated during the remediation of a cluster, the remediation task is canceled. After the failover finishes, you must restart the remediation task on the new node.

1. In the vSphere Client, navigate to a cluster that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. In the **Image Compliance** card, click the **Remediate All** button.
The **Review Remediation Impact** dialog box appears. The dialog box contains detailed information about all changes that remediation will enforce on the hosts in the cluster.
4. In the **Review Remediation Impact** dialog box, review the impact summary, the applicable remediation settings, and the EULA.
5. To save and review the impact details later, click **Export Impact Details**.
6. Accept the EULA by selecting respective check box.
The check box is selected by default.
7. Click the **Start remediation** button.
The **Remediate Cluster** task appears in the **Recent Tasks** pane. You can also observe the progress of the remediation task in the **Image Compliance** card. If remediation fails, vSphere Lifecycle Manager gives information about the reasons for the failure.

Remediate a Single Host Within a Cluster or a Standalone Host Against an Image

When you remediate a single host within a cluster or a standalone host against the image for the cluster or host, vSphere Lifecycle Manager applies the image to that host only. Remediation is the operation that makes a non-compliant host in the cluster or a standalone host compliant with the image that you use for that cluster or that host.

- Verify that the host is non-compliant with the image for the cluster or the standalone host.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

1. In the vSphere Client, navigate to a cluster or host that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. Remediate a host within the cluster or a standalone host.

Option	Action
Remediate a single host within a cluster.	1. In the Image Compliance card, select a host from the Hosts list. A card with information about the host appears on the right. 2. Click Actions > Remediate > . The Review Remediation Impact dialog box appears. The dialog box contains detailed information about all changes that remediation will enforce on the host.
Remediate a standalone host.	1. In the Image Compliance card, click the Remediate button. The Review Remediation Impact dialog box appears. The dialog box contains detailed information about all changes that remediation will enforce on the host.

4. Review the impact summary, the applicable remediation settings, and the EULA.
5. If you remediate a single host in a cluster, to save and review the impact details later, click **Export Impact Details**.
6. Accept the EULA by selecting respective check box.
The check box is selected by default.

7. Click the **Start remediation** button.

The **Remediate Cluster** or the **Remediate Host** tasks appear in the **Recent Tasks** pane. You can also observe the progress of the remediation task in the **Image Compliance** card. If remediation fails, vSphere Lifecycle Manager gives information about the reasons for the failure.

View Last Remediation or Remediation Pre-Check Results for a Cluster or a Standalone Host that Uses a Single Image

You can view the remediation results from the last remediation or remediation pre-check that vSphere Lifecycle Manager performed.

Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

1. In the vSphere Client, navigate to a cluster or a standalone host that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. In the **Image Compliance** card, click the horizontal ellipsis icon and select your task.
 - To view the results from the last remediation pre-check performed on the cluster or the standalone host, select **Last pre-check results**.
 - To view the results from the last remediation of the cluster or the standalone host, select **Last remediation results**.

The **Image Compliance** displays detailed information about the last remediation or remediation pre-check task that ran on the cluster or the standalone host.

Manage Depot Overrides for a Cluster or a Standalone Host

Instead of accessing the vSphere Lifecycle Manager depot in vCenter Server, clusters or hosts in Remote Office and Branch Office (ROBO) deployments can download data from a depot that is local for them. You can configure vSphere Lifecycle Manager to use local depots for any cluster or standalone host that uses images.

- Set up an online depot to which the cluster or host can connect.
- Export an offline bundle with components from a vSphere Lifecycle Manager image and import the offline bundle to the target local depot.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

A ROBO cluster or a ROBO standalone host has limited or no access to the Internet or limited connectivity to vCenter Server. As a result, clusters and standalone hosts in ROBO deployments might have limited access to the vSphere Lifecycle Manager depot during the compliance check, remediation pre-check, and remediation operations.

With vSphere Lifecycle Manager images, you can use a local depot for ROBO clusters and hosts, and configure vSphere Lifecycle Manager to use the local depot during the compliance check, remediation pre-check, and the remediation tasks. The local depot overrides the vSphere Lifecycle Manager depot. Using local depots with ROBO clusters and ROBO standalone hosts saves time and network bandwidth.

Starting with vSphere 8.0 Update 1, you can manage local depot overrides for ROBO standalone hosts in the vSphere Client and by using the vSphere Automation API.

For each cluster or standalone host that you manage with a single image, you can add and use multiple local depots instead of the default vSphere Lifecycle Manager depot. You can also delete the depot overrides that you configure. If

depot overrides are not active for a cluster or a standalone host, the cluster or the host uses the general vSphere Lifecycle Manager depot in vCenter Server.

1. In the vSphere Client, navigate to a cluster or a standalone host that you manage with a single image.
2. On the **Updates** tab, select **Hosts > Image**.
3. In the **Image Compliance** card, click the horizontal ellipsis icon and select **Manage depot overrides**.

The **Manage depot overrides** dialog box appears.

4. Select your task.

Option	Description
Add depot overrides	Enter a URL or a file path to a local depot and click Add . The depot is added to the Depot override URL list.
Delete a depot override	Click the horizontal ellipsis icon for a depot override from the list and click Delete . The depot is removed from the Depot override URL list.

5. Click **Close**.

In the **Image Compliance** pane, you see a notification if depot overrides are active for the cluster or the standalone host.

vSphere Lifecycle Manager Recommended Images

For the clusters that you manage with images, vSphere Lifecycle Manager can generate and provide you with software recommendations in the form of pre-validated images that are compatible with the hardware of the hosts in a cluster. Recommended images are valid images that are based on the latest ESXi version in a series of releases. You can easily replace the current image of a cluster with the recommended image for that cluster.

When you set up or edit an image, you manually combine the image elements (ESXi version, vendor add-on, firmware add-on, and additional components) in such a way as to define the full software stack to run on all hosts in the cluster. You must manually check whether a particular image set-up is complete and valid, and suitable to your environment. The vSphere Lifecycle Manager recommendations save you the effort of exploring the possible and applicable combinations of image elements vSphere Lifecycle Manager generates recommendations automatically under certain conditions and upon certain events. Alternatively, you can trigger recommendation generation manually.

Recommended images are validated through a series of checks that ensure that a recommended image has no missing dependencies or conflicting components. For vSAN clusters, the validation also runs a hardware compatibility check against the vSAN Hardware Compatibility List (vSAN HCL). The extensive validation checks ensure that if you decide to use a recommended image for a cluster, the remediation against the recommended image is successful.

To generate recommendations, vSphere Lifecycle Manager checks what software is available in the vSphere Lifecycle Manager depot and what firmware is available in the depot that the selected hardware support manager makes available. Based on the available software, firmware, and, for vSAN clusters, the hardware compatibility checks, for each cluster that you manage with a single image, vSphere Lifecycle Manager provides you with the latest and most appropriate image for your environment.

You can manually cancel the recommendation generation task from the **Recent Tasks** pane.

What Does Latest in Current Series Mean

The recommendation that vSphere Lifecycle Manager generates is based on the major ESXi version in the image that the cluster uses. vSphere Lifecycle Manager recommends the most recent ESXi version that causes no hardware compatibility issues or regressions. The ESXi version in the recommended image is always within the major release series of the ESXi version of the image for the cluster.

The latest image in a series contains the latest minor ESXi version or patch for that version. For example, if the current image for a cluster contains a base ESXi image of version 7.0 and base images of version 7.0 Update 1, 7.0 Update 2, 7.0 Update 3, and 8.0 are available in the vSphere Lifecycle Manager depot, the latest image in the current series recommendation contains ESXi version 7.0 Update 3.

In some cases, the ESXi version in a recommended image might be the same as the ESXi version in the current image for a cluster, but the recommended image might contain a later version of any other image element, for example the vendor add-on, a component, or the firmware add-on.

Why Are Some ESXi Versions not Recommended

The latest ESXi major version available in the depot might not be recommended, because it might be in conflict with upgrade path restrictions or because it causes hardware compatibility issues. Images that are not recommended contain ESXi versions that, if used for a cluster, might lead to regressions.

For example, the current image for a cluster contains a base ESXi image of version 7.0 Update 1 and the vSphere Lifecycle Manager depot contains ESXi images of version 7.0 Update 2, a hot patch for 7.0 Update 2, and 8.0. Upgrading to the ESXi 7.0 Update 2 with the hot patch is an allowed upgrade path. Upgrading further to ESXi 8.0 causes regression, because the recommended image that contains ESXi 8.0 does not contain the hot patch. So, vSphere Lifecycle Manager lists ESXi Update 2 with the hot patch as the latest in current series recommendation. ESXi 8.0 is listed as a not recommended ESXi version.

In some cases, vSphere Lifecycle Manager might list as not recommended an ESXi version within the current release series. The reason is possible hardware compatibility issues.

In the vSphere Client, you can view a list of all ESXi versions that are not recommended for use together with an explanation about the reason why they are not recommended.

Automatically Triggered Recommendation Generation

vSphere Lifecycle Manager generates a new image recommendation automatically in the following cases.

- The vSphere Lifecycle Manager depot gets updated.
By default, the depot updates every 24 hours. Also, the content of the depot changes when you import an offline bundle to the depot or you manually trigger synchronization to the configurable download sources.

Note:

If the depot gets updated with solution components only, vSphere Lifecycle Manager does not generate new recommendations. Similarly, if the depot gets updated by adding or removing base images that are of a major ESXi version as compared to the current image of a cluster, vSphere Lifecycle Manager does not start a recommendation generation task.

- You edit the image that you use for a cluster and save the new image set-up.

Note:

If you edit an image by only adding solution components to the image, vSphere Lifecycle Manager generates no new recommendation.

The automated recommendation generation is available only for clusters that already have recommended images generated. When vSphere Lifecycle Manager starts the generation of a new recommendation automatically, the Compute Image recommendations for cluster task appears in the **Recent tasks** pane. You can observe the progress of the task or cancel it. vCenter Server issues an event when a recommendation generation task starts or ends. If the task fails, vCenter Server issues an alarm of the warning type. In cases of failure, you must check for recommended images for the cluster manually. The recommendation generation task cannot run simultaneously with other vSphere Lifecycle Manager operations, for example remediation and compliance checks. If you need to start another operation immediately, you can cancel the Compute Image recommendations for cluster task at any time.

In ROBO deployments, the automatically triggered recommendation generation is possible only if the local depot and the central vSphere Lifecycle Manager depot are synchronized.

Check for Recommended Images

By using recommended images for your clusters or standalone hosts, you ensure that your environment runs the latest verified software. Although vSphere Lifecycle Manager generates image recommendation automatically when the cluster or the host changes or when new software is available in the vSphere Lifecycle Manager depot, you can manually check for recommendations at any time.

- Verify that remediation is not running for the cluster or the standalone host.
- Verify that you have connection to the Internet.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

A recommended image contains updates for your cluster or host. The recommendations is based on the ESXi versions available in the vSphere Lifecycle Manager depot. When you trigger the **Compute image recommendations** task, vSphere Lifecycle Manager first determines the recommended ESXi version for the cluster or host. After that, vSphere Lifecycle Manager checks sequentially for newer versions of the vendor add-on, additional components, and firmware add-on that are compatible with the recommended ESXi version and the hardware of the hosts in the cluster or the standalone host. So, sometimes, a recommended image might contain the same ESXi version as the ESXi version in the current image for the cluster or host but combined with an updated vendor add-on, component, or firmware add-on.

The **Compute image recommendations** task is non-cancellable. You must rerun the task periodically to ensure that the recommendations are valid and still suitable to the cluster or host.

Before you check for recommended images, you must ensure that the cluster or host is not being remediated. Recommendations generation and remediation are mutually exclusive operations. They cannot run simultaneously.

1. In the vSphere Client, navigate to a cluster or a standalone host that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Image**.
3. In the **Image** card, click the horizontal ellipsis icon and select **Check for recommended images**.
When the task finishes, a blue badge appears in the **Image** card.
4. To view the recommended images, click the horizontal ellipsis icon and select **View recommended images**.

vSphere Lifecycle Manager suggests a recommended image for the cluster or for the standalone host or reports that no recommended images are available for the cluster or the host. When no recommended images are available, vSphere Lifecycle Manager displays detailed information about why no recommendations are available or why certain ESXi versions are not recommended.

View the recommendations. You can import a recommended image to the cluster or host and replace the current image that the cluster or the host uses. See [Use the vSphere Lifecycle Manager Recommended Image as the New Image for a Cluster or a Standalone Host](#).

View the Recommended or Not Recommended Images for a Cluster or a Standalone Host

You can view the pre-validated image that vSphere Lifecycle Manager recommends for a cluster or a standalone host and also the ESXi versions that vSphere Lifecycle Manager does not recommend for that cluster or host because of possible hardware compatibility issues or regression.

- Verify that the **Compute image recommendations** task is performed for the cluster or the host or check for recommended images for the cluster or the host manually. See [Check for Recommended Images](#).

- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).
1. In the vSphere Client, navigate to a cluster or a standalone host that you manage with a single image.
 2. On the **Updates** tab, select **Hosts** > **Image**.
 3. If no recommendations are available for the cluster or host at all, click **View Details** in the banner at the top of the pane to see details.
 4. If recommendations are available, in the **Image** card, click the horizontal ellipsis icon and select **View recommended images**.
 - To view the image that vSphere Lifecycle Manager recommends for the cluster or the host, check the **Latest in current series** column in the **Recommended Images** dialog box.
 - To view the ESXi versions that vSphere Lifecycle Manager doesn't recommend for the cluster or the host, in the **Recommended Images** dialog box, click **Show versions that are not recommended**.

Note: The **Show versions that are not recommended** link is visible only if the vSphere Lifecycle Manager depot contains more recent ESXi versions than the ESXi version in the current cluster or host image, but those more recent versions are not within the major release series of the current base image or might lead to hardware compatibility issues. For example, if the image for your cluster or host contains ESXi version 7.0, and ESXi 8.0 is available in the vSphere Lifecycle Manager depot, the latter is listed as a not recommended version.

If the **View recommended images** option is dimmed, no recommended images are available for this cluster or host. If recommendations are available, in the **Image** card, you see a label indicating the number of recommendations.

The **Recommended Images** dialog box appears.

You can decide to use the recommended image and set it as the new image for the cluster or the host.

Use the vSphere Lifecycle Manager Recommended Image as the New Image for a Cluster or a Standalone Host

For any cluster or host that you manage with a single image, you can view the images that vSphere Lifecycle Manager recommends and you can replace the current image for the cluster or the host with a recommended image. Using recommended images saves you the time and effort of identifying valid images that are applicable to all hosts in a cluster or to a standalone host.

- Check for recommended images for a cluster or a host. See [Check for Recommended Images](#).
 - Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).
1. In the vSphere Client, navigate to a cluster or to a standalone host that you manage with a single image.
 2. On the **Updates** tab, select **Hosts** > **Image**.
 3. In the **Image** card, click the horizontal ellipsis icon and select **View recommended images**.
If the **View recommended images** option is dimmed, no recommended images are available for this cluster or host.
The **Recommended Images** dialog box appears.
 4. In the **Recommended Images** dialog box, select a recommended image by clicking the respective radio button and click **Continue**.

The recommended image is based on the latest ESXi version within the current release series. For example, if the ESXi version in your current image is 7.0, this option might include ESXi version 7.0 Update 3 and a related vendor add-on. The recommended image cannot contain ESXi 8.0 or any other minor release version within the 8.0 release series.

The selected image is imported to the cluster or the standalone host as a draft. The **Edit Image** card appears.

5. Optional: Edit the image and validate the new image set-up.

6. Click **Save**.

If you do not save the image, it is saved as a draft. The next time you start editing the image for that cluster or host, you can use the draft as a starting point.

The recommended image is saved for that cluster or host. If a draft exists for the cluster or the host, the draft is overridden by the recommended image. No software is installed on the hosts in the cluster or on the standalone host at this stage.

To apply the software specification that the image defines, remediate the cluster or the host against the new image. See [Run a Remediation Pre-Check for a Cluster, a Host Within a Cluster, or a Standalone Host](#) and [Remediating a Cluster or a Standalone Host Against a Single Image](#).

Firmware Updates with vSphere Lifecycle Manager

You can use vSphere Lifecycle Manager images to perform firmware updates on the ESXi hosts in a cluster or on the standalone hosts. With a single operation, you update both the software and the firmware on the host.

In earlier vSphere releases, you could perform firmware updates on vSAN clusters by using system-managed baselines. For non- vSAN clusters, firmware updates had to be manual.

vSphere Lifecycle Manager enables you to easily update the firmware in any cluster or host that you manage with a single image. Firmware updates are not available for clusters or hosts that you manage with baselines.

To apply firmware updates to the hosts in a cluster or to the standalone host that you manage with a single image, you must include a special type of add-on, the firmware and drivers add-on, in the image and remediate the cluster or the host to apply the image to the hosts. The firmware and drivers add-on is a vendor-provided add-on that contains the components that encapsulate firmware update packages. The firmware and drivers add-on might also contain the necessary drivers.

Unlike vendor add-ons, firmware and drivers add-ons are not distributed through the official VMware online depot or as offline bundles available at <https://my.vmware.com>. For a given hardware vendor, firmware updates are available in a special vendor depot, whose content you access through a software module called a hardware support manager. The hardware support manager is a plug-in that registers itself as a vCenter Server extension. Each hardware vendor provides and manages a separate hardware support manager that integrates with vSphere. For each cluster or host that you manage with a single image, you select the hardware support manager that provides the firmware updates for the cluster or the host. After you determine the hardware support manager that you want to use for a cluster or a standalone host, the hardware support manager provides you with a list of the available firmware updates. When you select and include a firmware add-on to an image, that add-on might modify the specified image by adding or removing components. The firmware add-on also defines the firmware versions to be installed on the hosts. During remediation, vSphere Lifecycle Manager applies the image to the hosts and requests the selected hardware support manager to update the firmware on the hosts in accordance with the firmware add-on specified in the image.

Selecting a hardware support manager and including a firmware add-on in your image guarantees that during a compliance check, vSphere Lifecycle Manager also determines the firmware compliance for the cluster or the host. So, you can easily detect and remediate any unwanted drifts. The hardware support manager is also responsible for retrieving the firmware versions on the host hardware, and, in some cases, determining the appropriate drivers for the updated firmware version.

For vSAN clusters, the hardware support manager inspects the hosts in the cluster to determine their current I/O device controllers and firmware. During a hardware compatibility check for the cluster, vSphere Lifecycle Manager checks whether the firmware in the image is compatible with the hardware in the cluster as per vSAN Hardware Compatibility List (vSAN HCL). The hardware compatibility check ensures that when vSphere Lifecycle Manager remediates the cluster and applies the image to all hosts, the firmware and drivers on the hosts are certified for use with vSAN.

Firmware Updates on DPU Devices

In a DPU-based environment, you can update the firmware on the DPU devices of the hosts through the vSphere Lifecycle Manager remediation operation only if the server vendor has an integrated hardware support manager. For servers that don't provide a hardware support manager, firmware updates are manual.

Deploying Hardware Support Managers

The deployment method and the management of a hardware support manager plug-in are determined by the respective OEM.

Several of the major OEMs develop and supply hardware support managers. For example:

- **Dell**
The hardware support manager that Dell provides is part of their host management solution, OpenManage Integration for VMware vCenter (OMIVV), which you deploy as an appliance.
- **HPE**
The hardware support managers that HPE provides are part of their management tools, iLO Amplifier and OneView, which you deploy as appliances.
- **Lenovo**
The hardware support manager that Lenovo provides is part of their server management solution, Lenovo XClarity Integrator for VMware vCenter, which you deploy as an appliance.
- **Hitachi**
The hardware support manager that Hitachi provides, Hitachi Unified Compute Platform Advisor, is infrastructure automation and management software for all Hitachi converged, hyperconverged, and integrated systems, which you deploy as an appliance.
- **Cisco**
The hardware support manager that Cisco provides is integrated with Cisco Intersight Infrastructure Service, which is part of Cisco Intersight and you activate the hardware support manager from within the Cisco Intersight SaaS-based management platform. No additional appliances are required on your vCenter Server instance.

You can find the full list of all VMware-certified hardware support managers in the VMware Compatibility Guide at <https://www.vmware.com/resources/compatibility/search.php?deviceCategory=hsm>.

Deploying and Configuring Hardware Support Managers

Regardless of the hardware vendor, you must deploy the hardware support manager appliance on a host with sufficient memory, storage, and processing resources. Typically, hardware support manager appliances are distributed as OVF or OVA templates. You can deploy them on any host in any vCenter Server instance.

Note: If vCenter Server is configured with a proxy for internet access, the proxy must be able to reach any hardware support manager registered with that vCenter Server instance. You must either assign the hardware support manager a private IP address within the 10.x.x.x range, which is automatically exempt from proxy use, or enable a direct access to the registered hardware support manager by configuring the proxy settings with an exception for its IP address.

After you deploy the appliance, you must power on the appliance virtual machine and register the appliance as a vCenter Server extension. You might need to log in to the appliance as an administrator. Each hardware support manager might register with only one or multiple vCenter Server systems.

A vCenter Server plug-in user interface might become available in the vSphere Client after you deploy a hardware support manager appliance, but the hardware support manager might also have a separate user interface of its own. For example, OMIVV, iLO Amplifier, and Lenovo XClarity Integrator for VMware vCenter all have a vCenter Server plug-in user interface, which helps you configure and work with the respective hardware support manager.

Each hardware support manager has its own mechanism of managing the actual firmware packages and making firmware add-ons available for you to choose.

The successful integration between the hardware support manager and vSphere Lifecycle Manager might require a specific configuration of the hardware support manager. For example, with OMIVV, you must first create a connection profile. Then, you must create a cluster profile and associate it with a cluster before you can add a firmware add-on from Dell to the image for that cluster.

With vSphere 8.0 Update 3, vSphere Lifecycle Manager and vSAN integrate to support proactive hardware management. You can deploy and register a hardware support manager that has proactive hardware support management functionality. All hosts in a vSAN cluster associated with such hardware support manager, start receiving proactive notifications about predictive disk failures. You can use these notifications to take early remediation steps before the actual disk failure occurs. Each vendor has a predefined set of health statuses that trigger the respective vCenter Server event. Upon receiving an event regarding disk failure, vSAN raises a health alarm and provides some troubleshooting guidance from the vendor. For more information about proactive hardware management, see the *Managing Proactive Hardware* section in the *vSAN Monitoring and Troubleshooting* documentation. For more information about how to replace failed disks, see the *Replace a Flash Caching Device on a Host in vSAN Cluster* and *Replace a Capacity Device in vSAN Cluster* procedures in the *vSAN Monitoring and Troubleshooting* documentation.

For detailed information about deploying, configuring, and managing hardware support managers, refer to the respective OEM-provided documentation.

Use an Image for Firmware Updates

vSphere Lifecycle Manager allows you to manage the firmware lifecycle on ESXi hosts that are part of a cluster or on a standalone host that you manage with a single image.

- Deploy the vendor-provided hardware support manager and register it as a vCenter Server extension. For more information about deploying and managing a hardware support manager, see the respective OEM documentation.
- If you use the hardware support manager provided by Dell, create a cluster profile and associate it with the cluster. For more information, review the OpenManage Integration for VMware vCenter (OMIVV) documentation.
- Verify that all hosts in the cluster are from the same vendor.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

1. In the vSphere Client, navigate to a cluster or a standalone host that you manage with a single image.

2. On the **Updates** tab, select **Hosts** > **Image**.

3. In the **Image** card, click the **Edit** button.

4. In the **Edit Image** card, for the **Firmware and Drivers Addon**, click **Select**.

The **Firmware and Drivers Addon** dialog box appears.

5. In the **Firmware and Drivers Addon** dialog box, select a hardware support manager from the drop-down menu.

The selected hardware support manager must be from the same hardware vendor as the hosts in the cluster or as the standalone host. Otherwise, during a compliance check, the hardware support manager reports the selected firmware and drivers add-on to be incompatible with the host or hosts that are from a different vendor. Firmware remediation fails.

A list of all available firmware add-ons appears.

6. Select a firmware add-on from the list.

An information panel appears on the right. The panel contains information about the supported ESXi versions and whether the selected add-on contains the necessary drivers.

7. Click **Select**.

The selected firmware and drivers add-on is included in the image.

8. In the **Image** card, validate and save the image.

After the image is saved, a compliance check against the new image is triggered for the cluster or host.

9. In the **Image Compliance** card, review the compliance check results for the cluster and for the host.
10. If any host in the cluster or the standalone host has firmware that is non-compliant with the new image firmware, remediate the respective host or the cluster.
 - a) Optional: In the **Image Compliance** card, run a remediation pre-check to ensure that remediation finishes successfully.
 - To run a pre-check for all hosts in the cluster, click the **Run Pre-check** button.
 - To run a pre-check for a single host in the cluster, click the vertical ellipsis icon for the host and select **Run Pre-check**.
 - To run a pre-check for a standalone host, click the **Run Pre-check**.
 - b) In the **Image Compliance** card, initiate remediation.
 - To remediate all hosts in the cluster, click the **Remediate All** button.
During cluster remediation, if the remediation of a single host fails, the remediation for the cluster ends prematurely.
 - To remediate a single host, click the vertical ellipsis icon for the host and select **Remediate**.
 - To remediate a standalone host, click the **Remediate** button.

You are not obliged to start remediation immediately after setting up an image for a cluster or a standalone host. However, nothing is installed on the hosts unless you remediate them against the image for the cluster or the host. The firmware on the hosts is actually updated only after successful remediation. You can remediate the objects in your environment at any time that is convenient for you.

The firmware on the hosts in the cluster or the standalone host is updated to the firmware version specified in the firmware add-on for the image.

vSphere Lifecycle Manager Hardware Compatibility Checks for Clusters and Hosts

vSphere Lifecycle Manager automates the process of validating the hardware compliance of hosts and clusters against a selected ESXi version. Hardware compatibility checks ensure that the host or cluster hardware is compliant with the VMware Compatibility Guide (VCG) and vSAN Hardware Compatibility List (vSAN HCL).

Hardware Compatibility Lists

Hardware compatibility lists are lists of hardware certified for use with various vSphere releases. The VCG contains information about server models and I/O devices that are certified for use with particular vSphere releases. Besides VCG, vSAN maintains a separate hardware compatibility list that lists all I/O and networking device controller hardware and the respective firmware versions certified for use with vSAN. The vSAN HCL also contains information about the disk drives that a specific vSphere release supports and the earliest disk drive firmware version certified for use with vSAN.

With vSphere Lifecycle Manager, you can perform the following tasks.

- Check the hardware compatibility of a single host.
- Check the hardware compatibility of a vSAN cluster.

In general, hardware incompatibilities do not prevent remediation and are not resolved upon remediation. However, you can configure vSphere Lifecycle Manager to prevent remediation when hardware compatibility issues exist for a cluster. For information about configuring the global vSphere Lifecycle Manager remediation settings, see [Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image](#). For information about configuring the remediation settings for a particular cluster, see [Override the Global vSphere Lifecycle Manager Remediation Settings for a Cluster or Host That You Manage with a Single Image](#).

Cluster-Level Hardware Compatibility Checks

Running a hardware compatibility check on your clusters before remediating them helps you ensure good vSAN cluster health and avoid entering into unsupported and unwanted configurations after remediation.

Cluster-level hardware compatibility checks are available only for vSAN clusters that you manage with a single image. If a vSAN cluster uses baselines, hardware compatibility checks are not available. Also, if a cluster uses a single vSphere Lifecycle Manager image but vSAN is not enabled for that cluster, hardware compatibility checks for that cluster are not available. Hardware compatibility checks for vSAN clusters are performed against the vSAN HCL.

To perform a hardware compatibility check for a vSAN cluster, the vSAN HCL data that is available to vSphere Lifecycle Manager must be up to date. vSAN HCL data is synchronized automatically or, in environments without connection to the Internet, manually. For more information about maintaining the vSAN HCL data up-to-date, see the vSAN documentation.

When you initiate a hardware compatibility check for a cluster, vSphere Lifecycle Manager scans the image and verifies that all elements of the image are compatible with the hardware of all hosts within the cluster. vSphere Lifecycle Manager validates only those hardware devices that vSAN uses. Because cluster-level hardware compatibility checks validate the compatibility between the cluster hardware and the cluster image, the compatibility results might not be accurate unless you remediate the cluster successfully and apply the image to all hosts in the cluster.

Note: vSphere Lifecycle Manager performs the full driver and firmware verification of PCI devices and disk drives only if you configure vSphere Lifecycle Manager with a hardware support manager and add a firmware add-on to the vSphere Lifecycle Manager image. Without using a hardware support manager, vSphere Lifecycle Manager only validates the PCI and NIC devices and their respective driver versions, and the disk drive version.

Hardware compatibility issues are reported as warnings, and as such, they do not prevent you from remediating the hosts in the cluster against the image unless you change that behavior by editing the remediation settings for the cluster. For information about configuring remediation settings for a cluster, see [Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image](#) and [Override the Global vSphere Lifecycle Manager Remediation Settings for a Cluster or Host That You Manage with a Single Image](#).

Depending on the type of device, during a hardware compatibility check for a cluster, vSphere Lifecycle Manager performs different validations and verifications during a hardware compatibility check.

Device Type	vSphere Lifecycle Manager Validation
Storage I/O controllers	For storage I/O controllers, vSphere Lifecycle Manager performs the following verifications: - Verifies that all storage device controllers are certified for use with the ESXi version specified in the image. - Verifies that the image contains the correct storage device driver and firmware versions as per the vSAN HCL. - Suggests a compatible storage device driver and firmware version combination as per the vSAN HCL.
Network controllers	For NIC devices, vSphere Lifecycle Manager performs the following verifications: - Verifies that the physical RDMA-enabled NICs on the hosts in the cluster are certified for use with the ESXi version specified in the image as per the vSAN HCL. - Verifies that the image that the cluster uses contains the correct NIC driver and firmware versions as per the vSAN HCL. - Suggests a compatible NIC driver and firmware version combination for the cluster as per the vSAN HCL.

Device Type	vSphere Lifecycle Manager Validation
	Verifies that the device is certified for all active features that are enabled on that device. For example, if vSAN uses a RDMA-capable NIC, vSphere Lifecycle Manager verifies that the NIC is certified for RDMA.
Disk drives	For disk drives, vSphere Lifecycle Manager performs the following verifications: - Verifies that all disk drives in the cluster are certified for use with the ESXi version specified in the image as per vSAN HCL. - Verifies that the disk drive firmware version specified in the image for the cluster is equal to or higher than the earliest supported firmware version for the device as per vSAN HCL. Note: vSphere Lifecycle Manager performs those verifications not only on directly connected disk drives, but also on physical disk drives behind RAID-0 logical volumes and behind a VMD controller.

What Are the Currently Active Features

The list of active features for a device is a list of the software features that are enabled on the device and that the device can support. The active features that a device has enabled are the result of enabling or using a solution or a service. In the vSphere Client, information about the currently active features on a device is only available for PCI devices - storage controllers and network controllers. A list of the currently active features is not displayed for disk drives.

When a solution or service, for example vSAN or NSX, enables a set of features on a device, those features impose certain hardware compatibility constraints upon the device. The device must be certified for those active features that vSAN uses. During a hardware compatibility check, in addition to validating the driver and firmware combination for a PCI device, vSphere Lifecycle Manager also validates the constraints that the active features impose on that device. So, by running a hardware compatibility check, you ensure that only certified features are enabled on the devices in the vSAN cluster. You don't get compatibility issues for features that are not enabled on the device.

When the set of active features for a device changes, you must re-run the hardware compatibility check to get correct compliance results.

Hardware Compatibility Check Results

After a hardware compatibility check, vSphere Lifecycle Manager lists a device as compliant or non-compliant. In some cases, the compatibility status is unavailable.

Compliant

A device is compliant when it is compatible with the ESXi version and the driver and firmware version defined in the cluster for the image. For PCI devices, the driver-firmware version combination must also be certified for all active features enabled on the respective device.

Non-compliant

A device is non-compliant when it is incompatible with the ESXi, driver, or firmware version defined in the cluster for the image. Additionally, PCI devices are non-compliant when the driver-firmware version combination is not certified for all active features enabled on the respective device.

If a device is not present in the vSAN HCL at all, vSphere Lifecycle Manager marks the device as non-compliant.

Unavailable

The hardware compatibility status for a NIC device is unavailable in the following cases:

- Compatibility information for PCI devices is unavailable if the vSAN HCL data is not updated and does not contain information about the constraints imposed upon a device by the active features that are enabled on that device.
- The cluster has no hosts or some hosts in the cluster are inaccessible.
- No hardware support manager is registered in vCenter Server or the image for the cluster does not contain a firmware and drivers add-on.

When Does the Hardware Compatibility Check Task Run

vSphere Lifecycle Manager performs regular hardware compatibility checks for the vSAN clusters that you manage with a single image. In addition, certain vSphere Lifecycle Manager operations also trigger an automatic hardware compatibility check. Automated hardware compatibility checks are available for vSAN clusters that use a single image.

For information about hardware compatibility checks and instructions how to manually perform a hardware compatibility check for a cluster or for a single host, see [vSphere Lifecycle Manager Hardware Compatibility Checks for Clusters and Hosts](#).

The Hardware Compatibility Check Task

The vSAN Hardware Compatibility List (vSAN HCL) database changes regularly. For example, when VMware certifies new OEM devices, drivers, or firmware, those become part of the vSAN HCL database. Similarly, devices, drivers, or firmware that are not supported anymore are removed from the vSAN HCL database.

Changes in the vSAN HCL database might make your hardware compatibility results invalid and outdated. To provide you with valid hardware compatibility information, vSphere Lifecycle Manager runs a periodic hardware compatibility check against the latest vSAN HCL data.

The periodic hardware compatibility check is a preconfigured scheduled task that you can edit and force to run at any time. By default, the task runs every 24 hours. The scheduled task is configured at a vCenter Server level. If a vCenter Server system contains no vSAN clusters that you manage with a single image, vSphere Lifecycle Manager skips the scheduled hardware compatibility check. This periodic task runs only for vSAN clusters that you manage with a single image.

vSphere Lifecycle Manager Operations That Trigger a Hardware Compatibility Check

vSphere Lifecycle Manager performs an automatic hardware compatibility check for any vSAN cluster that you manage with a single image in the following cases.

- You edit the image for the cluster and save the image.
When you edit and save an image, vSphere Lifecycle Manager starts the Check hardware compatibility of cluster's host with image task even for clusters without vSAN. In such case, vSphere Lifecycle Manager only returns a warning that image hardware compatibility is not verified in non- vSAN clusters.
If the automatically triggered hardware compatibility task fails, you can still save the new image for the cluster.
- You initiate a remediation pre-check or remediation.
The hardware compatibility check is a part of the remediation pre-check and remediation task for vSAN clusters. If a cluster is not vSAN-enabled, vSphere Lifecycle Manager does not perform a hardware compatibility check when you initiate a remediation pre-check or remediation.
You can configure how vSphere Lifecycle Manager behaves in case of hardware compatibility issues by editing the remediation settings for the cluster. For information about configuring remediation settings for a cluster, see [Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single](#)

[Image](#) and [Override the Global vSphere Lifecycle Manager Remediation Settings for a Cluster or Host That You Manage with a Single Image](#).

- You add or remove a host to and from the cluster.

When you add or remove a host to and from the cluster, vSphere Lifecycle Manager invalidates the hardware compatibility check results for the cluster and issues a warning. You must rerun a hardware compatibility check to obtain valid information about potential hardware compatibility issues. Alternatively, you can remediate the cluster or run a remediation pre-check, which both automatically trigger a hardware compatibility check.

NIC validation

Starting with vSphere 8.0, hardware compatibility checks for vSAN clusters are enhanced to incorporate NIC validation against the vSAN HCL. Determining the compatibility between the NICs in a vSAN cluster and the software defined in the image that the cluster uses is crucial to a successful cluster upgrade.

During a hardware compatibility check, for each NIC device, vSphere Lifecycle Manager verifies that the driver and firmware version combination defined in the image for the cluster is certified for use with all active features enabled on the device.

For NIC devices, vSphere Lifecycle Manager checks the exact firmware version during the hardware compatibility check.

In a vSAN cluster, during a hardware compatibility check, vSphere Lifecycle Manager validates only the RDMA-capable NICs that vSAN uses. That is, if a host has a RDMA-capable NIC but the NIC is not in use, vSphere Lifecycle Manager does not calculate the hardware compatibility of that device. vSphere Lifecycle Manager does not validate non-RDMA NICs. Non-RDMA NICs in use by vSAN need no certification and therefore vSphere Lifecycle Manager does not validate them during a hardware compatibility check.

System Requirements for NIC Validation

- vCenter Server 8.0 and later
- ESXi 8.0 and later

Disk Drive Validation

During a cluster-level hardware compatibility check, vSphere Lifecycle Manager verifies that the disk drives that vSAN uses are supported and certified as per the vSAN Hardware Compatibility List (vSAN HCL). vSphere Lifecycle Manager also ensures that the disk drive firmware version specified in the cluster image is compatible with the cluster hardware.

The disk drives in a vSAN cluster and the firmware installed on the drives are of paramount importance for the overall vSAN cluster health. For example, a faulty disk drive firmware might cause performance issues and unexpected vSAN input-output behavior. You can use vSphere Lifecycle Manager hardware support managers to perform disk drive firmware upgrades. Before you upgrade the disk drive firmware, however, you must ensure that the target firmware version is supported as per the vSAN HCL.

Note: For SAS and SATA disk drives, the vSAN HCL lists the earliest supported firmware version. All firmware versions later than the specified in the vSAN HCL are supported. For NVMe devices behind a VMD controller, the vSAN HCL lists the exact driver-firmware combination certified for use with a given ESXi base image version. Only that specific driver-firmware version combination listed in the vSAN HCL is supported.

Supported Disk Drives Types

vSphere Lifecycle Manager validates the following types of disk drives and storage device configurations:

- HDD (SAS/SATA)
- SSD (SAS/SATA)
- SAS/SATA disk drives behind single-disk RAID-0 logical volumes

- NVMe devices behind a VMD controller

Note: vSphere Lifecycle Manager treats NVMe devices that are not VMD-enabled as PCI storage controllers. In the vSphere Client, hardware compatibility information about NVMe devices that are not VMD-enabled is available on the **PCI Devices** tab and not on the **Disks** tab.

System Requirements for Disk Drive Validation

- vCenter Server 7.0 Update 3 and later
- ESXi 7.0 and later
- For validation of NVMe devices behind a VMD controller, vCenter Server 8.0 and later and ESXi 8.0 and later

Important: Hardware compatibility checks do not validate NVMe devices behind a VMD controller, if the ESXi version of the hosts is earlier than 8.0. If you want to check the hardware compatibility of NVMe devices behind a VMD controller for hosts running an earlier version of ESXi, for example 7.0 Update 3, you have the following workarounds:

- Override manually the compliance status of the NVMe device behind a VMD controller.
- Upgrade your hosts to ESXi 8.0 or later.

RAID-0 Logical Volumes

vSphere Lifecycle Manager can validate the physical SAS/SATA disk drives behind single-disk RAID-0 logical volumes. The following requirements exist:

- The RAID controller is in a RAID or mixed mode.

For more information about the RAID and mixed mode, see the VMware knowledge base article at <https://kb.vmware.com/s/article/53573>.

- vCenter Server 7.0 Update 3 and later
- ESXi 7.0 and later
- The hardware support manager must be upgraded and certified to work with vSphere 7.0 Update 3.

If the you do not use an upgraded version of the hardware support manager, the compliance status of the physical drives behind RAID-0 logical volumes is unknown. In this case, you must manually validate the disk drives and the target firmware version and override the compliance status for those disks.

Disk Drive Validation Results

vSphere Lifecycle Manager does not display a disk drive compatibility status and compatibility information for every single disk in the vSAN disk group. vSphere Lifecycle Manager groups the disk drives that vSAN uses by vendor, model, target firmware version, capacity, and part number. That is, all disk drives by the same vendor, the same model, and with the same target firmware version form one entry in the list of disk devices.

Disk drives can be compliant or non-compliant. In the cases when vSphere Lifecycle Manager cannot find a unique match for a disk device in the vSAN HCL, vSphere Lifecycle Manager prompts you to manually specify the exact device that you want to validate. vSphere Lifecycle Manager then calculates the compliance status based on your selection.

When vSphere Lifecycle Manager is unable to determine the disk drive compliance, the respective devices are listed as non-compliant. You can manually validate those devices and set the compliance status to compliant or non-compliant. For more information, see [Change the Compliance Status of a Disk Device Manually](#).

For each entry in the disk device list, you can view summarized information about the disk, the compliance status, the number of affected hosts, and a label that shows whether the compliance status is manually

set or whether the device is certified. The **Used by vSAN** label is attached to all disk devices that vSAN

The screenshot shows the vSphere Hardware Compatibility pane. On the left, a sidebar lists 'Hosts', 'Image', 'Hardware Compatibility' (selected), 'VMware Tools', and 'VM Hardware'. The main pane is titled 'Hardware Compatibility' and includes a 'RUN CHECKS' button. A message states: 'There were changes made since last check. Re-run checks for up to date compliance results'. Below this, it says 'Discover and resolve potential hardware compatibility issues between your hosts and the cluster image.' and 'Checked on 08/03/2021, 7:30:02 AM'. There are tabs for 'PCI DEVICES' and 'DISKS'. Under 'Non-compliant Disks', there is a 'Hide All' button. A list of devices is shown, with one entry highlighted: 'VMware, Virtual disk, 286.102 GB'. This entry has two buttons: 'Used by vSAN' and 'User Reviewed'. Below the list, a table provides device information:

Vendor	Model	Capacity	Firmware Version	Part #
VMware	Virtual disk	286.102 GB	Unknown	Unknown

Below the table, it says 'Hosts Affected' followed by IP addresses and a 'View All' link. At the bottom, there is a 'CHANGE CLASSIFICATION' button.

uses.

If you expand the entry, you can view detailed compliance information about the respective disk device and the affected hosts.

When a new disk is added to a vSAN cluster, you must manually re-run the check to obtain the new compliance information about the cluster. Similarly, if you remove a disk from the vSAN disk group, you must re-run the hardware compatibility check to obtain an updated compliance information about the cluster.

Check the Hardware Compatibility of a Cluster

For a vSAN cluster that you manage with a single image, you can check the compliance between the image components and the hardware in the cluster. The check is performed against the vSAN Hardware Compatibility List (vSAN HCL) and ensures that if the image is applied to the hosts, the result after remediation is in accordance with the vSAN HCL.

- Verify that vSAN is enabled for the cluster.
- Verify that the cluster uses a single image.
- Verify that all hosts in the cluster are from the same vendor.
- To validate the compatibility between the PCI device and disk device hardware and target firmware version, verify that the image for the cluster includes a firmware add-on.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

1. In the vSphere Client, navigate to a vSAN cluster that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Hardware Compatibility**.
In the **Hardware Compatibility** pane, you see the results from the previous compatibility check.
3. In the **Hardware Compatibility** pane, click the **Run Checks** button.

vSphere Lifecycle Manager displays all compatibility information and issues in the **Hardware Compatibility** pane. You can see detailed compatibility information for each PCI device or disk drive. Review the result from the hardware compatibility check.

Resolve any issues before you remediate the cluster.

Change the Compliance Status of a Disk Device Manually

You can manually change the compliance status of a disk device and mark it as compliant or non-compliant.

- vCenter Server 7.0 Update 3
- Verify that vSAN is enabled for the cluster.
- Verify that the cluster uses a single image.
- Verify that all hosts in the cluster are from the same vendor.
- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).

In cases when the compliance status of a disk drive is unavailable, you must perform the hardware compatibility check manually and mark the device as compliant or non-compliant.

1. In the vSphere Client, navigate to a vSAN cluster that you manage with a single image.
2. On the **Updates** tab, select **Hosts** > **Hardware Compatibility**.
In the **Hardware Compatibility** pane, you see the results from the previous compatibility check.
3. Click the **Disks** button.
You see a list of all non-compliant and compliant disk devices.
4. Click the disk device whose compliance status you want to override.
5. Click the **Change Classification** button.
A drop-down menu appears.

The screenshot shows the vSphere Client interface for Hardware Compatibility. The left sidebar has a menu with 'Hosts', 'Image', 'Hardware Compatibility' (selected), 'VMware Tools', and 'VM Hardware'. The main pane is titled 'Hardware Compatibility' and includes a 'RUN CHECKS' button. Below the title, it says 'Discover and resolve potential hardware compatibility issues between your hosts and the cluster image.' and 'Checked on 08/03/2021, 7:30:02 AM'. There are two tabs: 'PCI DEVICES' and 'DISKS' (selected). Below the tabs, it says 'Non-compliant Disks 01 Hide All'. A dark blue banner shows 'VMware, Virtual disk, 286.102 GB' with a 'Used by vSAN' tag and the message 'Device compliance status cannot be computed. Firmware version unknown. 4 Hosts'. Below this is a 'Device Info' table:

Vendor	Model	Capacity	Firmware Version	Part #
VMware	Virtual disk	286.102 GB	Unknown	Unknown

Below the table is a red error message: 'The system has insufficient information to find a corresponding VCG listing. Manually verify device compliance.' At the bottom, it says 'Hosts Affected 10.41.78.226, 10.41.78.225, 10.41.78.228, 10.41.78.229 View All'. A red box highlights the 'CHANGE CLASSIFICATION' button, which has a dropdown menu open with three options: 'Mark Issue As Verified', 'Flag as Non-Compliant', and 'Don't Override'.

6. Select the compliance status to apply to the disk device.
 - To mark the issue as compliant, select **Mark Issue As Verified**.
 - To mark the issue as non-compliant, select **Flag as Non-Compliant**.

The **Mark Issue As Verified** option is not available for compliant devices. The **Flag as Non-Compliant** option is unavailable if the disk device is non-compliant.

7. Optional: To undo the override selection, select **Don't Override** from the **Change Classification** drop-down menu.

You changed the compliance status of a disk device. However, the disk device remains in its original list until you run a new hardware compatibility check.

If you marked the device as verified, a **User Reviewed** label appears for the disk group.



If you marked the device as non-compliant, a **Flagged** label appears for the disk



group.

Run a new hardware compatibility check so that the new compliance status for a disk is saved.

Host-Level Hardware Compatibility Checks

You can run a hardware compatibility check for any host to determine which ESXi version the host hardware is compatible with. The hardware compatibility check ensures that the host hardware, that is server model and I/O devices, is certified for use with a selected ESXi version. Host-level hardware compatibility checks are performed against the VMware Compatibility Guide (VCG), even if the host is in a vSAN cluster.

You can check the hardware compatibility of any host, whether it is in a cluster that uses a single image or baselines. You can also check the hardware compatibility of a standalone host. Host-level hardware compatibility checks don't validate the firmware on the host, even if it is part of a cluster that you manage with a single image or if it is a standalone host that you manage with a vSphere Lifecycle Manager image.

After the hardware compatibility check, vSphere Lifecycle Manager shows the compliance status for the server and hardware devices. The server and devices might have one of the three different states: compatible, incompatible, and unknown. For more information about compatibility statuses, see [Hardware Compatibility Report for a Host](#).

If the server status is incompatible, vSphere Lifecycle Manager does not proceed with checking the compatibility for the hardware devices.

Check the Hardware Compatibility of a Host

You can check the hardware compatibility of a host to determine whether the host hardware is certified for use with a selected ESXi version. The hardware compatibility check is performed against the VMware Compatibility Guide (VCG) or, if the host is in a vSAN cluster, against the vSAN Hardware Compatibility List (HCL)

- If needed, synchronize hardware compatibility data. See [Sync Hardware Compatibility Data](#).
- Verify that the Customer Experience Improvement Program is enabled.
- Verify that vCenter Server is connected to the Internet.
- Verify that the host is not part of an VxRail environment.

- Verify that you have the proper privileges. See [vSphere Lifecycle Manager Privileges For Using Images](#).
1. In the vSphere Client, navigate to a standalone host or a host in a cluster.
 2. On the **Updates** tab, select **Hosts** > **Hardware Compatibility**.
 3. In the **Hardware Compatibility** pane, select your task.
 - To run a hardware compatibility check for the host for the first time, select a target ESXi from the drop-down menu and click **Apply**.
 - To check the hardware compatibility between the host and the already selected target ESXi version, click **Re-run Checks**.
 - To choose a new target ESXi version for the hardware compatibility check, click **Edit** and select a new target ESXi version.
 - To export the hardware compatibility report in a CSV format, click the **Export** button.

vSphere Lifecycle Manager displays the result from the compatibility check. You can see a list of the compatible, incompatible, and unknown devices. For each device, you can see full details by clicking the expand button.

Hardware Compatibility Report for a Host

The hardware compatibility report gives you information whether for a selected server model and hardware devices, vSphere Lifecycle Manager finds records for a target ESXi version in the VMware Compatibility Guide (VCG).

Server Hardware Compatibility

Host Model is Not Compatible

This compatibility status indicates that there are no records for the selected ESXi version in the VCG. If the host is not compatible with the selected ESXi version, vSphere Lifecycle Manager does not proceed to checking the compatibility of the devices.

In the **Host Model Compatibility** card, you can see details about the host: server model name, CPU model, and the BIOS version running on the host. At the bottom of the card, you see a list of all certified CPU series for the target ESXi version.

Host Model is Compatible

This compatibility status indicates that the host is certified for use with the selected ESXi version as per VCG. When the host is compatible, vSphere Lifecycle Manager proceeds with the device validation.

In the **Host Model Compatibility** card, you can see details about the host: server model name, CPU model, and the BIOS version running on the host. Because in VCG the information about CPUs is based on CPU series, and not specific models, you might need to manually check if the CPU of the host is part of the supported CPU series. You might also need to manually check if the BIOS version on the host matches any of the compatible BIOS versions for the CPU series as per VCG.

Hardware Compatibility Checks Not Supported for the Host Vendor Model

When the server model is not part of the list of certified OEMs, vSphere Lifecycle Manager does not perform a hardware compatibility check and you do not see a hardware compatibility report for the selected host.

Device Hardware Compatibility

The compatibility statuses for devices are: compatible, incompatible, and unknown.

Unknown

Unknown devices are devices for which no records exist in VCG. When you click the expand button for the device, you see the following device information: device IDs, driver and firmware currently running on the device. No compatibility data is generated and displayed.

The unknown status might also indicate that multiple matches exist in the VCG for the respective device. In such cases, use the device ID to manually check if the hardware device matches any of the supported devices for the target ESXi version in the VCG.

Incompatible

The incompatible status indicates that no records exist in the VCG for the selected ESXi version. When you click the expand button for the device, you see information about the ESXi versions that are compatible with the device as per VCG.

Compatible

The compatible status indicates that the device is compatible with the selected ESXi version as per VCG. When you click the expand button for the device, you see the following device information: device IDs, driver and firmware currently running on the device. For compatible devices, you might need to manually confirm that the driver-firmware combination running on the device is supported as per VCG.

Sync Hardware Compatibility Data

To initiate a hardware compatibility check for a host, the hardware compatibility data from VMware Compatibility Guide (VCG) must become available to vSphere Lifecycle Manager.

Verify that vCenter Server can access the following sites:

- vvs.esp.vmware.com
- auth.esp.vmware.com

Synchronizing compatibility ensures that the compatibility information from VCG becomes available to vSphere Lifecycle Manager. The synchronization task is not automated. When no compatibility data is available for use to vSphere Lifecycle Manager, you must trigger the compatibility data synchronization manually.

vSAN HCL data is not updated through synchronization. If you want to check the hardware compatibility of a host that is in a vSAN cluster, you must first verify that vSAN HCL data is up to date. For more information about updating vSAN HCL data, see the vSAN documentation.

1. In the vSphere Client, navigate to a standalone host or a host in a cluster.
2. On the **Updates** tab, select **Hosts** > **Hardware Compatibility**.
3. In the **Hardware Compatibility** pane, click **Sync compatibility data**.
4. In the **Sync hardware compatibility data** dialog box, click **Go to Lifecycle Manager**.
You are redirected to the vSphere Lifecycle Manager home view.
5. Select **Actions** > **Sync HCL**.
The Update HCL data task appears in the **Recent Taks** pane.

After the Update HCL data task finishes, the compatibility data from the VCG becomes available to vSphere Lifecycle Manager.

Check the hardware compatibility of your hosts against VCG before you update or upgrade them to a later ESXi version.

Using vSphere Lifecycle Manager Baselines and Baseline Groups

Using baselines and baseline groups to manage the updates and upgrades of ESXi hosts is a multi-stage process.

1. Populate the vSphere Lifecycle Manager repository with patches, extensions, and updates.
The vSphere Lifecycle Manager repository contains software updates that you can use with both vSphere Lifecycle Manager baselines and vSphere Lifecycle Manager images. On the **Updates** tab of the vSphere Lifecycle Manager home view, you see all software updates available in the vSphere Lifecycle Manager depot as bulletins.
Updates get into the vSphere Lifecycle Manager local depot through synchronization with configurable download sources. By default, vSphere Lifecycle Manager is configured to synchronize updates from the official VMware depot. You can also import updates into the depot manually.
For host upgrade operations through baselines, you use ESXi ISO images, which you must also import to the vSphere Lifecycle Manager depot manually.
For more information about working with the vSphere Lifecycle Manager depot, see [The vSphere Lifecycle Manager Depot](#).
2. Create baselines by combining bulletins from the depot and using manually uploaded ESXi ISO images.
You can also combine several non-conflicting baselines to create a baseline group. Baseline groups can contain different types of baselines. If a baseline group contains both upgrade and patch or extension baselines, the upgrade runs first.
For more information about creating baselines and baseline groups, see [Types of vSphere Lifecycle Manager Baselines and Baseline Groups](#).
3. Attach the baselines to individual ESXi hosts or container objects for ESXi hosts.
For more information, see [Attach Baselines and Baseline Groups to Objects](#).
4. Check the compliance of ESXi hosts against a selected baseline or baseline group.
You can run a compliance check on an individual ESXi host or a container object.
For more information about compliance checks against baselines and baseline groups, see [Checking Compliance Against vSphere Lifecycle Manager Baselines and Baseline Groups](#).
5. Review the compliance status of the scanned object.
For more information about compliance states, see [Viewing Compliance Information About ESXi Hosts and Updates](#).
6. Optionally, you can stage the patches and extensions to ESXi hosts before remediation. Staging is not a mandatory step, it is a step that you can skip.
For more information about staging updates before remediation, see [Stage Patches and Extensions to ESXi Hosts](#).
7. Remediate the non-compliant objects. After remediation, you can review the compliance status again to make sure that the updates are installed.
For more information about remediating objects against baselines and baseline groups, see [Remediating ESXi Hosts Against vSphere Lifecycle Manager Baselines and Baseline Groups](#).

Types of vSphere Lifecycle Manager Baselines and Baseline Groups

You use baselines and baseline groups to update the ESXi hosts in your vSphere inventory. The vSphere Lifecycle Manager baselines are three types: predefined baselines, recommendation baselines, or custom baselines, which you create manually. Depending on their content, baselines can be patch, extension, or upgrade baselines.

When you initiate a compliance check for an ESXi host, you evaluate it against baselines and baseline groups to determine its level of compliance to those baselines or baseline groups.

If your vCenter Server system is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, the baselines and baseline groups that you create and manage are applicable only to the inventory objects managed by the vCenter Server system where the selected vSphere Lifecycle Manager instance runs.

In the vSphere Client, the baselines and baseline groups are displayed on the **Baselines** tab of the vSphere Lifecycle Manager home view.

Predefined, Recommendation, and Custom Baselines

Predefined baselines

Predefined baselines cannot be edited or deleted, you can only attach or detach them to inventory objects.

Predefined baselines can be of the following type:

- **Host Security Patches**
The Host Security Patches baseline checks ESXi hosts for compliance with all security patches.
- **Critical Host Patches**
The Critical Host Patches baseline checks ESXi hosts for compliance with all critical patches.
- **Non-Critical Host Patches**
The Non-Critical Host Patches baseline checks ESXi hosts for compliance with all optional patches.

The Host Security Patches, and Critical Host Patches predefined baselines are attached by default to the vCenter Server instance where vSphere Lifecycle Manager runs.

Recommendation Baselines

Recommendation baselines are predefined baselines that vSAN generates.

You use recommendation baselines to update your vSAN clusters with recommended critical patches, drivers, updates, or the latest supported ESXi host version for vSAN.

These baselines appear by default when you use vSAN clusters with ESXi hosts of version 6.0 Update 2 and later in your vSphere inventory. If your vSphere environment does not contain any vSAN clusters, no recommendation baselines are created.

Recommendation baselines update their content periodically, which requires vSphere Lifecycle Manager to have constant access to the Internet. The vSAN recommendation baselines are typically refreshed every 24 hours.

Recommendation baselines cannot be edited or deleted. You do not attach recommendation baselines to inventory objects in your vSphere environment. You can create a baseline group by combining multiple recommendation baselines, but you cannot add any other type of baseline to that group. Similarly, you cannot add a recommendation baseline to a baseline group that contains upgrade, patch, and extension baselines.

Custom Baselines

Custom baselines are the baselines that you create. You can create custom patch, extension, and upgrade baselines to meet the needs of your specific deployment.

Upgrade, Patch, and Extension Baselines

Upgrade Baselines

Host upgrade baselines define the version to which you upgrade the hosts in your environment. With vSphere Lifecycle Manager 8.0, you can upgrade ESXi hosts from version 6.7 and 7.0 to ESXi 8.0. Host upgrades to ESXi 5.x, ESXi 6.7, or ESXi 7.0 are not supported. In case of an unsuccessful upgrade from ESXi 6.7 or ESXi 7.0 to ESXi 8.0, you cannot roll back to your previous ESXi 6.7 or ESXi 7.0 instance.

To create an upgrade baseline, you must first import an ESXi ISO image to the vCenter Server inventory.

You can use the ESXi installer image distributed by VMware with the name format `VMware-VMvisor-Installer-8.0-build_number.x86_64.iso` or a custom image created by using vSphere ESXi Image Builder. You can also use ISO images created and distributed by OEMs.

Patch Baselines

Patch baselines define a number of patches that must be applied to a given host. Patch baselines can be either dynamic or fixed.

Baseline	Description
Dynamic Patch Baseline	A dynamic baseline is a set of patches that meet certain criteria. You specify the criteria for patch inclusion in the baseline. Only the patches that meet the criteria are included in the baseline. As the set of available patches in the vSphere Lifecycle Manager depot changes, dynamic baselines are updated as well. You can manually include or exclude patches from the baseline.
Fixed Patch Baseline	A fixed baseline is a set of patches that does not change as patch availability in the depot changes. You manually select the patches from the total set of patches available in the vSphere Lifecycle Manager depot.

Extension Baselines

Extension baselines contain additional software modules for ESXi hosts, for example device drivers. This additional software might be VMware software or third-party software. All third-party software for ESXi hosts is classified as host extension, but extensions are not restricted to just third-party software. You can install additional modules by using extension baselines, and update the installed modules by using patch baselines.

Extensions deliver additional host features, updated drivers for hardware, Common Information Model (CIM) providers for managing third-party modules on the host, improvements to the performance or usability of the existing host features, and so on.

The host extension baselines that you create are always fixed. You must carefully select the appropriate extensions for the ESXi hosts in your environment.

You use extension baselines to install extensions on the ESXi hosts in your environment. After an extension is installed on a host, you can update the extension module through either patch, or extension baselines.

Note: When you use extension baselines, you must be aware of the functional implications that the installation of new modules on

the host might have. Extension modules might alter the behavior of ESXi hosts. During the installation of extensions, vSphere Lifecycle Manager only performs the checks and verifications expressed at the package level.

Baseline Groups

You create a baseline group by assembling existing and non-conflicting baselines. Baseline groups allow you to scan and remediate objects against multiple baselines at the same time.

The following are valid combinations of baselines that can make up a baseline group:

- Multiple host patch and extension baselines.
- One upgrade baseline, multiple patch, and extension baselines.

To create, edit, or delete baselines and baseline groups, you must have the **Manage Baseline** privilege. To attach baselines and baseline groups to target inventory objects, you must have the **Attach Baseline** privilege. The privileges must be assigned on the vCenter Server system where vSphere Lifecycle Manager runs.

For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation.

For a list of all vSphere Lifecycle Manager privileges and their descriptions, see [vSphere Lifecycle Manager Privileges For Using Baselines](#).

Creating Baselines in vSphere 7.0 and Later Releases

Because in vSphere 7.0 and later releases the official VMware online depot hosts certified partner content in addition to VMware content, a broader set of OEM bulletins are available in the vSphere Lifecycle Manager depot. As a result, in the **Create Baseline** and **Edit Baseline** wizards, you also see a broader set of OEM bulletins. Some of these bulletins might have dependencies that must be pulled into the baselines that you create, so that the remediation against those baselines is successful. Always consult the KB article for an individual bulletin before you include it in a baseline. The KB article contains information about the bulletin deployment specifics and required dependencies. You must include in the baseline, only bulletins compatible with the hardware on which the host is running. Otherwise, remediation might fail.

Starting with vSphere 7.0, some changes are also introduced in the way VMware content is packaged. As a result, at patch and update releases, you might see additional bulletins on the patch selection page of the **Create Baseline** and **Edit Baseline** wizards. Those bulletins are usually of the Enhancement or BugFix category. When you include those bulletins in a baseline, you might need to also include base ESXi bulletins in that baseline. To ensure successful application of VMware patches and updates, always include the appropriate roll-up bulletin into your baselines. Otherwise, remediation might fail.

Create, Edit, or Delete vSphere Lifecycle Manager Baselines and Baseline Groups

You can create upgrade baselines for ESXi hosts with ESXi8.0 images, which you must first import to the vSphere Lifecycle Manager depot.

- Required privileges: **VMware vSphere Lifecycle Manager > Manage Baselines**.

- To create an upgrade baseline, verify that you have an ESXi8.0 image available in inventory. For more information, see [Import an ISO Image to the vSphere Lifecycle Manager Depot](#).
1. In the vSphere Client, click the bars icon and select **Lifecycle Manager**.
 2. Optional: Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.
The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
 3. On the **Baselines** tab, select **New > Baseline**.
If you want to create an upgrade baseline, you can start the **Create Baseline** wizard from the **Imported ISOs** tab. Select an ESXi image from the list and click **New Baseline**.
 4. On the **Name and Description** page, enter information about the baseline, select the baseline type and click **Next**.
 5. Follow the prompts to finish creating the baseline.

Baseline Type	Steps
Patch Baseline	To create a fixed patch baseline, perform the following steps: 1. On the Select Patches Automatically page, stop the automatic updates by deselecting the respective check box and click Next. 2. On the Select Patches Manually page, select the patches that you want to include in the baseline and click Next. – To view only the rollup bulletins in the list, turn on the Show only rollup updates toggle switch. – To filter the patches that are available in the vSphere Lifecycle Manager depot and find specific patches to include in the baseline, use the filter icon next to each column header. If you use several criteria to filter the patches, the relationship between those filter criteria is defined by the Boolean operator AND. 3. On the Summary page, review your selections and click Finish. To create a dynamic patch baseline, perform the following steps: 1. On the Select Patches Automatically page, set the criteria for adding patches to the baseline. a. Enable the automatic update of the baseline by selecting the respective check box. b. On the Criteria tab, specify the criteria that a patch must meet to be added to the baseline and click Next. Note: The relationship between the selected criteria is defined by the Boolean operator AND. For example, when you select a product and severity option, the patches are restricted to the ones that are applicable for the selected product and are of the specified severity level. • Patch Vendor specifies which patch vendor to use. Note: In vSphere 7.0 and later releases, the vendor name of VMware for inbox components is changed from VMware, Inc to VMware. As a result, if you use the filter to see only components by VMware, the

Baseline Type	Steps
	filtered list contains both VMware, Inc for 6.x patches and VMware for 7.x and 8.x patches. • Product restricts the set of patches to the selected products or operating systems. Note: The asterisk at the end of a product name is a wildcard character for any version number. • Severity specifies the severity of patches to include. • Category specifies the category of patches to include. • Release Date specifies the range for the release dates of the patches. c. On the Matched tab, deselect patches from the ones that matched your criteria to exclude them permanently from the baseline. d. On the Excluded and Selected tabs, view the patches that are excluded from the baseline and the ones that are included in the baseline. You can use the filter icon next to each column header on the Matched, Excluded, and Selected tabs to filter the patches that are available in the vSphere Lifecycle Manager depot. This way, you can easily find specific patches to exclude from or include in the baseline. If you use several criteria to filter the patches, the relationship between those filter criteria is defined by the Boolean operator AND. 2. On the Select Patches Manually page, select individual patches to include in the baseline and click Next. The patches that are displayed on this page are patches that do not meet the criteria you set on the Select Patches Automatically page. You can use the filter icon next to each column header to filter the patches that are available in the vSphere Lifecycle Manager depot and find specific patches to include in the baseline. If you use several criteria to filter the patches, the relationship between those filter criteria is defined by the Boolean operator AND. The patches that you add manually to the dynamic baseline stay in the baseline regardless of the automatically downloaded patches. 3. On the Summary page, review your selections and click Finish.
Upgrade Baseline	1. On the Select ISO page, select an ESXi image from the list and click Next. 2. On the Summary page, review your selections and click Finish.
Extension Baseline	1. On the Select Extensions page, select individual extensions to include in the baseline and click Next. You can use the filter icon next to each column header to filter the extensions that are available in the vSphere Lifecycle Manager depot and find specific extensions to include in the baseline. If you use several criteria to filter the patches, the relationship between those filter criteria is defined by the Boolean operator AND. 2. On the Summary page, review your selections and click Finish.

The new baseline appears in the baselines list on the **Baselines** tab. You can attach the baseline to a data center, a cluster, or a host.

Create a Host Baseline Group

You can combine multiple baselines of different types into a baseline group. For example, you can combine one host upgrade baseline with multiple patch or extension baselines, or you can combine multiple patch and extension baselines.

Required privileges: **VMware vSphere Lifecycle Manager > Manage Baselines**

A baseline group might contain a single host upgrade baseline and multiple patch or extension baselines, or a combination of host patch and host extension baselines.

You can create a baseline group and add baselines to it later.

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Baselines** tab, select **New > Baseline Group**.
The **Create Baseline Group** wizard opens.
3. On the **Name and Description** page, enter a unique name and, optionally, a description for the baseline group, and click **Next**.
4. Optional: On the **Upgrade Baseline** page, select an upgrade baseline to include in the baseline group and click **Next**.
 - a) Select the **Add the following Upgrade Baseline to the Group** check box.
 - b) Select an upgrade baseline from the list.
5. Optional: On the **Patch Baselines** page, select patch baselines to include in the baseline group and click **Next**.
6. Optional: On the **Extension Baselines** page, select extension baselines to include in the baseline group and click **Next**.
7. On the **Summary** page, review your selections and click **Finish**.

The new host baseline group appears in the baselines list on the **Baselines** tab. You can attach the baseline group to a data center, a cluster, or a host.

Edit a vSphere Lifecycle Manager Baseline or Baseline Group

You can edit existing patch, upgrade, or extension baselines.

Required privileges: **VMware vSphere Lifecycle Manager > Manage Baselines**

1. In the vSphere Client, click the bars icon and select **Lifecycle Manager**.
2. Optional: Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
3. On the **Baselines** tab, select a baseline or a baseline group from the list and click **Edit**.
The **Edit Baseline** or **Edit Baseline Group** wizard appears.

4. Follow the prompts to finish editing the selected baseline or baseline group.
- Attach the edited baseline or baseline group to a data center, a cluster, or a host.

Add or Remove a Single Update from a Custom Baseline

You can edit the content of a custom baseline by adding or removing individual patches or extension from the baseline.

Required privileges: **VMware vSphere Lifecycle Manager > Manage Baselines**

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. Click the **Updates** tab.

You see a list of all updates in the vSphere Lifecycle Manager depot.
3. Select a patch or extension from the list, and click **Add/Remove Baselines**.

The **Add/Remove baselines** dialog box opens.
4. Select your task.
 - To add the patch to a baseline, select that baseline in the **Custom Patch Baselines** list.
 - To remove the patch from a baseline, deselect that baseline in the **Custom Patch Baselines** list.
5. Click **OK**.

Delete Baselines and Baseline Groups

You can delete the baselines and baseline groups that you no longer need. Deleting a baseline detaches it from all the objects to which it is attached. You cannot delete predefined and system-managed baselines.

Required privileges: **VMware vSphere Lifecycle Manager > Manage Baselines**

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Baselines** tab, select a baseline or a baseline group from the list and click **Delete**.

The **Delete Baseline** dialog box opens.
3. Click **Yes** to confirm the deletion of the selected baseline or baseline group.

Attach Baselines and Baseline Groups to Objects

To check the compliance status of the objects in your inventory against selected baselines and baseline groups, you must first attach the respective baselines and baseline groups to the objects.

- If you want to attach a baseline or a baseline group to a cluster, verify that the cluster is not configured to use a single image.

- Required privileges: **VMware vSphere Lifecycle Manager > Manage Baselines > Attach Baseline**.

You attach baselines and baseline groups to individual hosts or objects that contain hosts, such as clusters, data centers, and vCenter Server instances. In the vSphere infrastructure hierarchy, the baseline and baseline groups that you attach to container objects are also attached to the child objects. For example, if you attach a baseline or baseline group to a folder, the baseline or the baseline groups is inherited by all the objects in the folder, including subfolders.

You cannot use vSphere Lifecycle Manager to update the hosts in a cluster that uses a single vSphere Lifecycle Manager image. For more information about using vSphere Lifecycle Manager images to manage hosts in clusters collectively, see [How to Use vSphere Lifecycle Manager Images](#).

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Hosts > Baselines**.
3. In the **Attached Baselines** pane, click **Attach > Attach Baseline or Baseline Group**.
The **Attach** dialog box opens.
4. Select one or more baselines or baseline groups and click **Attach**.
If you select a baseline group, all the baselines in it are attached to the object.
The selected baselines or baseline groups are attached to the object. They appear on the list in the **Attached Baselines** pane. If the selected object is a container object, the selected baselines or baseline groups are attached to all the child objects.

Scan the selected object against the attached baselines.

Detach Baselines and Baseline Groups from Objects

You can detach baselines and baseline groups from the objects to which the baselines or baseline groups are directly attached.

- Verify that the cluster is not configured to manage all its hosts collectively.
- Required privileges: **VMware vSphere Lifecycle Manager > Manage Baselines > Attach Baseline**.

vSphere inventory objects might have inherited properties, so instead of detaching baselines and baseline groups directly from an object, you might need to select its container object and detach the baselines or baseline groups from the container object. For example, if you want to detach a baseline or a baseline group from a host that is a part of a cluster, you must select the cluster and not the host.

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Hosts > Baselines**.
3. In the **Attached Baselines** pane, select one or more baselines or baseline groups and click **Detach**.
The **Detach** dialog box opens.
4. Select the object to detach the baseline or baseline group from and click **Detach**.
The baseline is removed from the **Attached Baselines** list.

Checking Compliance Against vSphere Lifecycle Manager Baselines and Baseline Groups

Before you update or upgrade an ESXi host or a container object with vSphere Lifecycle Manager baselines, you must first check its compliance status.

You use vSphere Lifecycle Manager to check the compliance status of ESXi hosts against the baselines and baseline groups that you attach to the hosts or to a parent container object. You do a compliance check on hosts to determine whether they have the latest patches or extensions. During the compliance check, attributes of the host are evaluated against all patches, extensions, and upgrades from an attached baseline or baseline group.

You can check the compliance status of a single ESXi host or a valid container object. Supported groups of ESXi hosts include virtual infrastructure container objects such as folders, clusters, and data centers. When you initiate a compliance check for a container object, vSphere Lifecycle Manager scans all the ESXi hosts in that container object.

Note: If you initiate a compliance check for an inventory object, for example data center, that contains clusters that use vSphere Lifecycle Manager images, the compliance check is not performed for those clusters.

To generate compliance information, you can initiate compliance checks manually or you can schedule the compliance checks to run at regular periods. Schedule compliance checks at a data center or vCenter Server system level to make sure that the objects in your inventory are up-to-date.

You check the compliance status of vSphere objects from the vSphere Lifecycle Manager compliance view.

To initiate or schedule compliance checks, you must have the **Scan for Applicable Patches, Extensions, and Upgrades** privilege.

For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation.

For a list of all vSphere Lifecycle Manager privileges and their descriptions, see [vSphere Lifecycle Manager Privileges For Using Baselines](#).

Initiate a Compliance Check for ESXi Hosts Manually

Before remediation, you must check the compliance of the vSphere objects against the attached baselines and baseline groups. To check the compliance status of hosts in the vSphere inventory immediately, initiate a compliance check manually.

If you want to check the compliance status of a cluster, verify that the cluster is not configured to use a single image.

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Hosts > Baselines**.

The **Baselines** pane shows three panels. In those panels, you obtain host information about the selected object, host compliance information, and remediation information.
3. In the compliance information panel, click **Check Compliance**.

The selected inventory object and all child objects are scanned against all attached patch, extension, and upgrade baselines. The larger the virtual infrastructure and the higher up in the object hierarchy that you initiate the scan, the longer the scan takes.

Schedule Regular Compliance Checks for ESXi Hosts

You can configure vSphere Lifecycle Manager to check the compliance status of ESXi hosts at specific times or at intervals that are convenient for you.

If you want to check the compliance status of a cluster, verify that the cluster is not configured to use a single image.

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Hosts > Baselines**.
The **Baselines** pane shows three panels. In those panels, you obtain host information about the selected object, host compliance information, and remediation information.
3. In the compliance information panel, click **Schedule**.
The **Automatic compliance check** dialog box opens.
4. Configure the compliance check schedule.
 - a) Set the frequency and the starting point of the compliance check.
 - b) Enter a unique name, and optionally, a description for the scan task.
 - c) Optional: Specify one or more email addresses to receive notification after the scan task is complete.
You must configure mail settings for the vCenter Server system to enable this option.
5. Click **Save** to exit the **Automatic compliance check** dialog box.

View Information About the Patches, Extensions, and ISO Images in a Baseline

You can view information about the patches, extensions, and upgrades included in a baseline or a baseline group.

- Verify that the host for which you want to view compliance information uses baselines and is not managed with a single vSphere Lifecycle Manager image.
- Review the [The vSphere Lifecycle Manager Compliance View](#) topic.

For information about the different compliance statuses that an update might have, see [Compliance Statuses of Updates](#).

1. In the vSphere Client, navigate to a single ESXi host, cluster, or a valid container object.
2. On the **Updates** tab, select **Hosts > Baselines**.
3. In the **Attached Baselines and Baseline Groups** pane, select a baseline.

A new pane appears below the **Attached Baselines and Baseline Groups** pane. Depending on the selected object, the bottom pane might contain information about the updates and ESXi images in the baseline that you select. If the selected object is a container for ESXi hosts, the bottom pane shows the compliance of each ESXi in the container object against the selected baseline.

Baseline Type	Available Information
Patch	The bottom pane contains a table that lists all patches in the baseline. For each update, you can see the following information. • Update Name • Update ID The update ID is a vendor-assigned identification code of the patch.

Baseline Type	Available Information						
	• Status The Status column shows the compliance status of the update. • Severity • Category • Impact The Impact column displays the actions that you must take to install the update. For example, rebooting the system or putting the host in maintenance mode. • ESXi Version						
Upgrade	The bottom pane displays the following information. • ESXi Version • Build • Status The Status column shows the compliance status of the update. • Release Date • Vendor • Details • Release Notes • Acceptance Level ESXi images can be either Signed or Unsigned, which indicates their level of acceptance by VMware. The software packages included in ESXi images might have any of the following acceptance levels. <table border="0"> <tr> <td data-bbox="857 1052 1055 1079">VMware Certified</td><td data-bbox="1192 1052 1521 1283">The package has gone through a rigorous certification program that verifies the functionality of the feature, and is signed by VMware with a private key. VMware provides customer support for these packages.</td></tr> <tr> <td data-bbox="857 1287 1055 1314">VMware Accepted</td><td data-bbox="1192 1287 1521 1629">The package has gone through a less rigorous acceptance test program that only verifies that the package does not destabilize the system, and is signed by VMware with a private key. The test regimen does not validate the proper functioning of the feature. VMware support hands off support calls directly to the partner.</td></tr> <tr> <td data-bbox="857 1633 1055 1661">Partner Supported</td><td data-bbox="1192 1633 1521 1799">The partner has signed an agreement with VMware and has demonstrated a sound test methodology. VMware provides a signed private/public key pair to the partner.</td></tr> </table>	VMware Certified	The package has gone through a rigorous certification program that verifies the functionality of the feature, and is signed by VMware with a private key. VMware provides customer support for these packages.	VMware Accepted	The package has gone through a less rigorous acceptance test program that only verifies that the package does not destabilize the system, and is signed by VMware with a private key. The test regimen does not validate the proper functioning of the feature. VMware support hands off support calls directly to the partner.	Partner Supported	The partner has signed an agreement with VMware and has demonstrated a sound test methodology. VMware provides a signed private/public key pair to the partner.
VMware Certified	The package has gone through a rigorous certification program that verifies the functionality of the feature, and is signed by VMware with a private key. VMware provides customer support for these packages.						
VMware Accepted	The package has gone through a less rigorous acceptance test program that only verifies that the package does not destabilize the system, and is signed by VMware with a private key. The test regimen does not validate the proper functioning of the feature. VMware support hands off support calls directly to the partner.						
Partner Supported	The partner has signed an agreement with VMware and has demonstrated a sound test methodology. VMware provides a signed private/public key pair to the partner.						

Baseline Type	Available Information
	to use for self-signing their packages. The VMware support team redirects support calls directly to the partner. Community Supported The package is either unsigned or signed by a key that is not cross-signed by VMware. VMware does not provide support for the package. For support, customers must either use the community or contact the author of the package.
Extension	• Update Name • Update ID The update ID is a vendor-assigned identification code of the extension. • Status The Status column shows the compliance status of the update. • Severity • Category • Impact The Impact column displays the actions that you must take to install the update. For example, rebooting the system or putting the host in maintenance mode. • ESXi Version
Baseline Group	To view information about the patches, extension, and ISO images in a baseline group, select the respective tab in the bottom pane. • Click Baselines for information about the baselines that a baseline group contains. • Click ISO for information about the ESXi image that a baseline group contains. • Click Updates for information about the patches and extensions that the baseline group contains.

Host Upgrade Compliance Messages

When you check the compliance of ESXi hosts against an upgrade baseline, vSphere Lifecycle Manager runs a precheck script and provides informative messages in the in the bottom pane of the vSphere Lifecycle Manager compliance view. The messages notify you about potential problems with hardware or third-party software on the host, and configuration issues, which might prevent a successful upgrade to ESXi8.0.

The messages that vSphere Lifecycle Manager provides correspond to error or warning codes from running the host upgrade precheck script.

For interactive installations and upgrades performed by using the ESXi installer, the errors or warnings from the precheck script are displayed on the final panel of the installer, where you are asked to confirm or cancel the installation or upgrade. For scripted installations and upgrades, the errors or warnings are written to the installation log.

vSphere Lifecycle Manager displays scan result information in the bottom pane of the vSphere Lifecycle Manager compliance view. To see the original errors and warnings returned by the precheck script during an vSphere Lifecycle Manager host upgrade scan operation, review the vSphere Lifecycle Manager log file.

Table 123: Scan Result Messages and Corresponding Error and Warning Codes

Scan Result Message in vSphere Lifecycle Manager	Description
Host CPU is unsupported. New ESXi version requires a 64-bit CPU with support for LAHF/SAHF instructions in long mode.	This message appears if the host processor is 32-bit and does not support required features. The corresponding error code is 64BIT_LONGMODESTATUS.
Trusted boot is enabled on the host but the upgrade does not contain the software package <code>esx-tboot</code> . Upgrading the host will remove the trusted boot feature.	This message indicates that the host upgrade scan did not locate the <code>esx-tboot</code> VIB on the upgrade ISO. The corresponding error code is TBOOT_REQUIRED
VMkernel and Service Console network interfaces are sharing the same subnet <code>subnet_name</code> . This configuration is not supported after upgrade. Only one interface should connect to subnet <code>subnet_name</code> .	Warning. An IPv4 address was found on an enabled Service Console virtual NIC for which there is no corresponding address in the same subnet in the vmkernel. A separate warning appears for each such occurrence. The corresponding error code is COS_NETWORKING.
New ESXi version requires a minimum of <code>core_count</code> processor cores.	The host must have at least two cores. The corresponding error code is CPU_CORES.
Processor does not support hardware virtualization or it is disabled in BIOS. Virtual machine performance may be slow.	Host performance might be impaired if the host processor does not support hardware virtualization or if hardware virtualization is not turned on in the host BIOS. Enable hardware virtualization in the host machine boot options. See your hardware vendor's documentation. The corresponding error code is HARDWARE_VIRTUALIZATION.
Insufficient memory, minimum <code>size_in_MB</code> required for upgrade.	The host requires the specified amount of memory to upgrade. The corresponding error code is MEMORY_SIZE.
Host upgrade validity checks for <code>file_name</code> are not successful.	This test checks whether the precheck script itself can be run. The corresponding error code is PRECHECK_INITIALIZE.
The host partition layout is not suitable for upgrade.	Upgrade is possible only if there is at most one VMFS partition on the disk that is being upgraded and the VMFS partition starts after sector 1843200. The corresponding error code is PARTITION_LAYOUT.
Unsupported configuration.	The file <code>/etc/vmware/esx.conf</code> must exist on the host. This message indicates that the file <code>/etc/vmware/esx.conf</code> is either missing, or the file data cannot be retrieved or read correctly. The corresponding error code is SANE_ESX_CONF.
The host does not have sufficient free space on a local VMFS datastore to back up current host configuration. A minimum of <code>size_in_MB</code> is required.	The host disk must have enough free space to store the ESXi 5.x configuration between reboots. The corresponding error code is SPACE_AVAIL_CONFIG.
The upgrade is not supported for current host version.	Upgrading to ESXi8.0 is possible only from ESXi6.7 and ESXi7.0 hosts. The corresponding error code is SUPPORTED_ESX_VERSION.
Unsupported devices <code>device_name</code> found on the host.	The script checks for unsupported devices. Some PCI devices are not supported with ESXi8.0.

Scan Result Message in vSphere Lifecycle Manager	Description
	The corresponding error code is <code>UNSUPPORTED_DEVICES</code> .
Host software configuration requires a reboot. Reboot the host and try upgrade again.	To ensure a good bootbank for the upgrade, you must reboot the hosts before remediation. The corresponding error code is <code>UPDATE_PENDING</code> .
In an environment with Cisco Nexus 1000V Distributed Virtual Switch, vSphere Lifecycle Manager displays different messages in different situations. For details, see Host Upgrade Compliance Messages When Cisco Nexus 1000V Is Present .	If Cisco's Virtual Ethernet Module (VEM) software is found on the host, the precheck script checks if the software is part of the upgrade as well, and that the VEM supports the same version of the Virtual Supervisor Module (VSM) as the existing version on the host. If the software is missing or is compatible with a different version of the VSM, the script returns a warning and the scan result indicates the version of the VEM software that was expected on the upgrade ISO, and the version, if any, that was found on the ISO. The corresponding error code is <code>DISTRIBUTED_VIRTUAL_SWITCH</code> .
The host uses an EMC PowerPath multipathing module <code>file_name</code> to access storage. The host will not be able to access such storage after upgrade.	The script checks for installation of EMC PowerPath software, consisting of a CIM module and a kernel module. If either of these components is found on the host, the script verifies that matching components (CIM, VMkernel module) also exist in the upgrade. If they do not, the script returns a warning that indicates which PowerPath components were expected on the upgrade ISO and which, if any, were found. The corresponding error code is <code>POWERPATH</code> .

Host Upgrade Compliance Messages When Cisco Nexus 1000V Is Present

When a host is managed by the Cisco Nexus 1000V virtual switch and you check the compliance of the host against an upgrade baseline, the scan messages provide information about problems with compliance between the VEM modules installed on the host and the modules available on the ESXi8.0 image.

vSphere Lifecycle Manager supports Cisco Nexus 1000V, a virtual access software switch that works with VMware vSphere and consists of two components.

Virtual Supervisor Module (VSM)

The control plane of the switch and a virtual machine that runs NX-OS.

Virtual Ethernet Module (VEM)

A virtual line card embedded in ESXi hosts.

vSphere Lifecycle Manager determines whether a host is managed by Cisco Nexus 1000V. vSphere Lifecycle Manager verifies whether the Cisco Nexus 1000V VEM VIBs in the ESXi upgrade image are compatible with the Cisco Nexus 1000V VSM that manages the host.

By using vSphere ESXi Image Builder, you can create custom ESXi images, which contain third-party VIBs that are required for a successful remediation operation.

Table 124: Compliance Check Results for the Cisco Nexus 1000V Network Switch

Compliance Check Message	Description
The upgrade does not contain any Cisco Nexus 1000V software package that is compatible with the Cisco Nexus 1000V software package on the host. Upgrading the host will remove the feature from the host.	A VEM VIB is not available on the ESXi8.0 upgrade image.

Compliance Check Message	Description
The host is currently added to a Cisco Nexus 1000V virtual network switch. The upgrade contains a Cisco Nexus 1000V software package <i>VIB_name</i> that is incompatible with the Cisco Nexus 1000V VSM. Upgrading the host will remove the feature from the host.	The VEM VIB on the ESXi8.0 upgrade image is not compatible with the version of the VSM.
The host is currently added to a Cisco Nexus 1000V virtual network switch. The upgrade does not contain any Cisco Nexus 1000V software package that is compatible with the Cisco Nexus 1000V VSM. Upgrading the host will remove the feature from the host.	The host and the image do not contain VEM VIBs, but the host is still listed in vCenter Server as managed by Cisco Nexus 1000V.
Cannot determine whether the upgrade breaks Cisco Nexus 1000V virtual network switch feature on the host. If the host does not have the feature, you can ignore this warning.	There was a problem with determining compatibility between the VEM VIB on the ESXi8.0 upgrade image and the VSM. Check whether the version of the VSM managing the host is certified as being compatible with vCenter Server8.0 and ESXi8.0.

Viewing Compliance Information About ESXi Hosts and Updates

vSphere Lifecycle Manager scans objects to determine how they comply with the baselines and baseline groups that you attach to those objects. You can view compliance information about a single ESXi host or a group of hosts in a container object.

Supported groups of ESXi hosts include virtual infrastructure container objects such as folders, clusters, and data centers.

The host or the container object must have an attached baseline or baseline group to be examined for compliance information. Compliance with baselines and baseline groups is assessed at the time of viewing.

The overall compliance status of an ESXi hosts depends on the compliance statuses of all baselines and baseline groups that are attached to the object. For information about the different compliance statuses that an object, a baseline or a baseline group might have, see [Compliance Statuses of ESXi Hosts, Baselines, and Baseline Groups](#).

The compliance status of a baseline depends on the compliance statuses of all updates in the baseline. For information about the compliance statuses that updates might have, see [Compliance Statuses of Updates](#).

The ability to view the compliance status of vSphere objects depends on the privileges that you have. To view the compliance status of an inventory object, you must have the **View Compliance Status** privilege. Users that have privileges to remediate against patches, extensions, and upgrades and to stage patches and extensions on a particular inventory object, can view the compliance status of the same object even if they do not have the **View Compliance Status** privilege.

- Users with the privilege to view a container, but not all the contents of the container, can view the aggregate compliance status of all objects in the container.
- If a user does not have the permission to view an object, its contents, or a particular virtual machine, the results of those scans are not displayed.

For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation.

For a list of all vSphere Lifecycle Manager privileges and their descriptions, see [vSphere Lifecycle Manager Privileges For Using Baselines](#).

For information about checking the compliance of hosts against an image, see [Check the Compliance of a Cluster Against a vSphere Lifecycle Manager Image](#).

The vSphere Lifecycle Manager Compliance View

Compliance information about inventory objects is displayed on the **Updates** tab for the object, in the so-called vSphere Lifecycle Manager compliance view.

The vSphere Lifecycle Manager compliance view for objects that you manage with baselines and baseline groups consists of three panes.

Table 125: vSphere Lifecycle Manager Compliance View

Pane	Description
Baselines	The Baselines pane has three information panels. • Host Information panel For individual ESXi hosts, this panel shows information about the ESXi version installed on the host. You can also view all updates that are installed on the host. For container objects, this panel shows information about the ESXi versions of all hosts in the container object. • Compliance Information panel For individual ESXi hosts, this panel shows the overall compliance status of the host against all attached baselines and baseline groups. You can also view compliance information about the baselines and baseline groups attached to the host. For container objects, this panel shows the overall number of compliant and non-compliant hosts. This panel also shows the last time a compliance check was completed. • Remediation Information panel This panel shows the result of the remediation pre-check and indicates whether the selected object is ready for remediation. The panel also contains information about the issues that require user attention or action. This panel also shows the last time a remediation pre-check was completed. The information in the Baselines pane changes dynamically depending on the inventory object, baselines, and baseline groups that you select.
Attached Baselines and Baseline Groups	Displays the baselines and baseline groups attached to the selected object.
Bottom pane	The bottom pane appears when you select a baseline or a baseline group from the Attached Baselines and Baseline Groups pane. The information in this pane depends on the type of inventory object that you select. For individual hosts, the bottom pane shows information about all updates in the baseline or baseline group that you select from the Attached Baselines and Baseline Groups pane. • If you select a patch baseline or extension baseline, the bottom pane shows a list of all updates that the selected baseline contains. • If you select an upgrade baseline, the bottom pane shows information about the ESXi image that the upgrade baseline contains. • If you select a baseline group, the bottom pane shows all baselines included in the group along with their compliance statuses. You can also view all the updates that the baseline group contains. If the baseline group contains an ESXi image, information about it is also displayed in the bottom pane.

Pane	Description
	For container objects, the bottom pane displays compliance information about the ESXi hosts that have the selected baseline or baseline group. When you select a baseline or a baseline group from the Attached Baselines and Baseline Groups pane, the bottom pane appears and displays all the hosts to which the selected baseline or baseline group is attached along with the individual compliance statuses of those hosts.

Compliance Statuses of Updates

When you work with vSphere Lifecycle Manager baselines, an update stands for all patches, extensions, and upgrades that you can apply with vSphere Lifecycle Manager baselines. The compliance status of the updates in the baselines and baseline groups that you attach to objects in your inventory is calculated after you check the compliance of the target object.

The compliance statuses of the updates in a baseline define the overall compliance status of that baseline. For more information about baseline compliance statuses, see [Compliance Statuses of ESXi Hosts, Baselines, and Baseline Groups](#).

Conflict	The update conflicts with either an existing update on the host or another update in the vSphere Lifecycle Manager depot. vSphere Lifecycle Manager reports the type of conflict. A conflict does not indicate any problem on the target object. It just means that the current baseline selection is in conflict. You can perform compliance checks, remediation, and staging operations. In most cases, you must resolve the conflict.
Conflicting New Module	The host update is a new module that provides software for the first time, but it is in conflict with either an existing update on the host or another update in the vSphere Lifecycle Manager depot. vSphere Lifecycle Manager reports the type of conflict. A conflict does not indicate any problem on the target object. It just means that the current baseline selection is in conflict. You can perform scan, remediation, and staging operations. In most cases, you must resolve the conflict.
Incompatible Hardware	The hardware of the selected object is incompatible or has insufficient resources to support the update. For example, when you perform a host upgrade scan against a 32-bit host or if a host has insufficient RAM.
Installed	The update is installed on the target object and no further user action is required.
Missing	The update is applicable to the target object, but it is not yet installed. You must perform a remediation on the target object with this update, so that the update becomes compliant.
Missing Package	The metadata for the update is in the depot, but the corresponding binary payload is missing. The reasons can be that the product might not have an update for a given locale; the vSphere Lifecycle Manager depot is corrupt, and vSphere Lifecycle Manager no longer has Internet access to download updates; or you have manually deleted an upgrade package from the vSphere Lifecycle Manager depot.
New Module	The update is a new module. An update with this compliance status cannot be installed when it is part of a host patch baseline. When it is part of a host extension baseline, the new module

	status indicates that the module is missing on the host and can be provisioned by remediation. The compliance status of the baseline depends on the type of baseline containing the update with the New Module status. If the baseline is a host patch baseline, the overall status of the baseline is compliant. If the baseline is a host extension baseline, the overall status of the baseline is non-compliant.
Not Applicable	The update is not applicable to the target object. A patch might have the not applicable compliance status for one of the following reasons: • There are other patches in the vSphere Lifecycle Manager depot that obsolete this patch. • The update does not apply to the target object.
Not Installable	The update cannot be installed. The compliance check might succeed, but the remediation of the target object cannot be performed.
Obsoleted By Host	This compliance status is mainly applicable to patches. The target object has a newer version of the patch. For example, if a patch has multiple versions, after you apply the latest version to the host, the earlier versions of the patch have the Obsoleted By Host compliance status.
Staged	This compliance status applies to host patches and host extensions. It indicates that the update is copied from the vSphere Lifecycle Manager depot to the host, but it is not yet installed. Staged compliance status might occur only when you check the compliance status of hosts running ESXi6.7 and later.
Unknown	A patch is in the unknown state for a target object until vSphere Lifecycle Manager successfully scans the object. A scan might not succeed if the target object is of an unsupported version, if vSphere Lifecycle Manager lacks metadata, or if the patch metadata is corrupt.
Unsupported Upgrade	The upgrade path is not supported. For example, the current hardware version of the virtual machine is later than the latest version supported by the host.

Compliance Statuses of ESXi Hosts, Baselines, and Baseline Groups

Compliance statuses are computed after you initiate a compliance check for an inventory object against the attached baselines or baseline groups. The compliance statuses that baselines and baseline groups might have depend on the applicability of the patches, extensions, and upgrades contained in the baseline or baseline group attached to an object. The compliance status of a single host depends on the compliance statuses of all baselines and baseline groups attached to the host.

Compliant

The compliant status indicates that a vSphere object is compliant with all baselines in an attached baseline group or with all patches, extensions, and upgrades in an attached baseline. The compliant state requires no further action. If a baseline contains patches or upgrades that are not relevant to the target object, the individual updates, and baselines or baseline groups that contain them, are treated as not applicable, and represented as compliant. Compliant are also hosts with attached patch baselines containing extensions or patches with the obsoleted by host status.

The compliant status occurs under the following conditions:

- Target objects are compliant with the baselines and baseline groups when all updates in the baseline or baseline group are either installed on the target object, obsoleted by host, or are not applicable to the target object.
- The updates in a baseline are compliant when they are installed on the target object, or are not applicable to the object.

Non-Compliant

The non-compliant status indicates that one or more baselines in a baseline group, or one or more patches, extensions, or upgrades in a baseline are applicable to the target object, but are not installed (missing) on the target. You must remediate the target object to make it compliant.

When a baseline contains a non-compliant update, the overall status of the baseline is non-compliant. When a baseline group contains a non-compliant baseline, the overall status of the baseline group is non-compliant. The non-compliant status takes precedence over the incompatible, unknown, and compliant states.

Unknown

When you attach a baseline or a baseline group to a vSphere object and you do not initiate a compliance check for the object, the status of the vSphere object against the baseline or baseline group is unknown. This status indicates that a compliance check is required, that the compliance check has failed, or that you initiated a compliance check on an unsupported platform.

When a baseline contains updates in the compliant and unknown states, the overall status of the baseline is unknown. When a baseline group contains unknown baselines and compliant baselines, the overall status of the baseline group is unknown. The unknown compliance status takes precedence over the compliant status.

Incompatible

The incompatible status requires attention and further action. You must determine the reason for incompatibility by probing further. You can remediate the objects that have this status, but the operation might not be successful. In most cases, vSphere Lifecycle Manager provides sufficient details for the incompatibility.

When a baseline contains updates in the incompatible, compliant, and unknown states, the overall status of the baseline is incompatible. When a baseline group contains incompatible, unknown, and compliant baselines, the overall status of the baseline group is incompatible. The incompatible compliance status takes precedence over the compliant and unknown compliance statuses.

View Compliance Information About ESXi Hosts

You can check how the ESXi hosts in your environment comply with the baselines and baseline groups that you attach to those hosts. You can check and view the compliance status of an individual host or a container object.

- Verify that the host for which you want to view compliance information uses baselines and is not managed with a single vSphere Lifecycle Manager image.
- Review the [The vSphere Lifecycle Manager Compliance View](#) topic.

Compliance checks provide information about the degree of compliance of an object with the attached baselines and baseline groups.

In the compliance view for an object, you can view information about the compliance of the object with the attached baselines and baseline groups. You can also view the individual compliance statuses of the attached baselines and baseline groups. The compliance view changes dynamically and depends on the object that you want to view compliance information about. For a full description of the compliance information that you can obtain about an object, see [The vSphere Lifecycle Manager Compliance View](#).

For information about the different compliance statuses that an object might have, see [Compliance Statuses of ESXi Hosts, Baselines, and Baseline Groups](#).

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Hosts > Baselines** and review the compliance information in the compliance view.

You see complete compliance information about the selected object.

Stage Patches and Extensions to ESXi Hosts

Staging is the process during which vSphere Lifecycle Manager downloads patches and extensions on the ESXi hosts. During staging, the patches and extensions are not installed on the host. Staging reduces the time that the host spends in maintenance mode during remediation.

- Attach a patch or extension baseline or a baseline group containing patches and extensions to the host.
- Required privileges: **VMware vSphere Lifecycle Manager > Manage Patches and Upgrades > Stage Patches and Extensions**.

To stage patches or extensions to hosts, first attach a patch or extension baseline or a baseline group containing patches and extensions to the host. Staging patches and extensions does not require that hosts enter maintenance mode.

With the vSphere Client, you can stage a single baseline, multiple baselines, or baseline groups to a single host or a group of hosts in a container object.

Some limitations exist depending on the compliance status of the patches or extensions that you want to stage.

Patches cannot be staged if they are obsoleted by other patches in the baselines or baseline groups for the same stage operation. vSphere Lifecycle Manager stages only the patches that it can install in a subsequent remediation process, based on the current compliance status of the host. If a patch is obsoleted by patches in the same selected patch set, the obsoleted patch is not staged.

If a patch is in conflict with the patches in the vSphere Lifecycle Manager depot and is not in conflict with a host, after a compliance check, vSphere Lifecycle Manager reports this patch as a conflicting one. You can still stage the patch to the host and after the stage operation, vSphere Lifecycle Manager reports this patch as staged.

During the stage operation, vSphere Lifecycle Manager performs pre-scan and post-scan operations and updates the compliance status of the baseline.

For more information about the different compliance statuses that an update might have, see [Compliance Statuses of Updates](#).

After you stage patches or extensions to hosts, you must remediate the hosts against all staged patches or extensions.

After remediation finishes, the host deletes all staged patches or extensions from its cache regardless of whether they were applied during the remediation. The compliance status of the patches or extensions that were staged but not applied to the hosts reverts from Staged to its previous value.

Important: Staging patches and extensions is supported for hosts that are running ESXi6.7 and later. You can stage patches to PXE booted ESXi hosts, but if the host is restarted before remediation, the staged patches are lost and you must stage them again.

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Hosts > Baselines**.
3. In the **Attached Baselines** pane, select one or more baselines.
4. Click **Stage**.
The **Stage Patches** dialog box opens.
5. Select hosts on which to stage patches and extensions.
The number of selected hosts is on the top of the list.
6. To view the patches or extensions that will download to the selected hosts, expand the **Stage** list.
7. Click **Stage**.

The staging operation starts. You can monitor the progress of the task in the **Recent Tasks** pane.

Remediate the host or hosts.

After remediation, all staged patches and extensions, whether installed or not during the remediation, are deleted from the host.

Remediating ESXi Hosts Against vSphere Lifecycle Manager Baselines and Baseline Groups

Remediation is the process during which vSphere Lifecycle Manager applies patches, extensions, and upgrades to ESXi hosts. Remediation makes the selected vSphere objects compliant with the attached baselines and baseline groups.

General Considerations

- vSphere Lifecycle Manager supports the remediation of ESXi hosts against patch, extension, and upgrade baselines.
- You can initiate remediation manually or schedule a regular remediation task to run at a time that is convenient for you.
- You can remediate a single ESXi host or multiple hosts in a container object. You can initiate remediation at a folder, a cluster, a data center, and even vCenter Server level.

Note: If you initiate remediation against a baseline for an object that contains clusters that use a single vSphere Lifecycle Manager image, remediation is not performed on those clusters.

- By default, the remediation process runs sequentially. That is, vSphere Lifecycle Manager remediates the hosts in a cluster or another container object one by one. However, you can configure vSphere Lifecycle Manager to remediate multiple hosts in parallel.
- If a vCenter Server instance is connected to other vCenter Server systems by a common vCenter Single Sign-On domain, you can remediate only the inventory objects managed by the vCenter Server system where the vSphere Lifecycle Manager instance that you use runs.
- To remediate vSphere objects against baselines or baseline groups, you must have the **Remediate to Apply Patches, Extensions, and Upgrades** privilege. For more information about managing users, groups, roles, and permissions, see the *vSphere Security* documentation.

For a list of all vSphere Lifecycle Manager privileges and their descriptions, see [vSphere Lifecycle Manager Privileges For Using Baselines](#).

- If a vCenter HA failover is initiated during the remediation of a cluster, the remediation task is canceled. After the failover finishes, you must restart the remediation task on the new node.

Orchestrated Upgrades of ESXi Hosts

You can use baseline groups to perform an orchestrated upgrade of the ESXi hosts in your environment. The baseline group might contain a single host upgrade baseline and multiple patch or extension baselines, or multiple patch and extension baselines. vSphere Lifecycle Manager first upgrades the hosts and then applies the patch or extension baselines. Because the upgrade runs first and patches are applicable to a specific host version, the orchestrated workflow ensures that no patches are lost during the upgrade.

Orchestrated upgrades can be performed at a host, cluster, folder, or a data center level.

Instead of creating a baseline group, you can select and work with multiple baselines instead of grouping them into a baseline group first.

Maintenance Mode

If the update requires it, vSphere Lifecycle Manager puts hosts into maintenance mode during remediation. Virtual machines cannot run when a host is in maintenance mode. To ensure a consistent user experience, vCenter Server migrates the virtual machines to other hosts within the cluster before the host is put in maintenance mode. vCenter Server can migrate the virtual machines if the cluster is configured for vMotion and if VMware Distributed Resource Scheduler (DRS) and VMware Enhanced vMotion Compatibility (EVC) are enabled. However, EVC is not a prerequisite for vMotion. EVC guarantees that the CPUs of the hosts are compatible. For container objects or individual hosts that are not in a cluster, migration with vMotion cannot be performed. After remediation, hosts exit maintenance mode. In case of failure during remediation, hosts might be unable to exit maintenance mode.

Parallel Remediation

You can enable vSphere Lifecycle Manager to remediate in parallel the hosts within a cluster that uses baselines. Parallel remediation reduces the time needed for patching or upgrading the hosts in your environment. You can remediate in parallel only ESXi hosts that are already in maintenance mode. During parallel remediation, hosts do not enter maintenance mode automatically. Similarly, after remediation finishes, the hosts do not exit maintenance mode automatically. To remediate hosts in parallel, you must manually enter and exit maintenance mode. If you enable parallel remediation, vSphere Lifecycle Manager does not remediate the ESXi hosts that are not in maintenance mode.

When you configure vSphere Lifecycle Manager to remediate hosts in parallel, you can set the maximum number of hosts to be remediated in a single remediation task. Alternatively, you can let vSphere Lifecycle Manager to remediate all hosts in maintenance mode in parallel.

When you remediate hosts in parallel, if the remediation of a single host fails, the remediation task for the entire cluster does not stop and the rest of the hosts are remediated successfully. After remediation finishes, vSphere Lifecycle Manager reports an error for the respective host.

Parallel remediation is deactivated by default, but you can enable it during remediation or in the vSphere Lifecycle Manager general remediation settings.

You cannot remediate in parallel hosts in a vSAN cluster.

Remediation Pre-Check

Before you remediate an object, you can perform a remediation pre-check on the object. During that check, vSphere Lifecycle Manager identifies possible issues that might prevent successful remediation and takes or suggests actions to fix the issues.

For more information about the possible issues that might prevent successful remediation, see [Remediation Pre-Check Report](#).

Remediation of PXE Booted ESXi Hosts

vSphere Lifecycle Manager lets you remediate PXE booted ESXi hosts. vSphere Lifecycle Manager does not apply the patches that require a reboot to PXE booted ESXi hosts.

If there is any additional software installed on the PXE booted ESXi host, the software might be lost if the host restarts. Update your image profile with the additional software so that it will be present after the reboot.

To patch PXE booted ESXi hosts, you must enable the respective setting in the **Edit Settings for Host Remediation** dialog box, which you open from the **Settings** tab in the vSphere Lifecycle Manager home view.

Understanding the Remediation Operation

For ESXi hosts, updates are all-inclusive. The most recent update contains the patches from all previous releases.

The ESXi image on the host maintains two copies. The first copy is in the active boot and the second one is in the standby boot. When you patch an ESXi host, vSphere Lifecycle Manager creates an image based on the content of the active boot and the content of the patch. The new ESXi image is then located in the standby boot and vSphere Lifecycle Manager designates the active boot as the standby boot and reboots the host. When the ESXi host reboots, the active boot contains the patched image and the standby boot contains the previous version of the ESXi host image.

When you upgrade an ESXi host, vSphere Lifecycle Manager replaces the backup image of the host with the new image and replaces the active boot and the standby boot. During the upgrade, the layout of the disk that hosts the boots changes. The total disk space for an ESXi host remains 1GB, but the disk partition layout within that 1GB disk space changes to accommodate the new size of the boots where the ESXi8.0 image is stored.

For rollback purposes, the term update refers to all ESXi patches, updates, and upgrades. Each time you update an ESXi host, a copy of the previous ESXi build is saved on your host.

If an update fails and the ESXi8.0 host cannot boot from the new build, the host reverts to booting from the original boot build. ESXi permits only one level of rollback. Only one previous build can be saved at a time. In effect, each ESXi8.0 host stores up to two builds, one boot build and one standby build.

The remediation of ESXi6.7 and 7.0 hosts to their respective ESXi update releases is a patching process, while the remediation of ESXi hosts from version 6.7 or 7.0 to 8.0 is an upgrade process.

From the vSphere Lifecycle Manager settings, you can configure the host remediation process to skip a host reboot during host patch and host upgrade operations. This configuration setting is called Quick Boot. For more information about configuring the vSphere Lifecycle Manager remediation settings, see [vSphere Lifecycle Manager Remediation Settings](#).

Types of Host Remediation

Host remediation runs in different ways depending on the types of baselines that you attach to an object and whether the remediated host is in a cluster or not.

Host Upgrade Remediation

When you upgrade an ESXi6.7 or ESXi7.0 host to ESXi8.0, all supported custom VIBs remain intact on the host after the upgrade, regardless of whether the VIBs are included in the installer ISO.

You can upgrade hosts by using custom ESXi images that contain third-party modules for ESXi8.0. In such a case, third-party modules that are compatible with ESXi8.0 stay available on the upgraded host.

Host upgrade in a high-latency network in which vSphere Lifecycle Manager and the hosts are at different locations might take a few hours because the upgrade file is copied from the vSphere Lifecycle Manager depot to the host before the upgrade. During this time, the host stays in maintenance mode.

vSphere Lifecycle Manager 8.0 supports upgrade from ESXi6.7 and ESXi7.0 to ESXi8.0.

Upgrading to ESXi8.0 requires a boot device that is a minimum of 4 GB. When booting from a local disk, SAN or iSCSI LUN, up to 128 GB of disk space is used to create ESXi system partitions. You can create a VMFS datastore on a boot disk larger than 128 GB.

Note: After you upgrade your host to ESXi8.0, you cannot roll back to the previous ESXi versions, ESXi6.7, ESXi7.0. So, back up your host configuration before performing an upgrade. If the upgrade fails, you can reinstall the ESXi6.7 or ESXi7.0 software that you upgraded from and restore your host configuration. For more information about backing up and restoring your ESXi configuration, see the *VMware ESXi Upgrade* documentation. To upgrade ESXi hosts, you must first import ESXi ISO images to the vSphere Lifecycle Manager depot. You then create baselines and baseline groups to manage the upgrades for the ESXi hosts.

Host Patch Remediation

Patching is the process of remediating ESXi hosts against patch baselines.

The remediation of ESXi6.7 and 7.0 hosts to their respective ESXi update releases is a patching process, while the remediation of ESXi hosts from version 6.7 or 7.0 to 8.0 is an upgrade process.

vSphere Lifecycle Manager handles host patches in the following ways:

- If a patch in a patch baseline requires the installation of another patch, vSphere Lifecycle Manager detects the prerequisite in the depot and installs it together with the selected patch.
- If a patch is in a conflict with other patches that are installed on the host, the conflicting patch might not be staged or installed. However, if another patch in the baseline resolves the conflicts, the conflicting patch is installed. For example, consider a baseline that contains patch A and patch C, and patch A conflicts with patch B, which is already installed on the host. If patch C obsoletes patch B, and patch C is not in a conflict with patch A, the remediation process installs patches A and C.
- If a patch is in a conflict with the patches in the vSphere Lifecycle Manager depot and is not in a conflict with the host, after a compliance check, vSphere Lifecycle Manager reports this patch as a conflicting one. You can stage and apply the patch to the host.
- When multiple versions of the same patch are selected, vSphere Lifecycle Manager installs the latest version and skips installing the earlier versions.

During patch remediation, vSphere Lifecycle Manager automatically installs the prerequisites of the patches.

With vSphere Lifecycle Manager 8.0, you can remediate hosts of version ESXi6.7 and ESXi7.0 against patches from offline bundles, which you import to the vSphere Lifecycle Manager depot manually.

You can stage patches before remediation to reduce host downtime.

Host Extension Remediation

During extension remediation, vSphere Lifecycle Manager does not automatically install the prerequisites of the extension. Missing extension prerequisites cause some remediation operations to fail. If the missing prerequisite is a patch, you can add it to a patch baseline. If the missing prerequisite is an extension, you can add it to the same or another extension baseline. You can then remediate the host against the baseline or baselines that contain the prerequisite and the original extension baseline.

Remediating Hosts in a Cluster

For ESXi hosts in a cluster, the remediation process is sequential by default.

When you remediate a cluster of hosts sequentially and one of the hosts fails to enter maintenance mode, vSphere Lifecycle Manager reports an error and the remediation process stops and fails. The hosts in the cluster that are remediated stay at the updated level. The ones that are not remediated after one host fails remain unupdated.

The host upgrade remediation of ESXi hosts in a cluster proceeds only if all hosts in the cluster can be upgraded.

If you initiate remediation at a data center level, the remediation processes for the clusters run in parallel. Clusters that you manage with a single vSphere Lifecycle Manager image are not remediated against the attached baselines or baseline groups. If the remediation process fails for one of the clusters within a data center, the remaining clusters are still remediated.

Before you start remediation, you can generate a report that shows which cluster, host, or virtual machine has the cluster features enabled. For more information, see [Remediation Pre-Check Report](#).

Remediation and Cluster Settings

If a host in a DRS-enabled cluster runs a virtual machine on which vCenter Server is installed, DRS first attempts to migrate the virtual machine running vCenter Server to another host, so that the remediation succeeds. If the virtual machine cannot be migrated to another host, the remediation fails for the host, but the remediation process for the cluster does not stop. vSphere Lifecycle Manager proceeds to remediate the next host in the cluster.

Remediation of hosts in a cluster requires that you temporarily deactivate cluster features such as VMware DPM and HA admission control. Also, you must turn off Fault Tolerance if it is enabled on any of the virtual machines on a host, and disconnect the removable devices connected to the virtual machines on a host, so that they can be migrated with vMotion. For more information about configuring the vSphere Lifecycle Manager remediation settings, see [vSphere Lifecycle Manager Remediation Settings](#).

If a vCenter HA failover is initiated during the remediation of a cluster, the remediation task is canceled. After the failover finishes, you must restart the remediation task on the new node.

When you perform remediation on a cluster that consists of not more than two hosts, deactivating HA admission control might not be enough to ensure successful remediation. You might need to deactivate vSphere High Availability (HA) for the cluster. If you keep HA enabled, the remediation attempts on hosts in the cluster fail, because HA cannot provide recommendation to vSphere Lifecycle Manager to place any of the hosts into maintenance mode. The reason is that if one of the two hosts is placed into maintenance mode there is no failover host left available in the cluster. To ensure successful remediation on a two-node cluster, you must deactivate HA for the cluster or place the hosts in maintenance mode manually and then remediate the two hosts in the cluster.

vSAN Clusters

vSphere Lifecycle Manager remediates hosts that are part of a vSAN cluster sequentially. The reason is that by design only one host from a vSAN cluster can be in a maintenance mode at any time. For more information about using vSphere Lifecycle Manager with vSAN clusters, see [vSAN Clusters and vSphere Lifecycle Manager](#).

Remediating Hosts That Contain Third-Party Software

Hosts might contain third-party software, such as Cisco Nexus 1000V VEMs or EMC PowerPath modules. When you upgrade an ESXi6.7 or ESXi7.0 host to ESXi8.0, all supported custom VIBs are migrated and remain intact, regardless of whether the VIBs are included in the installer ISO.

If the host or the installer ISO image contains a VIB that creates a conflict and prevents the upgrade, an error message identifies the VIB that creates the conflict.

To discover potential problems with third-party software before an upgrade operation, scan the hosts against an upgrade baseline and review the scan messages in the vSphere Lifecycle Manager compliance view. See [Host Upgrade Compliance Messages](#) and [Host Upgrade Compliance Messages When Cisco Nexus 1000V Is Present](#).

For information about upgrading with third-party customization, see the *VMware ESXi Upgrade* documentation.

For information about using vSphere ESXi Image Builder to make a custom ISO, see the *VMware ESXi Installation and Setup* documentation.

Remediating ESXi6.7 or ESXi7.0 Hosts Against an ESXi8.0 Image

When you upgrade an ESXi6.7 or ESXi7.0 host to ESXi8.0, all supported custom VIBs remain intact on the host after the upgrade, regardless of whether the VIBs are included in the installer ISO.

When you perform a compliance check, the target host is scanned against a set of VIBs from the upgrade image. If you check the compliance of a host against an upgrade baseline that contains an ISO image of the same version as the target host, vSphere Lifecycle Manager displays Compliant or Non-compliant compliance status. If the upgrade image is the basic one distributed by VMware, or is a custom ISO image that contains the same set of VIBs as the ones already installed on the target host, the scan result is Compliant. If the upgrade ISO contains VIBs that are of different kind or version than the VIBs that are already on the target host, the scan result is Non-compliant.

The remediation process of an ESXi6.7 or ESXi7.0 host against an ESXi8.0 image is an upgrade process.

Note: Upgrading to ESXi8.0 requires a boot device that is a minimum of 4 GB. When booting from a local disk, SAN or iSCSI LUN, up to 128 GB of disk space is used to create ESXi system partitions. You can create a VMFS datastore on a boot disk larger than 128 GB.

You can use an ISO 8.0 image in an upgrade operation of an ESXi8.0 host. The remediation process of ESXi8.0 host by using ESXi8.0 image with additional VIBs is equivalent to a patching process. Because the upgrade image is of the same version as the target host, upon completing the upgrade operation, the additional VIBs are added to the target host.

Table 126: Scan and Remediation Situations for ESXi6.7 and ESXi7.0 Hosts Against ESXi8.0 Images

Action	Description
Compliance check and remediation of ESXi6.7 or ESXi7.0 hosts against an ESXi8.0 image that contains additional non-conflicting and non-obsolete VIBs with the target host.	vSphere Lifecycle Manager displays a Non-Compliant compliance status for the host. Remediation succeeds. All VIBs on the target host before remediation remain on the host. All VIBs from the upgrade image that are not present on the target host before remediation are added to the host.
Compliance check and remediation of ESXi6.7 or ESXi7.0 hosts against an ESXi8.0 image that contains VIBs of a version later than the version of the same VIBs on the target host.	vSphere Lifecycle Manager displays a Non-Compliant compliance status for the host. Remediation succeeds. VIBs on the target host are updated to the later version.
Compliance check and remediation of ESXi6.7 or ESXi7.0 hosts against an ESXi8.0 image that contains conflicting VIBs with the target host.	vSphere Lifecycle Manager displays an Incompatible compliance status for the host. Remediation fails. The host remains intact.
Scan and remediation of ESXi6.7 or ESXi7.0 hosts against an ESXi8.0 image that contains vendor-tagged VIBs.	- If the vendor-tagged VIBs do not match the host hardware, vSphere Lifecycle Manager displays an Incompatible compliance status for the host. Remediation fails. - If the vendor-tagged VIBs match the host hardware, vSphere Lifecycle Manager displays a Non-Compliant compliance status for the host and remediation succeeds.
Scan and remediation of ESXi6.7 or ESXi7.0 hosts against an ESXi8.0 image that contains VIBs that obsolete the VIBs installed on the host.	Remediation succeeds. All VIBs that have been installed on the target host before remediation are replaced by the newer VIBs from the ESXi image.

Generate a Pre-Remediation Check Report

When you generate a pre-remediation check report, vSphere Lifecycle Manager generates a list with actions that you must perform to ensure successful remediation of the hosts in your cluster.

The remediation pre-check report contains information about issues at the cluster, host, and VM level that might prevent the completion of remediation.

For information about the possible issues that might prevent successful remediation, see [Remediation Pre-Check Report](#).

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Host > Baselines**.
3. In the **Baselines** pane, click **Pre-Check Remediation**.
The **Remediation Pre-check** dialog box opens.
4. Review the results from the pre-check and click **Done**.
In the bottom pane of the **Remediation Pre-check** dialog box, you see a list of issues at the host and virtual machine level.

The **Remediation Pre-check** dialog box lists the issues with cluster, hosts, and virtual machines that might prevent successful remediation of the selected object.

In the upper pane of the **Remediation Pre-check** dialog box, you see a list of issues at a cluster level.

In the bottom pane of the **Remediation Pre-check** dialog box, you see a list of issues at the host and virtual machine level.

Fix all issues that vSphere Lifecycle Manager identifies during the pre-remediation check and remediate the selected object.

Remediation Pre-Check Report

The remediation pre-check report shows the results of a check that is performed on a cluster or a host before remediation. During that check, vSphere Lifecycle Manager identifies possible issues that might prevent successful remediation. Depending on the type of issue, vSphere Lifecycle Manager suggests actions that you must take to fix the issue or resolves the issue automatically.

You can generate a pre-check remediation report in the vSphere Lifecycle Manager compliance view for an object.

Table 127: Cluster Issues

Current Configuration/Issue	Recommended Action	Details
DRS is deactivated on the cluster.	Enable DRS on the cluster.	DRS enables vCenter Server to place and migrate virtual machines automatically on hosts to attain the best use of cluster resources.
vSAN health check fails during the pre-check.	Navigate to the vSAN Health page and address any health issues before proceeding with remediation.	The vSAN health check performs a series of tests on the hosts in the vSAN cluster. The vSAN health check must succeed to ensure the hosts are successfully remediated. If you start a remediation task in a vSAN cluster that failed the vSAN health check during the remediation pre-check, the hosts enter maintenance mode, get upgraded, but might fail to exit maintenance mode. The remediation eventually fails.

Current Configuration/Issue	Recommended Action	Details
Insufficient licenses for one or multiple ESXi hosts in the cluster.	Ensure that you have multiple licenses for the ESXi hosts that have more than 32 cores per CPU.	One CPU license covers up to 32 physical cores. If a CPU has more than 32 cores, you must assign additional CPU licenses to the respective ESXi host. For more information, see https://www.vmware.com/company/news/updates/cpu-pricing-model-update-feb-2020.html .
DPM is enabled on the cluster.	None. vSphere Lifecycle Manager deactivates DPM automatically.	If a host has no running virtual machines, DPM might put the host in standby mode before or during remediation and vSphere Lifecycle Manager cannot remediate them.
HA admission control is enabled on the cluster.	None. vSphere Lifecycle Manager deactivates HA admission control automatically.	HA admission control prevents the migration of virtual machines with vSphere vMotion and the hosts cannot enter maintenance mode.
EVC is deactivated on the cluster.	None. vSphere Lifecycle Manager enables EVC automatically, although no notification or message appears in the vSphere Client.	If EVC is deactivated for a cluster, the migration of virtual machines with vSphere vMotion cannot proceed. The result is downtime of the machines on the hosts that you remediate with vSphere Lifecycle Manager.

Table 128: Host Issues

Current Configuration/Issue	Recommended Action	Details
A CD/DVD drive is attached to a virtual machine on the ESXi host.	Disconnect the CD/DVD drive.	Any CD/DVD drives or removable devices connected to the virtual machines on a host might prevent the host from entering maintenance mode. When you start a remediation operation, the hosts with virtual machines to which removable devices are connected are not remediated.
A floppy drive is attached to a virtual machine on the ESXi host.	Disconnect the floppy drive.	Any floppy drives or removable devices connected to the virtual machines on a host might prevent the host from entering maintenance mode. When you start a remediation operation, the hosts with virtual machines to which removable devices are connected are not remediated.
Fault Tolerance (FT) is enabled for a virtual machine on the ESXi host.	Deactivate FT for the virtual machine.	If FT is enabled for any of the virtual machines on a host, vSphere Lifecycle Manager cannot remediate that host.

Current Configuration/Issue	Recommended Action	Details
A powered on virtual machine is configured to use Virtual Flash Read Cache.	Deactivate Virtual Flash Read Cache before proceeding with the upgrade.	Virtual Flash Read Cache is not supported. During an upgrade operation, vSphere Lifecycle Manager removes Virtual Flash Read Cache for all virtual machines on the host. Before remediation, consult https://kb.vmware.com/s/article/2057840 .
VMware vCenter Server is installed on a virtual machine on the ESXi host and DRS is deactivated on the cluster.	Enable DRS on the cluster and ensure that virtual machines can be migrated with vSphere vMotion.	One of the virtual machines in the cluster runs the vCenter Server instance that you currently use. If you enable DRS on the cluster, vSphere vMotion can migrate the virtual machine where vCenter Server runs to ensure that the remediation of the hosts is successful.
An ESXi host in the cluster has a CPU with more than 32 cores and requires multiple licenses.	Assign as many licenses as the host needs.	One CPU license covers up to 32 physical cores. If a CPU has more than 32 cores, you must obtain additional CPU licenses. For more information, see https://www.vmware.com/company/news/updates/cpu-pricing-model-update-feb-2020.html .

Remediate ESXi Hosts Against a Single Baseline or Multiple Baselines

Remediation makes the remediated ESXi host compliant with the attached baselines and baseline groups. You can remediate a host against a single or multiple baselines, or against a baseline group. Baseline groups might contain multiple patch and extension baselines, or an upgrade baseline combined with multiple patch and extension baselines. You can remediate ESXi hosts against a single attached upgrade baseline at a time.

- Required privileges: **VMware vSphere Lifecycle Manager > Manage Patches and Upgrades > Remediate to Apply Patches, Extensions, and Upgrades**.
- Attach a patch, upgrade, or extension baseline or a baseline group containing patches, upgrades, and extensions to the host.
- Resolve any issues that occur during Remediation Pre-check.
- In upgrade scenarios, verify that the ESXi hosts to upgrade have a boot disk of at least 4 GB. When booting from a local disk, SAN or iSCSI LUN, up to 128 GB of disk space is used to create ESXi system partitions. You can create a VMFS datastore on a boot disk larger than 128 GB.
- To enable Quick Boot, verify that the ESXi host is compatible with the feature. For more information, see [Quick Boot](#).

You can upgrade all hosts in your vSphere inventory by using a single upgrade baseline that contains an ESXi image. You can remediate a single ESXi host or a group of ESXi hosts in a container object, such as a folder, a cluster, or a data center. You can also initiate remediation at a vCenter Server level.

Note: Because the official VMware online depot hosts certified partner content in addition to VMware content, a broader set of OEM bulletins are available in the vSphere Lifecycle Manager depot. As a result, a broader set of OEM bulletins are included in the vSphere Lifecycle Manager predefined bulletins. During remediation, always inspect the contents of those baselines to exclude the bulletins that you do not need in the baseline. For the bulletins that you do need, consult the

corresponding KB articles for information about deployment specifics and dependencies. Verify that dependent bulletins are also included in the baselines that you use for remediation.

1. In the vSphere Client, navigate to the vSphere Lifecycle Manager compliance view for an individual host or a container object.
 - a) Navigate to a host, cluster, or a container object.
 - b) Click the **Updates** tab.
2. Select **Hosts > Baselines**.
3. In the **Attached Baselines and Baseline Groups** pane, select the baselines and baseline groups to use for remediation.

You can select a single baseline or baseline group. You can also select multiple baselines and baseline groups. Your selection must contain no more than one upgrade baseline.
4. Click **Remediate**.

If the selected baselines and baseline groups do not contain an upgrade image, the **Remediate** dialog box opens.

If the selected baselines and baseline groups contain an upgrade image, the **End User License Agreement** dialog box opens.
5. To proceed to remediation, accept the terms and the license agreement in the **End User License Agreement** dialog box.

After you accept the agreement and click **OK** to close the dialog box, the **Remediate** dialog box opens.
6. Expand the list of pre-check issues and review the actions that vSphere Lifecycle Manager must perform to ensure successful remediation.
7. Optional: To generate a full remediation pre-check report, click **Show Full Remediation Pre-Check Report**.

If you select this option, the **Remediate** dialog box closes and vSphere Lifecycle Manager does not proceed with the remediation process. Instead, the **Remediation Pre-Check** dialog box opens. After you review the results from the remediation pre-check, you must initiate remediation again.
8. Expand the list of hosts to be remediated and deselect any host that you do not want to remediate.

The list contains all the hosts to which the selected baselines and baseline groups are attached. Even if you navigated to a single host before initiating remediation, the list might still display multiple hosts to be remediated. All hosts in the list are selected by default. Deselecting hosts from the list changes the overall number of hosts to be remediated.
9. Optional: To view information about the updates that will be installed during the remediation, expand the list of updates.

If the selection of baselines and baseline groups contains an upgrade baseline, information about the ESXi image is also displayed.
10. Optional: To schedule the remediation task for a later time, expand **Scheduling Options** and configure a scheduled remediation task.

By default, the remediation task starts immediately after closing the **Remediate** dialog box.

11. Expand **Remediation settings** and review and edit the remediation settings.

- To turn on or turn off Quick Boot, select or deselect the respective check box in the **Remediation settings** table.
- To allow or disallow health checks after remediation, select or deselect the respective check box in the **Remediation settings** table.
- To ignore warnings about unsupported hardware devices, select the respective check box in the **Remediation settings** table.
- To configure parallel remediation for the selected hosts, expand **Parallel remediation**, select the respective check box and configure the maximum number of concurrent remediations.

Note: vSphere Lifecycle Manager remediates in parallel only the ESXi hosts that are in maintenance mode. Hosts that are not in maintenance mode are not remediated. If you do not set the maximum number of concurrent remediations, vSphere Lifecycle Manager remediates all hosts that are in maintenance mode.

If the hosts have NSX virtual distributed switches that are ready to be migrated to vSphere Distributed Switches, you must set the maximum number of parallel remediations to no more than 4. In cases when host switch migration is needed, if more than 4 hosts are remediated in parallel, the remediation might fail, because the host switch migration takes more time than the time vSphere Lifecycle Manager needs to complete the parallel remediation.

- To change any other of the remediation settings, click the **Close Dialog And Go To Settings** link above the table.

If you select this option, the **Remediate** dialog box closes and vSphere Lifecycle Manager does not proceed with the remediation process. Instead, you are redirected to the **Baselines Remediation Settings** pane on the **Settings** tab of the vSphere Lifecycle Manager home view. To change any of the remediation settings, click the **Edit** button. Remediation does not resume automatically. After you make the desired changes, you must initiate remediation again.

12. Click **Remediate**.

Depending on the remediation schedule you configure, the remediation task starts immediately or runs later.

Using vSphere Lifecycle Manager to Migrate an NSX Virtual Distributed Switch to a vSphere Distributed Switch

Starting with vSphere 7.0, the vSphere Distributed Switch supports the NSX functionality. For clusters enabled with NSX, you can migrate the NSX-managed Virtual Distributed Switches on the hosts to vSphere Distributed Switches during an upgrade of a cluster against a vSphere Lifecycle Manager baseline group that contains an ESXi image and NSX VIBs.

Migrating your host switch to vSphere Distributed Switch 7.0 or later ensures optimal pNIC usage, and lets you manage the networking for your NSX hosts from vCenter Server.

During an upgrade remediation, vSphere Lifecycle Manager checks if an NSX virtual distributed switch is present on each of the hosts in the cluster and if it is ready to be migrated. To prepare the NSX virtual distributed switch for migration, you must run the Upgrade Readiness Tool before upgrading the cluster. If the NSX distributed switch on any of hosts in the cluster is not ready to be migrated to a vSphere Distributed Switch, you cannot proceed with the remediation process. In this case, you need to go to NSX Manager and run the Upgrade Readiness Tool.

Requirements

- ESXi 7.0 Update 2
- vCenter Server 7.0 Update 2
- NSX 3.1.1

As a good practice, contact VMware support to assess the impact of migrating to vSphere Distributed Switch 7.0 or later.

Workflow

1. In NSX Manager, use the Upgrade Readiness Tool to run the migration readiness pre-check, address any configuration issues, review the recommended topology, and apply the new topology.
For more information about the steps you need to perform in NSX Manager, see "Migrate Host Switch to vSphere Distributed Switch" in the *NSX Center Administration* documentation.
2. Upgrade vCenter Server to version 7.0 Update 2.
For more information about upgrading vCenter Server, see the *vSphere Upgrade* documentation.
3. Create a baseline group that contains an ESXi image version 7.0 Update 2 and the NSX kernel module for ESXi 7.0 .
 - a. Import an ESXi 7.0 Update 2 ISO image to the vSphere Lifecycle Manager depot.
For more information, see [Import an ISO Image to the vSphere Lifecycle Manager Depot](#).
 - b. Create an upgrade baseline that contains the imported ISO image.
For more information, see [Create, Edit, or Delete vSphere Lifecycle Manager Baselines and Baseline Groups](#) .
 - c. From customerconnect.vmware.com, download the NSX kernel module for ESXi 7.0.
 - d. Import the downloaded NSX bundle to the vSphere Lifecycle Manager depot.
For more information, see [Import Updates to the vSphere Lifecycle Manager Depot](#).
 - e. Create an extension baseline that contains the imported kernel module.
For more information, see [Create, Edit, or Delete vSphere Lifecycle Manager Baselines and Baseline Groups](#) .
 - f. Create a baseline group that contains the ESXi upgrade baseline and the extension baseline with NSX VIBs.
For more information, see [Create a Host Baseline Group](#).
4. Attach the baseline group to a cluster.
For more information, see [Attach Baselines and Baseline Groups to Objects](#).
5. Remediate the cluster against the attached baseline group.
During the remediation, vSphere Lifecycle Manager upgrades the ESXi version first and then migrates the host switch, if migration is needed.
For more information, see [Remediate ESXi Hosts Against a Single Baseline or Multiple Baselines](#).
6. If you use host profiles to configure the hosts in your environment, re-extract a new host profile from the reference host in the cluster.
After the upgrade, due to the migration of the host switch, the existing host profiles become invalidated and inapplicable.
For more information, see the *vSphere Host Profiles* documentation.

Upgrading Virtual Machines with vSphere Lifecycle Manager

You can use vSphere Lifecycle Manager to upgrade virtual machine hardware and the VMware Tools version of a virtual machine.

Whether you perform an upgrade of the virtual machine hardware version or the VMware Tools version, the upgrade is a multi-stage process.

1. You check the status of individual virtual machines or a container object.
vSphere Lifecycle Manager checks the status of a virtual machine against the latest virtual machine hardware version supported by the host on which the virtual machine runs. Similarly, vSphere Lifecycle Manager checks the status of the virtual machine against the latest VMware Tools version supported by the host on which the virtual machine runs.
For more information about checking virtual machine status, see [Checking the Status of Virtual Machines](#).
2. You review the status of the scanned virtual machines.

3. You upgrade the virtual machine to match the host where it resides.

With vSphere Lifecycle Manager, you can upgrade the virtual machine hardware version and the VMware Tools version that a virtual machine has. You can use vSphere Lifecycle Manager to upgrade the virtual machine hardware version to the latest hardware version, vmx-19, and to the latest VMware Tools version on the hosts.

For more information about upgrading virtual machines, see [Upgrading Virtual Machines](#).

Configure Virtual Machine Rollback Settings

By default, vSphere Lifecycle Manager takes snapshots of virtual machines before upgrading them. If the upgrade fails, you can use the snapshot to return a virtual machine to its state before the upgrade.

Required privileges: **VMware vSphere Lifecycle Manager > Configure**

You can configure vSphere Lifecycle Manager to keep snapshots for an indefinite or fixed period of time. Use the following guidelines when managing snapshots.

- Keeping snapshots indefinitely might consume a large amount of disk space and degrade virtual machine performance.
- Keeping no snapshots saves space, ensures best virtual machine performance, and might reduce the remediation time. However, keeping no snapshots limits the availability of a rollback.
- Keeping snapshots for a fixed period of time uses less disk space and offers a backup for a short time.

vSphere Lifecycle Manager does not take snapshots of fault tolerant virtual machines and virtual machines of virtual machine hardware version 3. If you decide to take snapshots of such virtual machines, the upgrade might fail.

If you configure vSphere Lifecycle Manager to automatically upgrade VMware Tools on power cycle for selected virtual machines, vSphere Lifecycle Manager does not take snapshots of the virtual machines before upgrading them and you cannot roll back.

1. Navigate to the vSphere Lifecycle Manager home view.
 - a) In the vSphere Client, select **Menu > Lifecycle Manager**.
 - b) Select a vCenter Server system from the **Lifecycle Manager** drop-down menu.

The drop-down menu is available only when multiple vCenter Server systems are connected by a common vCenter Single Sign-On domain. By selecting a vCenter Server system, you specify which vSphere Lifecycle Manager instance you want to administer.
2. On the **Settings** tab, select **Cluster lifecycle > VMs**.
3. Click the **Edit** button.

The **Edit Default Settings for VM Rollback** dialog box opens.
4. Select the **Take snapshot of VMs** check box.
5. Select the time period for keeping the snapshots.
 - Do not delete snapshots
 - Keep snapshots for a configurable, fixed period of time
6. Click **Save** to save your changes and close the **Edit Default Settings for VM Rollback** dialog box.

These settings become the default rollback option settings for virtual machines. You can specify different settings when you configure individual remediation tasks.

Checking the Status of Virtual Machines

You use vSphere Lifecycle Manager to check the status of virtual machines before you upgrade them. The status check shows if the virtual machine is up-to-date or can be upgraded.

With vSphere Lifecycle Manager, you can check the status of a single virtual machine or a group of virtual machines in a parent container object.

Supported groups of virtual machines or ESXi hosts include virtual infrastructure container objects such as folders, vApps, clusters, and data centers.

vSphere Lifecycle Manager checks the status of virtual machines in two aspects.

- You can use vSphere Lifecycle Manager to check the status of the virtual machines in respect with the VMware Tools version that they have installed.
The status check is performed against the latest VMware Tools version that the parent host supports.
- vSphere Lifecycle Manager checks the status of the virtual machines in respect with their VM hardware compatibility.
vSphere Lifecycle Manager compares the hardware compatibility of the virtual machines with the default VM hardware compatibility configured for the host.

Check the Status of an Individual Virtual Machine

You check the status of virtual machines to see whether the VMware Tools version they have is up to date and whether their hardware compatibility matches the default VM hardware compatibility for the host.

1. In the vSphere Client, navigate to a virtual machine.

2. On the **Updates** tab, click **Check Status**.

The **Scan entity** task appears in the **Recent Tasks** pane. After the task finishes, status information appears in the **VMware Tools** and **VM Hardware Compatibility** panels.

The virtual machines are scanned for VMware Tools and VM hardware compliance.

Check the Status of the Virtual Machines in a Container Object

You check the status of virtual machines to see whether the VMware Tools version they have is up-to-date and whether their hardware compatibility matches the default VM hardware compatibility for the host where they reside.

When you perform a status check for a container object, vSphere Lifecycle Manager checks the VMware Tools and VM Hardware Compatibility statuses for all child virtual machines. The larger the virtual infrastructure and the higher up in the object hierarchy you initiate the status check, the longer the task takes.

- In the vSphere Client, navigate to a virtual machine container object, such as a virtual machine folder, host, cluster, and so on.
- Click the **Updates** tab.
- Select your task.

Option	Action
Check the VMware Tools status of the virtual machines in the container object.	- Select Hosts > VMware Tools > . - Click Check Status. The information about the VMware Tools status appears in the Tools Status column in the table that lists all virtual machines in the selected container object. If the container object is a data center or a vCenter Server instance, you must first specify the cluster that you want to see results for.
Check the VM Hardware compatibility status of the virtual machines in the container object.	- Select Hosts > VM Hardware. - Click Check Status. The information about the VM Hardware Compatibility status appears in the Status column in the table that lists all virtual machines in the selected container object. For each virtual

Option	Action
	machine in the object, you can also see the VM hardware compatibility and the host compatibility. If the container object is a data center or a vCenter Server instance, you must first specify the cluster that you want to see results for.

The VMware Tools Status

Check the VMware Tools status for information whether the current version of VMware Tools is installed, supported, or whether upgrades are available.

Table 129: VMware Tools Status

VMware Tools Status	Description
Up to Date	VMware Tools is installed, supported, and the version is compliant.
	VMware Tools is installed, supported, and the version is newer than the version available on the ESXi host.
Upgrade Available	VMware Tools is installed, but the version is old.
	VMware Tools is installed and supported, but a newer version is available on the ESXi host.
Version Unsupported	VMware Tools is installed, but the version is old.
	VMware Tools is installed, but the version has a known issue and must be immediately upgraded.
	VMware Tools is installed, but the version is too new to work correctly with this virtual machine.
Not Installed	VMware Tools is not installed on this virtual machine.
Guest Managed	vSphere does not manage VMware Tools.
Unknown	The status of the virtual machine is not checked.

Upgrading Virtual Machines

With vSphere Lifecycle Manager, you can upgrade the VMware Tools version and the hardware version of a virtual machine. You can also upgrade multiple virtual machines simultaneously if the virtual machines are in a container object, such as a folder or vApp. You can also upgrade simultaneously all virtual machines that run on a host, in a cluster, or in a data center.

vSphere Lifecycle Manager supports upgrading powered on, suspended, and powered off virtual machines.

During the upgrade of VMware Tools, the virtual machines must be powered on. If a virtual machine is in the powered off or suspended state before remediation, vSphere Lifecycle Manager powers it on. After the upgrade completes, vSphere Lifecycle Manager restarts the machine and restores the original power state of the virtual machine.

During the virtual hardware upgrade, the virtual machines must be powered off. If a virtual machine is powered on, vSphere Lifecycle Manager powers the machine off, upgrades the virtual hardware, and then powers the virtual machine on.

You can also upgrade VMware Tools and the hardware version of a virtual machine template. A template is a copy of a virtual machine that you can use to create and provision new virtual machines.

You can set up automatic upgrades of VMware Tools on power cycle. For more information, see [Automatically Upgrade VMware Tools on Reboot](#).

You can configure vSphere Lifecycle Manager to take snapshots of virtual machines and to keep the snapshots indefinitely or for a specific period of time. By using snapshots, you can roll back a virtual machine to its previous state if upgrading the virtual machine with vSphere Lifecycle Manager fails. After the upgrade finishes, you can delete

the snapshots if you do not need them. For more information about configuring virtual machine rollback settings, see [Configure Virtual Machine Rollback Settings](#).

You can upgrade virtual machines immediately or schedule an upgrade operation to run at a convenient time.

If a host is connected to vCenter Server by using an IPv6 address, you cannot scan and remediate virtual machines that run on the host.

Upgrade the VM Hardware Compatibility of Virtual Machines

You can upgrade the hardware of virtual machines to the latest hardware version that the host supports. You can upgrade immediately or schedule an upgrade at a time that is convenient for you.

With vSphere Lifecycle Manager, you can upgrade the hardware compatibility version of a single virtual machine or multiple virtual machines simultaneously. Supported container objects for virtual machines in the vSphere inventory are folders, vApps, data centers.

1. In the vSphere Client, navigate to a single virtual machine or a container object.

You can also initiate upgrade at the level of any inventory object where virtual machines run. For example, you can start the upgrade operation at a host or cluster level.

2. Open the **Upgrade VM Hardware to Match Host** dialog box.

Inventory Object	Steps
Virtual Machine	1. Click the Updates tab. 2. In the VM Hardware Compatibility panel, click Upgrade to Match Host.
Container Object, Host, Cluster, Data Center, or vCenter Server Instance	1. Click the Updates tab. 2. Select Hosts > VM Hardware. 3. If the selected inventory object is a data center or a vCenter Server instance, select a cluster from the list. A list of all virtual machines in the cluster appears in the bottom pane. 4. Select the virtual machines to upgrade. 5. Click Upgrade to Match Host.

A list of the virtual machines selected for upgrading is visible in the **Upgrade VM Hardware to Match Host** dialog box.

3. Optional: To change the selection of the virtual machines to upgrade, select or deselect virtual machines from the list.
4. Optional: To schedule the upgrade for a specific date and time, expand **Scheduling Options** and configure the scheduled task.
 - a) Enter a name and, optionally, a description for the scheduled upgrade task.
 - b) Use the **Powered On VMs**, **Powered Off VMs**, and **Suspended VMs** drop-down menus to configure the upgrade to run immediately or at a specific date and time.
5. Optional: To configure the use of snapshots, expand **Rollback Options** and change the default settings.
 - a) To allow or disallow taking of snapshots of virtual machines before upgrading them, select or deselect the **Take snapshot of VMs** check box.
The option to take snapshots is selected by default.
 - b) Select a period for keeping the snapshots.
 - Keep the snapshots indefinitely.
 - Keep the snapshots for a fixed period.
 - c) Enter a snapshot name and, optionally, a description for the snapshot.
 - d) Include the virtual machine memory in the snapshot by selecting the respective check box.
6. Review your selections and click the **Upgrade to Match Host** button.

The hardware versions of the selected virtual machines are upgraded and the virtual machine status changes to Up to Date.

Upgrade the VMware Tools Version of Virtual Machines

You can upgrade the VMware Tools version of virtual machines to the latest version that the host supports. You can upgrade immediately or schedule an upgrade at a time that is convenient for you.

With vSphere Lifecycle Manager, you can upgrade the VMware Tools version of a single virtual machine or multiple virtual machines simultaneously. Supported container objects for virtual machines in the vSphere inventory are folders, vApps, data centers.

1. In the vSphere Client, navigate to a single virtual machine or a container object.
You can initiate the upgrade at the level of any inventory object where virtual machines run. For example, you can start the upgrade operation at a host or cluster level.
2. Open the **Upgrade VMware Tools to Match Host** dialog box.

Inventory Object	Steps
Virtual Machine	1. Click the Updates tab. 2. In the VMware Tools panel, click Upgrade to Match Host.
Container Object, Host, Cluster, Data Center, or Instance	1. Click the Updates tab. 2. Select Hosts > VMware Tools. 3. If the selected inventory object is a data center or a instance, select a cluster from the list. A list of all virtual machines in the cluster appears in the bottom pane. 4. In the VMs in Cluster pane, select the virtual machines to upgrade. 5. Click Upgrade to Match Host.

A list of the virtual machines selected for upgrading is visible in the **Upgrade VMware Tools to Match Host** dialog box.

3. Optional: To change the selection of the virtual machines to upgrade, select or deselect virtual machines from the list.
4. Optional: To schedule the upgrade for a specific date and time, expand **Scheduling Options** and configure the scheduled task.
 - a) Enter a name and, optionally, a description for the scheduled upgrade task.
 - b) Use the **Powered On VMs**, **Powered Off VMs**, and **Suspended VMs** drop-down menus to configure the upgrade to run immediately or at a specific date and time.
5. Optional: To configure the use of snapshots, expand **Rollback Options** and change the default settings.
 - a) To allow or disallow taking of snapshots of virtual machines before upgrading them, select or deselect the **Take snapshot of VMs** check box.
The option to take snapshots is selected by default.
 - b) Select a period for keeping the snapshots.
 - Keep the snapshots indefinitely.
 - Keep the snapshots for a fixed period.
 - c) Enter a snapshot name and, optionally, a description for the snapshot.
 - d) Include the virtual machine memory in the snapshot by selecting the respective check box.
6. Review your selections and click the **Upgrade to Match Host** button.

The VMware Tools version that runs on the selected virtual machines is upgraded and the VMware Tools status changes to Up to Date.

Automatically Upgrade VMware Tools on Reboot

You can automate the process of upgrading VMware Tools for the virtual machines in your inventory.

Verify that you have the **Manage Patches and Upgrades > Remediate to Apply Patches, Extensions, and Upgrades** privilege.

You can set up vSphere Lifecycle Manager to check the VMware Tools version of a virtual machine when the virtual machine is rebooted. If necessary, vSphere Lifecycle Manager upgrades VMware Tools to the latest version supported by the host on which the virtual machine runs.

Note:

When you perform a VMware Tools upgrade on power cycle, vSphere Lifecycle Manager does not take a snapshot of the virtual machine and you cannot roll back to the previous version of the virtual machine.

1. In the vSphere Client, navigate to a single virtual machine or an inventory object that contains virtual machines.
2. Configure vSphere Lifecycle Manager to automatically upgrade VMware Tools on reboot.

Inventory Object	Steps
Virtual machine	1. Click the Updates tab. 2. In the VMware Tools panel, click Turn On.
Container Object, Host, cluster, data center, or vCenter Server instance	1. Click the Updates tab. 2. Select Hosts > VMware Tools. 3. If the selected inventory object is a data center or a vCenter Server instance, select a cluster from the list. A list of all virtual machines in the cluster appears in the bottom pane. 4. In the VMs in cluster pane, select the virtual machines for which you want to enable the automatic upgrade of VMware Tools.

Inventory Object	Steps
	5. Click Set Auto Update and select On . The new status is visible in the Auto Update column.

The next time you power on or restart a virtual machine, vSphere Lifecycle Manager checks the version of VMware Tools installed on the virtual machines and performs an upgrade, if necessary.

vSphere Lifecycle Manager Images and Other VMware Products and Solutions

You can use a vSphere Lifecycle Manager image to manage a cluster if the cluster only contains solutions integrated to work with vSphere Lifecycle Manager. If an unintegrated solution is enabled on a cluster, you cannot use vSphere Lifecycle Manager images to manage that cluster, but you can still use baselines and baseline groups.

A solution is a VMware product that integrates with vCenter Server and adds some new functionality to the ESXi hosts in the inventory.

When you enable a solution for a cluster that uses a vSphere Lifecycle Manager image, the solution automatically uploads an offline bundle with components into the vSphere Lifecycle Manager depot and adds its component to all hosts in the cluster. You cannot control the lifecycle of solution components. For example, if you export the image, solution components are not part of the exported image.

Integrated Solutions

You can manage a cluster with a single image if the cluster has any of the following solutions enabled.

- vSphere High Availability
- vSAN

For more information about the integration between vSAN and vSphere Lifecycle Manager, see [vSAN Clusters and vSphere Lifecycle Manager](#) and the *Administering VMware vSAN* documentation.
- vSAN File Services
- vSphere Supervisor

For detailed information about the integration between vSphere Supervisor and vSphere Lifecycle Manager, see the documentation.
- NSX

For more information about the integration between NSX and vSphere Lifecycle Manager, see the *NSX Administration* documentation.
- VMware Cloud Foundation

vSphere Lifecycle Manager is available as an option in the VMware Cloud Foundation Workload Domain. For more information, see the *VMware Cloud Foundation Lifecycle Management* documentation.
- VMware vSphere Replication

For more information about the integration between vSphere Lifecycle Manager and vSphere Replication, see the *vSphere Replication* documentation.

You can also use baselines for clusters that have those solutions enabled.

Unintegrated Solutions

You cannot manage a cluster with a single image if the cluster has any of the following solutions enabled.

- NSX for vSphere®
- Dell EMC VxRail

You can use baselines and baseline groups to manage clusters that have those solutions enabled.

vSAN Clusters and vSphere Lifecycle Manager

You can manage a vSAN cluster by using vSphere Lifecycle Manager baselines and baseline groups or by using a single image for that cluster. Working with vSAN clusters has its specifics regardless of whether you manage the cluster with a single image or with baselines.

Managing a vSAN Cluster by Using Recommendation Baseline Groups

You can update and upgrade the hosts in a vSAN cluster by using automatically generated system-managed baseline groups. Those system-managed baseline groups are called recommendation baseline groups. Recommendation baseline groups do not contain firmware and driver updates. Recommendation baselines contain only patch or upgrade baselines.

If you want to switch to using images for a vSAN cluster that contains ESXi hosts of versions earlier than 7.0, you must first use an upgrade baseline to upgrade the hosts. Then, you can switch to using a vSphere Lifecycle Manager image for the cluster.

For more information about recommendation baselines, see [About Recommendation Baseline Groups](#).

For more information about using baselines to manage hosts and clusters, see [Using vSphere Lifecycle Manager Baselines and Baseline Groups](#).

Managing a vSAN Cluster with a Single Image

The image that you use for a cluster defines the full software stack to run on the hosts in that cluster: ESXi version, vendor customization, drivers, and firmware. When you manage a vSAN cluster with a single image, you can take advantage of the functionalities that vSphere Lifecycle Manager images provide.

- You can update the firmware on all hosts in the vSAN cluster.
You perform firmware update by setting up an image that contains a firmware add-on and remediating the vSAN cluster against that image. For more information about performing firmware updates by using vSphere Lifecycle Manager images, see [Firmware Updates with vSphere Lifecycle Manager](#).
- You can run a hardware compatibility check for the cluster.
The hardware compatibility check task verifies that the image for the cluster can be successfully applied to all hosts and that it is compliant with the vSAN Hardware Compatibility List (HCL). For more information about hardware compatibility checks, see [vSphere Lifecycle Manager Hardware Compatibility Checks for Clusters and Hosts](#).
- You can check firmware compliance with the image.
When you perform a compliance check against the image for a cluster, firmware compliance is also checked. As a result, you can easily notice if a driver or firmware in your cluster becomes non-compliant. For more information about checking the compliance of a cluster against an image, see [Check the Compliance Against a Single Image](#).
- You can use vSphere Lifecycle Manager recommended images.
When you manage a vSAN cluster with vSphere Lifecycle Manager images, the vSAN recommendation engine does not generate vSAN health alarms or recommendation baselines for that cluster. However, vSphere Lifecycle Manager generates pre-validated images include a recommended firmware version for the hosts in your vSAN cluster. For more information about vSphere Lifecycle Manager recommended images, see [vSphere Lifecycle Manager Recommended Images](#).

Managing a vSAN Witness Host with a Single Image

Starting with vSphere 8.0 Update 2, you can start managing a standalone vSAN witness host with a single image. The witness host can be dedicated to a vSAN stretched cluster or support up-to 64 two-node vSAN clusters. You upgrade a witness host that you manage with a vSphere Lifecycle Manager image in the same way you upgrade any other standalone host. For more information, see [How to Use vSphere Lifecycle Manager Images](#) and [vSphere Lifecycle Manager and the vSAN Witness Hosts](#).

Remediation Specifics of vSAN Clusters

Whether you manage a vSAN cluster with baselines or with a single image, remediating the hosts that are part of a vSAN cluster has its specifics.

When you remediate hosts that are part of a vSAN cluster, you must be aware of the following behavior:

- vSphere Lifecycle Manager puts only one host at a time in maintenance mode.
- vSphere Lifecycle Manager remediates hosts that are part of a vSAN cluster sequentially.
- Because vSphere Lifecycle Manager handles the remediation of the hosts sequentially, the host remediation process might take an extensive amount of time to finish.
- vSphere Lifecycle Manager remediates vSAN clusters with configured fault domains by upgrading all hosts from one fault domain first and then upgrading the hosts in the next fault domain.
- For a vSAN stretched cluster, vSphere Lifecycle Manager first remediates the hosts from the preferred site and then proceeds with remediating the hosts in the secondary site.

Host Maintenance Mode and vSAN Clusters

You can remediate a host that is in a vSAN cluster in two ways, depending on how you want to handle the virtual machines on the host:

- You can put the host in maintenance mode manually and remediate the host by using vSphere Lifecycle Manager.
- You can have the host enter maintenance mode during the vSphere Lifecycle Manager remediation process.

In the vSphere Client, when you put a host from a vSAN cluster into maintenance mode, you can choose between multiple options: Ensure accessibility, Full data evacuation, and No data evacuation. The Ensure accessibility option is the default option, and means that when you put a host in maintenance mode, vSAN ensures that all accessible virtual machines on the host remain accessible. To learn more about each of the options, see the "Place a Member of a vSAN Cluster in Maintenance Mode" topic in the *vSphere Storage* documentation.

During remediation, vSphere Lifecycle Manager puts the hosts from the vSAN cluster in maintenance mode and handles the virtual machines on the host in the manner of the default Ensure accessibility option.

If a host is a part of a vSAN cluster, and any virtual machine on the host uses a VM storage policy with the setting for "Number of failures to tolerate=0", the host might experience unusual delays when it enters maintenance mode. The delay occurs because vSAN has to migrate the virtual machine data from one disk on the vSAN datastore cluster to another. Delays might take up to hours. You can work around this by setting the "Number of failures to tolerate=1" for the VM storage policy, which results in creating two copies of the virtual machine files on the vSAN datastore.

vSAN Health Check

vSphere Lifecycle Manager performs a remediation pre-check of vSAN clusters to ensure successful remediation. The vSAN health check is part of the remediation pre-check.

The vSAN health check gives you information about the cluster state and whether you must take extra actions to ensure successful remediation. Even if you do not take the recommended actions, you can still remediate the vSAN cluster or a host from the cluster. vSphere Lifecycle Manager successfully puts the host in maintenance mode and applies software updates on the host successfully. However, the host might fail to exit maintenance mode, and the remediation process might fail. As a result, the host from the vSAN cluster is upgraded, but you must take manual steps to take the host out of maintenance mode.

Using vSphere Lifecycle Manager Images to Remediate vSAN Stretched Clusters

When you manage a vSAN stretched cluster or a two-node ROBO cluster with vSphere Lifecycle Manager, you can manage the hosts in the cluster with a single image that is different from the image used to upgrade the dedicated witness host. With vSphere 8.0 Update 2, you upgrade the vSAN witness host in the same way as you upgrade a standalone host.

What Is a Stretched Cluster

A stretched cluster is a deployment model in which two or more hosts are part of the same logical cluster but are located in separate geographical locations. Every vSAN stretched cluster or two-node ROBO cluster has a witness host, which is a standalone host that is not a member of the respective cluster but is associated with it. The witness host of a vSAN cluster is managed by the same vCenter Server where the respective stretched or ROBO cluster resides.

vSphere Lifecycle Manager and the vSAN Witness Hosts

The vSAN witness host is a physical or virtual ESXi host that contains the witness components of virtual machine objects stored in the vSAN cluster. The witness host does not support workloads and is not a data node. A single stretched or two-node ROBO cluster can have only one witness host.

You can use vSphere Lifecycle Manager images to manage a vSAN stretched cluster and its witness host. Starting with vSphere 8.0 Update 2, you define separate images for the vSAN cluster and for the witness host. See [Upgrading vSAN Stretched Clusters by Using a vSphere Lifecycle Manager Image](#). The following requirements exist:

- vCenter Server must be version 8.0 Update 2 and later.
- The witness host must be ESXi version 7.0 Update 2 and later.
- The witness host can be a virtual server or a physical server.
- The witness host can be a dedicated witness host or a shared witness host.
- The witness host must be upgraded before the hosts in the associated vSAN stretched or two-node cluster.
- The witness host and the associated vSAN clusters must not be upgraded in parallel.
- You cannot run virtual machines on a witness host. If vSphere Lifecycle Manager detects any stale virtual machines running on a witness host, during the remediation of the standalone host vSphere Lifecycle Manager sets the **VM power state** remediation setting to **Do not change power state**. For more information, see [Configure vSphere Lifecycle Manager Remediation Settings for Clusters or Standalone Hosts that You Manage with a Single Image](#).

You start using vSphere Lifecycle Manager images to manage the witness host by performing any of the following tasks:

- You switch from using vSphere Lifecycle Manager baselines to using vSphere Lifecycle Manager images for an existing vSAN stretched or two-node ROBO cluster and for the dedicated standalone host.
Note: The transition to using images is blocked if the witness host is of ESXi version earlier than 7.0 Update 2. In such cases, you can use baselines to upgrade the witness host to version 7.0 Update 2 or later, and then you can start managing the witness host with a single vSphere Lifecycle Manager image.
- You convert an existing vSAN cluster that uses a single image into a stretched cluster with a virtual witness host.
- You upgrade to version 8.0 Update 2 and later for vCenter Server and version 7.0 Update 2 or later for the witness host.

You stop using vSphere Lifecycle Manager images to manage the witness host in the following cases:

- You convert an existing vSAN stretched cluster that uses images into a regular vSAN cluster.
- You deactivate vSAN for an existing vSAN stretched cluster that you manage with a single image.

Important: With vSphere 8.0, you can use vSphere Lifecycle Manager images to manage standalone hosts in your vCenter Server inventory. Starting with vSphere 8.0 Update 2, you can apply a separate vSphere Lifecycle Manager image to the witness host of a vSAN cluster. You can start managing the witness host with a vSphere Lifecycle Manager image at the time of adding the host to the inventory or you can transition an existing standalone host that uses a single image to a witness host.

Upgrading vSAN Stretched Clusters by Using a vSphere Lifecycle Manager Image

You upgrade a stretched vSAN cluster along with the associated witness host by using a single vSphere Lifecycle Manager image, if the following requirements exist:

- vCenter Server must be version 7.0 Update 3 and later.

- The witness host must be ESXi version 7.0 Update 2 and later.
- The witness host is a virtual server and a dedicated witness host.

Starting with vSphere 8.0 Update 2, once you upgrade the vCenter Server instance to version 8.0 Update 2, the virtual dedicated witness host is no longer managed with the vSphere Lifecycle Manager image that you define for the vSAN cluster. You can use a full vSphere Lifecycle Manager image to upgrade a witness host in the same way you upgrade a standalone host. The desired image you apply on a witness host can contain a base ESXi image, and any user components, solution components, or OEM add-ons.

With vSphere 8.0 Update 2, for stretched vSAN clusters, you must first upgrade the witness host with the separate vSphere Lifecycle Manager image you configured and then proceeds to remediating the hosts in the preferred site and the secondary site. If all hosts in the preferred site are in a compliant state, then vSphere Lifecycle Manager skips the preferred site and starts remediating the hosts from the secondary site. If any host in the entire cluster is in an incompatible state, remediation stops. For more information about fault domain-aware remediation and the order in which vSphere Lifecycle Manager remediates the hosts in a vSAN cluster, see [Using vSphere Lifecycle Manager Images to Remediate vSAN Clusters with Configured Fault Domains](#).

To remediate the witness host against a single vSphere Lifecycle Manager image, the following requirements exist:

- vCenter Server must be version 8.0 Update 2 and later.
- The witness host must be ESXi version 7.0 Update 2 and later.
- The witness host can be a virtual server and a physical server.
- The witness host can be a dedicated witness host and a shared witness host.

Using vSphere Lifecycle Manager Images to Remediate vSAN Clusters with Configured Fault Domains

In vSAN clusters with configured fault domains, vSphere Lifecycle Manager remediates the hosts in an order that vSphere Lifecycle Manager calculates by factoring in the defined fault domains.

What Is a Fault Domain?

A fault domain consists of one or more vSAN hosts grouped according to their physical location in the data center. When configured, fault domains enable vSAN to tolerate failures of entire physical racks as well as failures of a single host, capacity device, network link, or a network switch dedicated to a fault domain. You can configure fault domains for non-stretched and stretched vSAN clusters. For more information about configuring fault domains, see the *Administering VMware vSAN* documentation.

Upgrading vSAN Clusters Configured with Multiple Fault Domains

vSphere Lifecycle Manager remediates vSAN clusters with configured fault domains by remediating all hosts in one fault domain at a time. To define the order of fault domains, vSphere Lifecycle Manager calculates and assigns priority to each fault domain for the vSAN cluster.

Remediation starts with the fault domain that has the highest priority. The priority of a fault domain is determined by the number of non-compliant hosts in that fault domain. The fewer non-compliant hosts in a fault domain, the higher the priority of that fault domain. However, if multiple fault domains have the same priority, vSphere Lifecycle Manager selects the first fault domain from the list of fault domains.

After vSphere Lifecycle Manager selects a fault domain, vSphere Lifecycle Manager uses DRS recommendations to select the optimal host within that domain to be remediated.

For fault domain-aware remediation of vSAN clusters, the following requirements exist:

- vCenter Server must be version 7.0 Update 1 and later
- The ESXi hosts must be version 7.0 and later

Upgrading vSAN Clusters Enabled with NSX or vSphere Supervisor

You can remediate a vSAN cluster against a vSphere Lifecycle Manager image that contains the same ESXi version as the ESXi version currently on the hosts, but the latest versions of NSX and vSphere Supervisor components. In that case, vSphere Lifecycle Manager upgrades only those components, without upgrading the ESXi version. Even in those cases, vSphere Lifecycle Manager still recognizes the configured fault domains for the vSAN cluster and performs the solution upgrade in accordance with the fault domain configuration.

For fault domain-aware remediation of vSAN clusters with enabled NSX or vSphere Supervisor, the following requirements exist:

- vCenter Server must be version 7.0 Update 2
- The ESXi hosts version 7.0 and later

Updating Firmware in vSAN Clusters

You can use vSphere vSphere Lifecycle Manager images to upgrade the firmware of the servers that run in your vSAN clusters.

In a vSAN cluster, the SCSI controller firmware and the physical drive firmware are handling most of the data communication. To ensure your vSAN cluster health, you must perform controller firmware updates, when necessary.

Because firmware updates affect the hardware layer in your vSphere environment, they usually are rare events. Firmware updates occur during initial ESXi host setup or during major updates of vSphere or vSAN.

In earlier vSphere releases, firmware updates are delivered as baselines in the vSAN-managed baseline group. You must use a special vendor-provided tool that vSAN uses to detect, download, and install firmware updates. However, in vSphere 8.0, the recommendation baseline group that vSAN generates contains only patch updates and driver updates. It no longer contains firmware updates. As a result, you cannot use baselines to update the firmware in your vSAN clusters if the ESXi hosts are of version 7.0 and later. You can still use baselines to perform firmware updates on hosts of earlier versions, for example ESXi 6.7. But to perform firmware updates on hosts that are of ESXi version 8.0 and later and that are in a vSAN cluster, you must manage that cluster with a single image. You must also deploy an OEM-provided hardware support manager and register it as a vCenter Server extension. The hardware support manager inspects the hardware of the hosts in the cluster and lists available and compatible firmware versions, which you can add to the image for the cluster. The actual firmware update happens upon remediating the cluster against an image that contains a firmware add-on.

For more information about the requirements for using images, see [What Are the Requirements for Using vSphere Lifecycle Manager](#).

For more information about performing firmware updates by using images, see [Firmware Updates with vSphere Lifecycle Manager](#).

About Recommendation Baseline Groups

vSAN creates system-managed baseline groups called recommendation baseline groups. You use recommendation baseline groups to upgrade the hosts in a vSAN cluster to the latest supported ESXi version, to patch the hosts with critical patches, or to update drivers on the host.

vSAN generates recommendation baseline groups automatically. If your vSphere environment does not contain any vSAN clusters, no recommendation baseline groups are generated. For each vSAN cluster in the vSphere inventory, vSphere Lifecycle Manager displays a single recommendation baseline group. You cannot edit or delete a recommendation baseline group and you cannot add it to custom baseline groups.

Recommendation baseline groups can contain any of the following software updates:

- An upgrade baseline that contains an ESXi upgrade image by a certified vendor with the latest tested and recommended version for the vSAN cluster.
- One or multiple patch baselines that contain recommended critical patches for the ESXi version of the hosts in the vSAN cluster.
- Recommended drivers for the ESXi hosts in the vSAN cluster.

Note: Recommendation baseline groups no longer contain firmware updates. To update the firmware on your hosts, you must convert to using a single image for the vSAN cluster.

How Does vSphere Lifecycle Manager Generate Recommendation Baselines?

A vSAN recommendation engine regularly checks the current state of the software installed on the hosts in the vSAN cluster against the vSAN Hardware Compatibility List (HCL). If update recommendations are detected, the engine downloads all new critical patches and upgrade images and generates a vSAN cluster-level baseline. All available baselines are packed together in a recommendation baseline group and made available for use by vSphere Lifecycle Manager.

Every 24 hours, vSphere Lifecycle Manager runs an automatic check for a recommendation baseline group with build recommendations coming from vSAN. If a new recommendation baseline group is detected, vSphere Lifecycle Manager automatically attaches the vSAN recommendation baseline group to the vSAN cluster.

After refreshing the vSAN recommendation baseline group, vSphere Lifecycle Manager automatically performs a compliance check operation on the vSAN clusters against the updated recommendation baseline group. Operations such as adding and removing hosts from an existing vSAN cluster also trigger refresh of the attached recommendation baseline group, followed by a compliance check.

System Requirements for Using vSAN Recommendation Baseline Groups

- vCenter Server 7.0 or later.
vSphere Lifecycle Manager runs as a service in vCenter Server 7.0 and later.
- vSAN cluster that contains hosts of ESXi version 6.0 Update 2 and later.
- Constant access of the vSphere Lifecycle Manager host machine to the Internet.

vSphere Lifecycle Manager and vSphere Supervisor

You can enable vSphere Supervisor on vSphere clusters that you manage with a single vSphere Lifecycle Manager image.

vSphere Lifecycle Manager and vSphere Supervisor with vSphere Networking

You start using vSphere Supervisor with vSphere networking on a cluster that uses a single vSphere Lifecycle Manager image by enabling the cluster for **Workload Management**. A cluster enabled for **Workload Management** is called a Supervisor. You enable a cluster for **Workload Management** from the **Workload Management** user interface in the vSphere Client.

Requirements

- Verify that all ESXi hosts in the Supervisor are version 7.0 Update 1 and later.
- Verify that the ESXi hosts in the Supervisor are assigned the required license.
- Verify that vCenter Server is version 7.0 Update 1 and later.
- **Note:** You can manage the lifecycle of a Supervisor with either vSphere Lifecycle Manager baselines or vSphere Lifecycle Manager images. However, you cannot convert a Supervisor that uses vSphere Lifecycle Manager baselines to a Supervisor that uses vSphere Lifecycle Manager images. To use vSphere Lifecycle Manager images for a

Supervisor, you must first switch the cluster which is not yet enabled for **Workload Management** to using images and then enable vSphere Supervisor on that cluster.

- Review the configuration requirements and additional information in *Working with vSphere Lifecycle Manager* chapter in the documentation.

Supported Workflows

The following workflows are supported for any Supervisor that uses vSphere Lifecycle Manager images and is configured to use the vSphere networking stack.

- You can upgrade a Supervisor to the latest version of vSphere Supervisor. You can also upgrade the ESXi version of the hosts in the Supervisor.

You perform the upgrade of the Supervisor from the **Workload Management** user interface in the vSphere Client. You upgrade the ESXi version of the hosts in the Supervisor by remediating the cluster from the vSphere Lifecycle Manager user interface in the vSphere Client.

Note: You cannot perform a simultaneous upgrade of both vSphere Supervisor and ESXi.

- You deactivate vSphere Supervisor from the **Workload Management** user interface in the vSphere Client. You can deactivate vSphere Supervisor on a cluster that uses a single vSphere Lifecycle Manager image.

When you deactivate vSphere Supervisor, you can use the cluster for traditional virtual machine workloads.

- You can add and remove hosts to and from a cluster that has both vSphere Supervisor and vSphere Lifecycle Manager enabled.

For more information about adding and removing hosts to and from a cluster, see the *vCenter Server and Host Management* documentation.

For detailed information about working with a Supervisor that uses a single image, see the *Working with vSphere Lifecycle Manager* section in the documentation.

Upgrading the vSphere Supervisor Components in a vSAN Cluster

To upgrade the vSphere Supervisor components installed on the hosts in a cluster, you must remediate the cluster against a vSphere Lifecycle Manager image that contains the latest version of those vSphere Supervisor components.

In a vSAN cluster with configured fault domains, vSphere Lifecycle Manager recognizes the configured fault domains for the cluster and performs the solution upgrade in accordance with the fault domain configuration. If the vSAN cluster is a stretched cluster, you must upgrade the witness host separately, after vSphere Lifecycle Manager finishes remediating all fault domains. For more information about remediating vSAN stretched clusters and vSAN clusters configured with fault domains, see [Using vSphere Lifecycle Manager Images to Remediate vSAN Clusters with Configured Fault Domains](#). For more information about stretched clusters, see the *vSAN Planning and Deployment* documentation.

Scalability

For information about the scalability that vSphere Lifecycle Manager supports, visit the VMware Configuration Maximums Matrix at <https://configmax.vmware.com/>.

vSphere Lifecycle Manager and vSphere Supervisor with NSX Networking

You can start using vSphere Supervisor with NSX networking on a cluster that uses a single vSphere Lifecycle Manager image by enabling the cluster for **Workload Management**. A cluster enabled with **Workload Management** is called a Supervisor. You enable the cluster for **Workload Management** from the **Workload Management** user interface in the vSphere Client. Upon enabling **Workload Management** on a Supervisor configured with the NSX networking stack, vSphere Lifecycle Manager installs the Spherelet VIB on every ESXi host in the cluster.

Requirements

- Verify that all ESXi hosts in the Supervisor are version 7.0 Update 2 or later.
- Verify that all ESXi hosts in the Supervisor are assigned the required license.
- Verify that vCenter Server is version 7.0 Update 2 or later.
- **Note:** You can manage the lifecycle of a Supervisor with either vSphere Lifecycle Manager baselines or vSphere Lifecycle Manager images. However, you cannot convert a Supervisor that uses vSphere Lifecycle Manager baselines to a Supervisor that uses vSphere Lifecycle Manager images. To use vSphere Lifecycle Manager images for a Supervisor, you must first switch the cluster which is not yet enabled for **Workload Management** to using images and then enable vSphere Supervisor on that cluster.
- Review the configuration requirements and additional information in the *Working with vSphere Lifecycle Manager* chapter in the documentation.

Supported Workflows

The following workflows are supported for any Supervisor that uses vSphere Lifecycle Manager images and is configured to use the NSX networking stack.

- You can upgrade a Supervisor to the latest version of vSphere Supervisor. You can also upgrade the ESXi version of the hosts in the Supervisor.
You perform the upgrade of the Supervisor from the **Workload Management** user interface in the vSphere Client. During upgrade, vSphere Lifecycle Manager upgrades the Spherelet VIB on the hosts to make it compatible with the new version of vSphere Supervisor or the new version of ESXi.
You upgrade the ESXi version of the hosts in the Supervisor by remediating the cluster from the vSphere Lifecycle Manager user interface in the vSphere Client.
Note: You cannot perform a simultaneous upgrade of both vSphere Supervisor and ESXi.
- You deactivate vSphere Supervisor from the **Workload Management** user interface in the vSphere Client. You can deactivate vSphere Supervisor on a cluster that uses a single vSphere Lifecycle Manager image.
When you deactivate vSphere Supervisor, you can use the cluster for traditional virtual machine workloads.
- You can add and remove hosts to and from a Supervisor that uses vSphere Lifecycle Manager images.
When you add a host to a Supervisor that you manage with a single vSphere Lifecycle Manager image, vSphere Lifecycle Manager automatically installs the Spherelet VIB on the newly added host.
When you remove a host from a Supervisor that you manage with a single vSphere Lifecycle Manager image, vSphere Lifecycle Manager removes the Spherelet VIB from the host. vSphere Lifecycle Manager also deletes the Spherelet VIB from a host that you move to another Supervisor.
For more information about adding and removing hosts to and from a cluster, see the *vCenter Server and Host Management* documentation.

For detailed information about working with a Supervisor that uses a single image, see the *Working with vSphere Lifecycle Manager* chapter in the documentation.

Upgrading the vSphere Supervisor Components in a vSAN Cluster

To upgrade the vSphere Supervisor components installed on the hosts in a cluster, you must remediate the cluster against a vSphere Lifecycle Manager image that contains the latest version of those vSphere Supervisor components.

In a vSAN cluster with configured fault domains, vSphere Lifecycle Manager recognizes the configured fault domains for the cluster and performs the solution upgrade in accordance with the fault domain configuration. If the vSAN cluster is a stretched cluster, you must upgrade the witness host separately, after vSphere Lifecycle Manager finishes remediating all fault domains. For more information about remediating vSAN stretched clusters and vSAN clusters configured with fault domains, see [Using vSphere Lifecycle Manager Images to Remediate vSAN Clusters with Configured Fault Domains](#). For more information about stretched clusters, see the *vSAN Planning and Deployment* documentation.

Scalability

For information about the scalability that vSphere Lifecycle Manager supports, visit the VMware Configuration Maximums Matrix at <https://configmax.vmware.com/>.

vSphere Lifecycle Manager and VMware NSX®

You can use vSphere Lifecycle Manager baselines for upgrade operations in environments with NSX 3.0. Starting with vSphere 7.0 Update 1 and NSX 3.1, you can also use vSphere Lifecycle Manager images to manage clusters that have NSX enabled.

Using vSphere Lifecycle Manager Baselines to Upgrade ESXi Hosts in an Environment With VMware NSX® 3.0

You can use vSphere Lifecycle Manager to upgrade the ESXi hosts in an environment where you have NSX enabled.

Requirements

- Verify that the ESXi hosts to upgrade are version 6.7 or later.
- vCenter Server version 6.7 or earlier
- NSX 3.0
- Verify that the vmknics on the ESXi host are properly configured and the DHCP server works properly.

Workflow

1. Upgrade vCenter Server to version 7.0.
For more information about upgrading vCenter Server, see the *vSphere Upgrade* documentation.
2. Import an ESXi 7.0 ISO image to the vSphere Lifecycle Manager depot.
For more information, see [Import an ISO Image to the vSphere Lifecycle Manager Depot](#).
3. Download the NSX 3.0.0 NSX Kernel Module for VMware ESXi 7.0 from <http://customerconnect.vmware.com>.
4. Import the kernel module to the vSphere Lifecycle Manager depot.
For more information, see [Import Updates to the vSphere Lifecycle Manager Depot](#).
5. Create an upgrade baseline with the imported ESXi 7.0 ISO image.
For more information, see [Create, Edit, or Delete vSphere Lifecycle Manager Baselines and Baseline Groups](#).
6. Create an extension baseline with the uploaded NSX kernel module.
For more information, see [Create, Edit, or Delete vSphere Lifecycle Manager Baselines and Baseline Groups](#).
7. Create a baseline group that contains the newly created upgrade and extension baselines.
For more information, see [Create a Host Baseline Group](#).
8. Attach the baseline group to a cluster.
For more information, see [Attach Baselines and Baseline Groups to Objects](#).
9. Remediate the cluster against the attached baseline group.
For more information, see [Remediate ESXi Hosts Against a Single Baseline or Multiple Baselines](#).

Using vSphere Lifecycle Manager Images in an Environment With NSX 3.1

You can use NSX on a cluster that you manage with a single vSphere Lifecycle Manager image. As a result, you can use NSX Manager to install, upgrade, or remove NSX components on the ESXi hosts in a cluster that you manage with a single image.

Requirements

- Verify that all ESXi hosts in the cluster are version 7.0 Update 1 or later.
- Verify that vCenter Server is version 7.0 Update 1 or later.
- Verify that NSX is version 3.1 or later.
- Verify that a vSphere Distributed Switch (VDS) is configured to manage the NSX traffic.

Supported Workflows

The following workflows are supported for clusters that are enabled for both vSphere Lifecycle Manager images and NSX.

- You can enable NSX on a cluster that you manage with a single vSphere Lifecycle Manager image.
You perform the operation by configuring a transport node profile (TNP) for the cluster in the NSX Manager. In the NSX Manager, you can either manually add a TNP to the cluster, or automatically generate one in the **Getting Started** wizard. You can continue leveraging individual transport node configurations, but you must always use a TNP for the clusters that you manage with a single vSphere Lifecycle Manager image. You cannot enable NSX on a cluster that uses a single vSphere Lifecycle Manager image if you choose to only use individual transport node configurations for the hosts in the cluster.
- You can add hosts to a cluster that you manage with a single vSphere Lifecycle Manager image and that is enabled with NSX. You can also remove hosts from such a cluster.
You perform the add and remove host operations in the vSphere Client. When you add a host to the cluster, vSphere Lifecycle Manager automatically installs the NSX component to the newly added host. To add a host to a cluster that you manage with a single vSphere Lifecycle Manager image, the host must be added to the VDS associated with the TNP. Otherwise, the host cannot fully work with NSX.
When you move a host from one cluster that uses a single vSphere Lifecycle Manager image to another, vSphere Lifecycle Manager applies the target cluster's image together with the target NSX component to the newly added host. If a host is deleted from the vCenter Server inventory, the NSX component is uninstalled from the host.
For more information about adding and removing hosts to and from a cluster, see the *vCenter Server and Host Management* documentation.
- You can upgrade NSX 3.1 to a later version in a cluster that you manage with a single vSphere Lifecycle Manager image.
You perform the operation from the NSX Manager.
- You can upgrade both NSX and ESXi in a single vSphere Lifecycle Manager remediation task. The workflow is supported only if you upgrade from NSX version 3.1.
In the NSX Manager, you stage the NSX upgrade as part of the image that the cluster uses. From the vSphere Lifecycle Manager user interface in the vSphere Client, you can further edit the image and you initiate remediation of the cluster. During remediation, vSphere Lifecycle Manager applies both the NSX and ESXi upgrades to the hosts in the cluster. For more information, see the *Upgrading NSX Guide* documentation.
- You can switch from using vSphere Lifecycle Manager baselines to using a vSphere Lifecycle Manager image for a cluster that is enabled with NSX.
- You can uninstall NSX from a cluster that you manage with a single vSphere Lifecycle Manager image.
- You can check the compliance, generate a remediation pre-check report, and remediate a cluster that you manage with a single vSphere Lifecycle Manager image and that is enabled with NSX.
You perform the check compliance, generate a remediation pre-check, and remediation operations in the vSphere Client. Whenever you change the NSX configuration in the NSX Manager, the compliance state of the cluster that you see on the **Updates** tab for the cluster in the vSphere Client changes to non-compliant. You can remediate non-compliant hosts and clusters in the vSphere Client or you can solve the issues that cause non-compliance in the NSX Manager.
- You can back up and restore NSX.
- You can export the vSphere Lifecycle Manager image of a cluster that is enabled with NSX and import this image to another cluster that has both vSphere Lifecycle Manager images and NSX enabled.

For detailed information about all workflows that you perform in the NSX Manager, see the *NSX Administration* documentation.

Upgrading the NSX Components in a vSAN Cluster

To upgrade the NSX components installed on the hosts in a cluster, you must remediate the cluster against a vSphere Lifecycle Manager image that contains the latest version of those NSX components.

In a vSAN cluster with configured fault domains, vSphere Lifecycle Manager recognizes the configured fault domains for the cluster and performs the solution upgrade in accordance with the fault domain configuration. If the vSAN cluster is a stretched cluster, you must upgrade the witness host separately, after vSphere Lifecycle Manager finishes remediating all fault domains. For more information about remediating vSAN stretched clusters and vSAN clusters configured with fault domains, see [Using vSphere Lifecycle Manager Images to Remediate vSAN Clusters with Configured Fault Domains](#). For more information about stretched clusters, see the *vSAN Planning and Deployment* documentation.

Scalability

For information about the scalability that vSphere Lifecycle Manager supports, visit the VMware Configuration Maximums Matrix at <https://configmax.vmware.com/>.

Using vSphere Lifecycle Manager Images for Standalone Hosts with NSX 4.2 and later

You can use NSX on a standalone host that you manage with a vSphere Lifecycle Manager image. As a result, vSphere Lifecycle Manager becomes responsible for installing, upgrading, and removing NSX components on the host.

Requirements

- Verify that the standalone ESXi host and the vCenter Server instance are of version 8.0 Update 3 and later.
- Verify that NSX is of version 4.2 and later.

Supported Workflows

- You can enable NSX on a standalone host that uses vSphere Lifecycle Manager images.
You perform the operation by configuring NSX on the individual host in the NSX Manager UI. For more information about how to configure individual hosts as transport nodes, see the *Prepare ESXi Individual Hosts as Transport Nodes* procedure in the *NSX Installation Guide*.
- You can uninstall NSX from a standalone host that you manage with a vSphere Lifecycle Manager image.
- You can switch from using vSphere Lifecycle Manager baselines to using vSphere Lifecycle Manager images for an existing standalone host that is prepared as a transport node.
During the eligibility check on the transport node, vSphere Lifecycle Manager performs several additional checks and allows the operation to continue in case all of the following conditions are met:
 - The vCenter Server instance is of version 8.0 Update 3 and later.
 - The ESXi host is of version 8.0 Update 3 and later.
 - The NSX version on the standalone host is 4.2 and later.
 - The **Multi NSX** mode is not enabled on the vCenter Server instance. See *Multiple NSX Managers Managing a Single VMware vCenter* in the *NSX Installation Guide*.
- You can upgrade both NSX and ESXi on the standalone host by using a single vSphere Lifecycle Manager remediation task. You run the remediation operation in the vSphere Client.

- You can add a standalone host that has NSX already installed to a data center or a folder and enable vSphere Lifecycle Manager on the host during this process.
During the eligibility check on the host, vSphere Lifecycle Manager performs several additional checks and allows the operation to continue in case all of the following conditions are met:
 - The vCenter Server instance is of version 8.0 Update 3 and later.
 - The ESXi host is of version 8.0 Update 3 and later.
 - The NSX version on the standalone host is 4.2 and later.
 - The **Multi NSX** mode is not enabled on the vCenter Server instance. See *Multiple NSX Managers Managing a Single VMware vCenter* in the *NSX Installation Guide*.
- If you take a host out of a cluster that is prepared with a transport node profile (TNP) and is managed with a vSphere Lifecycle Manager image, the host becomes standalone. The host can retain the image of the cluster in case the following criteria are met:
 - The vCenter Server and ESXi host are of version vSphere 8.0 Update 3 and later.
 - NSX is of version 4.2 or later.
 - During the move operation the host stays connected or reachable.
 Otherwise, the image is removed from the host and the host is reverted back to using baselines.
- You can add a standalone host that used baselines to a cluster that you manage with a single vSphere Lifecycle Manager image and that is enabled with NSX.
When you add the standalone host to the cluster, vSphere Lifecycle Manager automatically installs the NSX component to the newly added host. To add a host to a cluster that you manage with a single vSphere Lifecycle Manager image, the host must be added to the VDS associated with the TNP. Otherwise, the host cannot fully work with NSX.

Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level

With vSphere Configuration Profiles, you can manage the configuration of all hosts in a cluster collectively. You set up a desired configuration that you can apply to all hosts in a single operation. You can also use a reference host and make its configuration the desired configuration for an entire cluster. Using vSphere Configuration Profiles ensures consistency in host configuration.

vSphere Configuration Profiles uses a declarative model to manage host configuration at a cluster level. With vSphere Configuration Profiles, you can perform the following tasks:

- Set a desired host configuration at a cluster level. The configuration is created and managed in the form of a human-readable JSON file backed by a JSON schema. Starting with vSphere 8.0 Update 2, you can edit the desired cluster configuration settings in the vSphere Client.
- Check the host compliance against the configuration for the cluster.
- Remediate the cluster to make non-compliant hosts compliant with the configuration set at the cluster level.

To perform all these tasks, you need the proper privileges. To view the full list of privileges required for using vSphere Configuration Profiles, see [Required Privileges for Using vSphere Configuration Profiles](#)

Starting with vSphere 8.0 Update 3, you can enable vSphere Configuration Profiles on a cluster that you manage with a single image and on a cluster that you manage with baselines. For clusters that you manage with baselines, you can also use vSphere Host Profiles, see the [vSphere Host Profiles](#) documentation.

You start using vSphere Configuration Profiles by enabling the feature during cluster creation or by transitioning to vSphere Configuration Profiles. You can switch from using host profiles to enabling and using vSphere Configuration Profiles at any time. The change is permanent. That is, if you switch to using vSphere Configuration Profiles, you cannot undo the transition.

For information about how to enable vSphere Configuration Profiles during cluster creation and set a desired configuration for the cluster, see [Enable vSphere Configuration Profiles During Cluster Creation](#).

For information about how to switch from using legacy configuration management tools, such as host profiles, to using vSphere Configuration Profiles, see [Transition to Using vSphere Configuration Profiles](#)

System Requirements for Using vSphere Configuration Profiles


- Your vCenter Server instance and all ESXi hosts in a cluster managed with a single image must be of version 8.0 Update 1 and later.
- Your vCenter Server instance and all ESXi hosts in a cluster managed with baselines must be of version 8.0 Update 3 and later.
- Required privileges.

Limitations to Using vSphere Configuration Profiles

- You can't deactivate vSphere Configuration Profiles.
- With vSphere 8.0 Update 1, you can enable vSphere Configuration Profiles for a cluster that uses vSphere Distributed Switch. Note that you can only view the vSphere Distributed Switch configuration of the hosts in the cluster. Both, your vCenter Server instance and all ESXi hosts in the cluster managed with a single image, must be of version 8.0 Update 1 and later.
- Starting with vSphere 8.0 Update 3, you can manage the vSphere Distributed Switch configuration of the hosts in a cluster with enabled vSphere Configuration Profiles. You can also use the workflows for managing and configuring the vSphere Distributed Switches described in the *vSphere Networking* documentation. Both, your vCenter Server instance and all ESXi hosts in the cluster managed with a single image, must be of version 8.0 Update 3 and later.
- You can't enable vSphere Configuration Profiles on a cluster with DPU-backed hosts.
- You can't enable vSphere Configuration Profiles if NSX is also enabled for the cluster.

How to Start Using vSphere Configuration Profiles for Clusters That You Manage with a Single Image

You can enable vSphere Configuration Profiles during the cluster creation or you can transition an existing cluster managed with a vSphere Lifecycle Manager image to vSphere Configuration Profiles.

Figure 62: Enable vSphere Configuration Profiles on a Cluster That You Manage with a Single Image

When you enable vSphere Configuration Profiles during the cluster creation, the cluster is created with a default configuration and you must set up the desired cluster configuration by editing its settings. For more information, see [Enable vSphere Configuration Profiles During Cluster Creation](#).

When you transition an existing cluster that you manage with a vSphere Lifecycle Manager image to vSphere Configuration Profiles, you set up a desired cluster configuration and as a result of the transitioning operation, all hosts in the cluster are compliant with the desired configuration. For more information, see [Transition to Using vSphere Configuration Profiles](#).

Starting with vSphere 8.0 Update 2, you can create a draft configuration and edit the configuration settings directly in the vSphere Client. For more information, see [Create a Draft Configuration in vSphere Client](#).

How to Start Using vSphere Configuration Profiles for Clusters That You Manage with Baselines

With vSphere 8.0 Update 3, you can enable vSphere Configuration Profiles during the cluster creation or you can transition an existing cluster managed with baselines to vSphere Configuration Profiles.

When you enable vSphere Configuration Profiles during the cluster creation operation, the cluster is created with an empty configuration. To create a desired cluster configuration, you must add at least one host to the cluster that you manage with baselines. Then you must select a reference host from which the configuration schema and the configuration settings are extracted and imported to the cluster as the desired cluster configuration. The reference host must be part of the cluster managed with baselines and must be of version 8.0 Update 3 and later. For more information, see [Enable vSphere Configuration Profiles During Cluster Creation](#).

When you transition an existing cluster that you manage with baselines to vSphere Configuration Profiles, you use a host from the cluster as a reference host to set up the desired cluster configuration. When the transitioning operation completes, all hosts in the cluster are compliant with the desired configuration. For more information, see [Transition to Using vSphere Configuration Profiles](#).

Enable vSphere Configuration Profiles During Cluster Creation

You can enable vSphere Configuration Profiles during the process of creating a cluster. However, in this case, there are some steps which you must perform to create a desired cluster configuration.

If you create a cluster that you manage with a single image, it is created with a default configuration. You must define the desired configuration for the cluster by editing the default settings after the cluster is successfully created. In case you create a cluster that you manage with baselines, it is created with an empty configuration and you must add at least one host to the cluster before setting up the desired cluster configuration. To define the desired configuration for the cluster that you manage with baselines, you must use a reference host from the cluster.

The overall workflow for enabling vSphere Configuration Profiles on a cluster that you manage with a single image and setting up a desired configuration is the following:

1. Create a cluster. For more information, see [Creating and Managing vSphere Lifecycle Manager Clusters](#).

2. Set up the desired configuration for the cluster.

Note: Starting with vSphere 8.0 Update 2, you can use vSphere Client to create a draft configuration and edit the host configuration settings for a cluster in the user interface.

- Option 1: Create a draft configuration by using one of the following options:
 - Use the settings of the default configuration for the cluster.
 - Reuse the configuration from an existing cluster by exporting its configuration document and importing it into the created draft configuration. See [Import a Configuration Document or a Configuration from a Reference Host Into an Existing Draft Configuration](#)
 - Reuse the configuration of one of the hosts in the cluster.

After creating the draft configuration, edit the configuration settings to achieve the desired configuration for the cluster. For more information, see [Create a Draft Configuration in vSphere Client](#) and [Edit the Host Settings of a Draft Configuration in vSphere Client](#).

- Option 2: Export the configuration schema for the newly created cluster and use it to create your own JSON configuration document, which you can import into the cluster.
The exported configuration schema document follows the JSON Schema format. You use it to create and edit your own JSON configuration document in an external JSON editor tool.
- Option 3: Export the default or the draft configuration document of the cluster, edit the JSON file manually, and import it back into the cluster. You can also export and use the default or the draft configuration schema for the cluster to facilitate yourself in editing the desired configuration document.
- Option 4: Extract the settings from a host in the vCenter Server inventory and import them into the newly created cluster to use them as the common desired configuration.

The overall workflow for enabling vSphere Configuration Profiles on a cluster that you manage with baselines and setting up a desired configuration is the following:

1. Create a cluster and add at least one host to the cluster. For more information about creating a cluster, adding and removing hosts to and from a cluster, see the *vCenter Server and Host Management* documentation.
2. Set up the desired configuration for the cluster.
 - Option 1: Create a draft configuration by reusing the configuration of one of the hosts in the cluster as reference.
After creating the draft configuration, edit the configuration settings to achieve the desired configuration for the cluster. For more information, see [Create a Draft Configuration in vSphere Client](#) and [Edit the Host Settings of a Draft Configuration in vSphere Client](#).
 - Option 2: Export the configuration schema for the newly created cluster and use it to create your own JSON configuration document, which you can import into the cluster.
The exported configuration schema document follows the JSON Schema format. You use it to create and edit your own JSON configuration document in an external JSON editor tool.
 - Option 3: Export the draft configuration document of the cluster, edit the JSON file manually, and import it back into the cluster. You can also export and use the default or the draft configuration schema for the cluster to facilitate yourself in editing the desired configuration document.

Transition to Using vSphere Configuration Profiles

For any cluster that you manage with a single image or with baselines, you can start using vSphere Configuration Profiles. The transition workflow allows you to enable vSphere Configuration Profiles on an existing cluster and set up the desired configuration for the cluster.

- Verify that vCenter Server and ESXi are of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image.
- Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
- Verify that if the cluster uses vSphere Distributed Switch, vCenter Server and all ESXi hosts in the cluster are of version 8.0 Update 1 and later.

- Verify that if the cluster uses vSphere Distributed Switch and you want to manage the vSphere Distributed Switch configuration with vSphere Configuration Profiles, vCenter Server and ESXi are of version 8.0 Update 3 and later.
- Verify that the cluster is not empty.
- Verify that all hosts in the cluster are compliant with the image for the cluster if the cluster is managed with a single image.
- Verify that NSX is not enabled on the cluster.
- Verify that the ESXi hosts in the cluster have the required license.
- Verify that you have the required privileges for the transitioning operation. See [Required Privileges for Using vSphere Configuration Profiles](#).

The transition workflow for a cluster that you manage with a single image starts with eligibility checks on the cluster and hosts. If the cluster is not eligible, the transition workflow cannot proceed. If the cluster is eligible, user can import a configuration from a host in the cluster or from a JSON file. The configuration is then validated. You can edit the configuration at this stage by using the export and import options. vSphere Configuration Profiles runs a pre-check to ensure that the configuration can be applied to all hosts. After you confirm the enablement of vSphere Configuration Profiles, eligibility checks are re-run and the cluster is remediated against the configuration that you set up.

Starting with vSphere 8.0 Update 3, you can enable vSphere Configuration Profiles on a cluster that you manage with baselines. The transition workflow starts with selecting a reference host whose configuration schema is imported and used as a desired cluster configuration schema.

Starting with vSphere 8.0 Update 1, you can enable vSphere Configuration Profiles on a cluster that uses a vSphere Distributed Switch.


Note: Starting with vSphere 8.0 Update 3, vSphere Configuration Profiles can manage the vSphere Distributed Switch configuration of the hosts in a cluster managed with a single image.

To manage any drifts related to the vSphere Distributed Switch configuration, you can use both options:

- The vSphere Distributed Switch workflows described in the *vSphere Networking* documentation
 - The vSphere Configuration Profiles workflows described in this chapter.
1. In the vSphere Client, navigate to a cluster that you manage with a single image or a cluster that you manage with baselines.
 2. On the **Configure** tab, click **Desired State > Configuration**.
 3. Click **Create configuration**.

The **Check Cluster Eligibility for Transition** task starts to ensure that the cluster is eligible for transitioning to vSphere Configuration Profiles.

4. In the **Manage configuration at cluster level** pane, follow the prompts to set up a desired configuration for the cluster and apply it to all hosts.



- a) Under **Create configuration**, select an import option and click **Next**.

- To import the configuration from a reference a host within the cluster, click **Import from reference host**, select the reference host, and confirm the import.

Note: This option is available for a cluster that you manage with baselines and for a cluster that you manage with a single image.

- To import the configuration from a JSON file that you created or edited, click **Import from file**, locate the JSON file with the configuration that you want to import and confirm the import.

Note: This option is only available for a cluster that you manage with a single image.

The **Validate Configuration for Transition** task starts to ensure that the imported configuration is valid and can be used as the desired configuration for the cluster.

b) Optional: Under **Validate configuration**, view the validation messages and edit the imported configuration.

- Click **Export configuration** to download the current configuration document and edit its settings.
- Click **Export configuration schema** to download the JSON file that represents the configuration schema.
- Click **Import configuration** to import the edited configuration document back into the cluster.

c) Click **Next**.

The **Precheck Configuration for Transition** task starts to ensure that the configuration can be applied to all hosts successfully.

d) Under **Pre-check and apply**, view pre-check details and remediation impact and click **Finish and apply**.

e) In the **Finish and apply** dialog box, confirm that you want to apply the configuration to all hosts in the cluster.

vSphere Configuration Profiles is enabled for the cluster and the cluster is remediated against the configuration you set up during this operation.

f) Optional: Click **Go to configuration** to view and edit the cluster configuration settings.

You can cancel the transition at any stage of the workflow. Later, if you did not discard the changes that you made during the transition workflow, you can resume the transition and use the configuration imported before canceling the transition operation. If you want to start the whole process anew, discard all your changes. When you discard your changes, the imported configuration is also deleted.

The cluster now uses vSphere Configuration Profiles to manage the configuration of its hosts. All hosts in the cluster are compliant with the desired configuration you set up during the transitioning operation.

View the current desired cluster configuration. For more information, see [View Host Settings in the Desired Configuration](#). You can also create a draft configuration and edit the host settings in the vSphere Client. For more information, see [Create a Draft Configuration in vSphere Client](#).

Working with Configuration Documents and Configuration Schemas

You can manually edit the configuration settings for a cluster by using the configuration document and the configuration schema available for the current configuration and for the draft configuration of the cluster.

What Are the Configuration Schema and the Configuration Document

You set up and manage the desired host configuration for an entire cluster by working with the configuration document for the respective cluster. The configuration document is a JSON file that you can download to your local machine and edit with a JSON editor tool.

Starting with vSphere 8.0 Update 2, you no longer need to download the configuration document to edit it. Instead, you can directly edit the current cluster configuration in the vSphere Client. You create a working copy of the cluster configuration, called a draft, edit the configuration settings, and apply your changes to the cluster in a single operation. For more information, see [Edit the Host Settings of a Draft Configuration in vSphere Client](#).

A valid configuration document contains a `profile` section and, optionally, a `host-specific` and `host-override` sections.

- The `profile` section contains common configuration that is applicable to all hosts in the cluster.
- The `host-specific` section represents configuration that can only be specified per host.
- The `host-override` section represents configuration that is overridden for a specific host in the cluster.

The configuration document is backed by a JSON schema, which is not editable. The configuration schema is a JSON file that represents the complete ESXi configuration. The schema follows the JSON Schema format and contains the default values for all host properties. The configuration schema is generated from the desired software specification defined in the image for the cluster. The configuration schema changes when you change the software specification for the cluster.

Export the Cluster Configuration or Cluster Configuration Schema

You can export the current cluster configuration in the form of a JSON file, which you can manually edit. You can also export the configuration schema for the cluster and use it to help you create your own configuration document from scratch or edit an exported configuration document.

- Verify that vCenter Server is of version 8.0 Update 1 or later.
 - Verify that ESXi is of version 8.0 Update 1 or later.
 - Verify that you have the required license.
 - Verify that you have the required privileges for the export operation. See [Required Privileges for Using vSphere Configuration Profiles](#).
1. In the vSphere Client, navigate to a cluster that you manage with a single image.
 2. On the **Configure** tab, click **Desired State > Configuration**.
 3. On the **Settings** tab, select your task.
 - To download the current configuration for the cluster, from the **Export** drop-down menu, click **Cluster configuration**.
 - To download the configuration schema for the cluster, from the **Export** drop-down menu, click **Cluster configuration schema**.

A dialog box appears that prompts you to download the JSON file that contains the current cluster configuration or the configuration schema.

4. Click **Download**.

The current configuration document or the configuration schema is saved locally on your machine.

Edit the configuration document or use the configuration schema to create a new JSON file with the desired host configuration for the cluster. Then, import the newly created or edited configuration document into the cluster.

Export the Draft Cluster Configuration or Draft Cluster Configuration Schema

You can export the draft cluster configuration in the form of a JSON file and manually edit it. You can also export the draft configuration schema for the cluster and use it to help you create your own configuration document from scratch or edit an exported configuration document.

- Verify that vCenter Server is of version 8.0 Update 2 or later.
- Verify that ESXi is of version 8.0 Update 1 or later.
- Verify that you have the required license.
- Verify that you have a working copy of the configuration created for the cluster.

- Verify that you have the required privileges for the export operation. See [Required Privileges for Using vSphere Configuration Profiles](#).
1. In the vSphere Client, navigate to a cluster that you manage with a single image.
 2. On the **Configure** tab, click **Desired State > Configuration**.
 3. On the **Draft** tab, select your task.
 - To download the current draft configuration document for the cluster, click the horizontal ellipses icon and select **Export**.
 - To download the draft configuration schema for the cluster, click the horizontal ellipses icon and select **Export configuration schema**.

A dialog box appears that prompts you to download the JSON file that contains the current draft cluster configuration or the draft configuration schema.

4. Click **Download**.

The current draft configuration document or the draft configuration schema is saved locally on your machine.

Edit the draft configuration document or use the draft configuration schema to create a new JSON file with the desired host configuration for the cluster. Then, import the newly created or edited draft configuration document into the cluster.

Extract The Configuration from a Reference Host

You can extract the configuration of any host in the vCenter Server inventory. The configuration is saved on your local machine as a JSON file, which you can edit. You can import the extracted configuration into any other cluster that has vSphere Configuration Profiles enabled.

- Verify that vCenter Server is of version 8.0 Update 1 or later.
- Verify that ESXi is of version 8.0 Update 1 or later.
- Verify that you have the required license.
- Verify that you have the required privileges for the extract operation. See [Required Privileges for Using vSphere Configuration Profiles](#).

1. In the vSphere Client, navigate to a cluster that you manage with a single image.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Settings** tab, click the **Export** drop-down menu and select **Reference host configuration**.
The **Export Configuration** wizard opens.

4. Follow the prompts to download the configuration of the selected reference host.

You can activate the **Show hosts from existing cluster** toggle to narrow down the list of hosts available in your vCenter Server to only the hosts available in the selected cluster.

The configuration of a selected reference host is exported and downloaded to your local machine as a JSON file.

Import the downloaded JSON configuration into another cluster to reuse the configuration. Alternatively, you can also edit the downloaded file first and then import it into the same or different cluster.

View Host Settings in the Desired Configuration

You can view information about how each host setting is configured in the current desired configuration for the cluster. For each host setting defined in the configuration document, you can see the values as defined for all hosts in the cluster and also any host overrides.

- Verify that vCenter Server and ESXi are of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image.
- Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
- Verify that you have the required license.

- Verify that you have the required privileges for viewing host settings in the desired configuration. See [Required Privileges for Using vSphere Configuration Profiles](#).

With vSphere 8.0 Update 2, the common settings, the host-specific settings, and the host overrides can be configured manually in the configuration document or in the vSphere Client. In the vSphere Client, you can view both, the settings of the currently applied desired configuration and the editable settings of the draft configuration of a cluster.

1. In the vSphere Client, navigate to a cluster that you manage with a single image or to a cluster that you manage with baselines.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Settings** tab, view information about each setting defined in the current desired configuration

The screenshot displays the vSphere Client interface for a cluster named 'My Cluster'. The 'Configure' tab is selected, and the 'Desired State > Configuration' path is followed. The 'Settings' sub-tab is active, showing a list of settings. The 'vmkernel' setting is highlighted with a green circle labeled 3.a. The right pane shows the 'vmkernel' configuration details, including 'module_settings' (highlighted with a green circle labeled 3.b) and 'options'. Below this, another 'Configuration' section is visible, showing a list of settings including 'network_vss', 'security', 'services', 'storage', and 'storage_iscsi'. A green circle labeled 3.c is placed over the 'security' setting in this section.

Create a Draft Configuration in vSphere Client

When you want to make changes to the desired configuration of a cluster that uses vSphere Configuration Profiles, you can create a draft configuration for that cluster in the vSphere Client.

- Verify that vCenter Server is of version 8.0 Update 2 or later and ESXi is of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image.
- Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
- Verify that you have the required license.
- Verify that you are using vSphere Configuration Profiles for the cluster.
- Verify that you have the required privileges for the operation. See [Required Privileges for Using vSphere Configuration Profiles](#).

You can create only one draft configuration at a time for a cluster. When you finish making changes to the host settings, you must run a remediation pre-check before applying the changes to all host in the cluster and making your draft configuration the new desired cluster configuration.

Note: More than one user with the required privileges can create a draft configuration for a cluster at the same time. When one of the users remediates the cluster with their draft configuration making it the desired configuration for the cluster, the other users editing their draft configurations for the same cluster will see an error message.

1. In the vSphere Client, navigate to a cluster that you manage with a single image or to a cluster that you manage with baselines.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Draft** tab, choose the method of creating a draft configuration for the cluster.

Option	Steps
Create a copy of the current cluster configuration	To create a working copy from the current cluster configuration, click Create draft. • If you enabled vSphere Configuration Profiles during cluster creation, the draft configuration contains the default settings of the cluster that you manage with a single image retrieved from the desired cluster image. • If you transition a cluster that you manage with a single image or a cluster that you manage with baselines to using vSphere Configuration Profiles, the draft configuration is created from the desired cluster configuration.
Import a configuration document into a cluster and edit it as a draft configuration	This option for creating a draft configuration is available only for clusters that you manage with a single image. 1. On the Draft tab, select Import from file. The Import Configuration from File dialog box appears. 2. Click Browse to locate the file that you want to import and click Import. 3. Click Close.
Import the configuration from one of the hosts in the cluster	Use this option to import the configuration from one of the hosts in the cluster managed with a single image or with baselines. Tip: If you enabled vSphere Configuration Profiles during cluster creation and the cluster is managed with baselines, a default draft configuration is not available for the cluster and you must use this option to import one from a reference host in the cluster 1. On the Draft tab, select Import from host.

Option	Steps
	The Select reference host dialog box appears. 2. Select one of the hosts listed for the cluster and click Import. Note: During the import operation, the configuration of the reference host is extracted and is compared to the current configurations of all hosts in the cluster. If a host in the cluster has configuration settings that differ from the settings of the reference host, these settings are marked in the draft configuration as host overrides. You can view the host overrides for a specific host setting on the Host overrides tab for that setting or under the <code>host-override</code> section of the draft configuration document.

You created a draft configuration of the cluster which you can use to edit the host settings in the vSphere Client.

You can edit the host settings within the draft configuration and remediate the cluster against the draft configuration in a single operation. For more information, see [Edit the Host Settings of a Draft Configuration in vSphere Client](#).

Edit the Host Settings of a Draft Configuration in vSphere Client


You can create a draft configuration for a cluster and edit the host settings in the vSphere Client.

- Verify that vCenter Server is of version 8.0 Update 2 or later and ESXi is of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image.
 - Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
 - Verify that you have the required license.
 - Verify that you have a draft configuration created for the cluster.
 - Verify that you have the required privileges for performing remediation. See [Required Privileges for Using vSphere Configuration Profiles](#).
1. In the vSphere Client, navigate to a cluster that you manage with a single image or to a cluster that you manage with baselines.
 2. On the **Configure** tab, click **Desired State > Configuration**.
 3. On the **Draft** tab, navigate to the host setting that you want to edit.
 4. Edit the host setting of the draft configuration.


Option	Steps
Configure a common setting that has no existing configuration	1. On the Common settings tab, click Configure settings. The Edit dialog box appears.

Option	Steps
--------	-------

Figure 63: Configure a common setting



Option	Steps
	2. Select the common setting from the left sidebar and configure its value. The settings of the selected item appear on the right of the Edit dialog box. 3. (Optional) Activate the Show advanced settings toggle, to view and edit the advanced settings. If there is at least one advanced setting configured, the Show advanced settings toggle is automatically activated. 4. (Optional) Deactivate the Show unconfigured settings toggle, to view and edit only the configured settings. By default, the toggle is activated. 5. Click Save. The configured settings appear under the Common settings tab.
Edit a common setting that is already configured	1. On the Common settings tab, navigate to the current configuration for the selected setting and select it from the tree on the left. 2. Click Edit. The Edit dialog box appears. 3. Navigate to the configured setting by clicking on the View settings links and edit the value of the setting. 4. Click Save.
Delete a common setting configuration	1. On the Common settings tab, select the setting you want to delete from the list of configured settings and click Delete. 2. Click Delete setting to confirm the deletion of the setting configuration.
Add a host override for a configured common setting	1. On the Host overrides tab, click the Add override drop-down menu and select an option. — To add a host override by filtering the available hosts in the cluster using their host names, click Add override by host name. The Add override for <setting_name> dialog box appears. Select a host and click Add override. — To add a host override by filtering the hosts using their BIOS UUID, click Add override by BIOS UUID. The Add override for <setting_name> dialog box appears. Enter the BIOS UUID of the host in the respective field and click Add override. You can use this option to add a host override for a host that is not currently part of the cluster. 2. From the Show overrides for drop-down menu, select the host override that you created. The configured common setting with its value appear under the drop-down menu.

Option	Steps
	3. Select from the options for editing. To edit the current configuration of the common setting, click Edit. The Edit - <BIOS UUID (<host_name>)> dialog box appears. Figure 64: Edit a host override for a configured common setting 

Option	Steps
	Navigate to the configured setting and change its value. – To apply the system default configuration which can be different from the common configuration of the setting, click Use system default and confirm your choice.

The settings of the draft configuration are configured.

You can view the changes between the current draft configuration and the desired configuration of the cluster by clicking **Show changes** on the **Draft** tab. You can run a draft pre-check to validate the draft configuration and ensure that it can be applied on all hosts in the cluster. For more information, see [Run a Remediation Pre-Check or a Draft Pre-Check](#). You can also remediate the cluster against the draft configuration and make it the new desired cluster configuration. For more information, see [Remediate a Cluster Against the Desired or Draft Configuration](#).

Import a Configuration Document or a Configuration from a Reference Host Into an Existing Draft Configuration

For any cluster that uses vSphere Configuration Profiles and has a draft configuration created, you can override and replace the configuration settings by importing a valid configuration document or using the configuration settings of a reference host in the cluster.

- Verify that vCenter Server is of version 8.0 Update 2 or later and ESXi is of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image.
- Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
- Verify that you have the required license.
- Verify that you have a draft configuration created for the cluster.
- Verify that you have the required privileges for the import operation. See [Required Privileges for Using vSphere Configuration Profiles](#).

You can import the configuration document of another cluster and reuse the configuration defined in that other cluster. You can also import a configuration document that you manually edited or created from scratch. You can also import the settings of one of the hosts available in the cluster and use them as common settings for all hosts in the cluster.

1. In the vSphere Client, navigate to a cluster that you manage with a single image or to a cluster that you manage with baselines.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Draft** tab, select your task.

Import Option	Steps
Import a JSON file	1. To import a configuration document, click the horizontal ellipses and select Import from file. The Import Configuration from File dialog box appears. 2. Click Browse to locate the file that you want to import and click Import and overwrite draft.
Import the configuration of a reference host in the cluster	1. To import the configuration settings from a reference host in the cluster, click the ellipses and select Import from host. The Select reference host dialog box appears. 2. Click the radio button for the respective host which configuration settings you want to use and click Import and overwrite draft.

4. Click **Close**.

The imported configuration settings become the new draft configuration for the entire cluster.

When you replace the settings in the draft configuration with the settings of a single host from the cluster, these settings become the common settings for all hosts in the cluster. If a host in the cluster has settings that differ from the new common settings for the cluster, the original host settings are saved in the new draft configuration as host overrides.

You can view the changes applied to the edited draft configuration of the cluster by clicking on **Show changes**. You can also start a draft pre-check to ensure that the configuration is valid and to compute the remediation impact on all hosts in the cluster. Then you can remediate the cluster to make all non-compliant hosts compliant.

Delete a Draft Configuration for a Cluster

You can discard all changes that you made to a draft configuration by deleting it in the vSphere Client. The current desired state of the cluster remains unaffected.

- Verify that vCenter Server is of version 8.0 Update 2 or later.
- Verify that ESXi is of version 8.0 Update 1 or later.
- Verify that you have a draft configuration created for a cluster.
- Verify that you have the required license.
- Verify that you have the required privileges for viewing host settings in the desired configuration. See [Required Privileges for Using vSphere Configuration Profiles](#).

When you create and edit a draft configuration, your changes are stored on vCenter Server. When you delete your draft configuration, the changes you made will be permanently lost. If another user creates a draft configuration for the same cluster, that draft configuration remains unaffected when you delete your draft configuration. You can preserve your current cluster configuration and continue editing the draft configuration settings only if you export the draft configuration document. For more information, see [Export the Draft Cluster Configuration or Draft Cluster Configuration Schema](#).

1. In the vSphere Client, navigate to a cluster that you manage with a single image.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Draft** tab, click the horizontal ellipses icon and select **Discard draft**.
The **Discard draft** dialog box appears.
4. Click **Discard draft**.

The draft configuration is permanently removed and you can no longer edit its settings.

Create a new draft configuration if you want to edit the cluster configuration in the vSphere Client. For more information, see [Create a Draft Configuration in vSphere Client](#).

Check Host Compliance Against the Desired Configuration for the Cluster

You can check the compliance of each host in a cluster only against the desired configuration set for the entire cluster. If a host is out of compliance with the desired configuration, you can see detailed information about the settings that make the host non-compliant.

- Verify that vCenter Server and ESXi are of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image.
- Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
- Verify that you have the required license.
- Verify that you have the required privileges for checking compliance. See [Required Privileges for Using vSphere Configuration Profiles](#).

The **Check Cluster Configuration Compliance** task runs automatically when you add a host to the cluster. You can manually run a compliance check only on the current desired configuration of the cluster.

Starting with vSphere 8.0 Update 1, you can enable vSphere Configuration Profiles on a cluster that contains hosts connected to a vSphere Distributed Switch. If you add a host connected to a vSphere Standard Switch to a cluster that contains hosts connected to a vSphere Distributed Switch, vSphere Configuration Profiles shows the added host as non-compliant. You must use the vSphere Distributed Switch workflows to remediate the host and make the host compliant with the desired cluster configuration. For more information about how to add a host to a vSphere Distributed Switch, see the *vSphere Networking* documentation.

Note: In vSphere 8.0 Update 3, vSphere Configuration Profiles can manage the vSphere Distributed Switch configuration of the hosts in a cluster managed with a desired configuration. You can also use the workflows for managing and configuring the vSphere Distributed Switches described in the *vSphere Networking* documentation.

1. In the vSphere Client, navigate to a cluster that you manage with a single image or to a cluster that you manage with baselines.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Compliance** tab, click **Check compliance**.
The **Check Cluster Configuration Compliance** task starts. After the task finishes, a list of all non-compliant hosts shows on the **Compliance** tab.
4. View the compliance results for each non-compliant host in the cluster.
 - a) On the **Compliance** tab, select a host from the **Hosts** list.
An information panel appears on the right. You see the host settings and the respective desired and current values that cause the noncompliance.

Remediate the cluster to make all non-compliant hosts compliant.

Set Up Alarms for Monitoring Configuration Drifts

You can create alarms that notify you when vSphere Configuration Profiles detects drifts from the desired cluster configuration for a cluster or a host. Receiving such notifications allows you to take the necessary actions and remediate the clusters which are out of compliance with the desired cluster configuration.

- Verify that the clusters and hosts you select have the vSphere Configuration Profiles enabled.
- Required privileges for creating alarms: **Alarms > Create alarm**.
- Verify that you have the required license.

vSphere Configuration Profiles generates events when a change in the configuration of a cluster or a host occurs. You can create alarms that activate in response of these events and notify you about the configuration drifts by sending emails or SNMP traps.

For detailed information about how to manage alarms in the vSphere Client, see the *vSphere Monitoring and Performance* documentation.

1. In the vSphere Client, navigate to a vCenter Server instance or an inventory object.
2. On the **Configure** tab, click **Alarm Definitions**.
Depending on the inventory object you select, the location of the **Alarm Definitions** node varies.
3. Click **Add** to create a new alarm definition.
The **New Alarm Definition** wizard opens.

4. On the **Name and Targets** page of the wizard, enter a name and use the **Target type** drop-down menu to select whether the alarm is triggered for clusters or hosts.
5. On the **Alarm rule 1** page of the wizard, select the event trigger and how you want to be notified about changes to the cluster or host configuration.
 - a) From the drop-down menu select the event trigger. You can choose from `com.vmware.vcIntegrity.ClusterConfigurationOutOfCompliance` or `com.vmware.vcIntegrity.HostConfigurationOutOfCompliance`.
 - a) From the severity drop-down menu, select the severity of the alarm.
 - b) To send email notifications when alarms are triggered, activate the **Send email notifications** toggle and configure the recipient addresses and the subject of the email notifications.
 - c) To send SNMP traps when alarms are triggered, activate the **Send SNMP traps** toggle.
 - d) Click **Next**.
6. On the **Reset Rule 1** page of the wizard, specify the alarm reset rule if needed and click **Next**
7. On the **Review** page of the wizard, review the alarm definition and activate the **Enable this alarm** toggle.
8. Click **Create**.

When vSphere Configuration Profiles detects that all hosts in a cluster or only some of the hosts in the cluster are out of compliance with the desired cluster configuration, an alarm is triggered. You receive email notifications or SNMP traps and can take the necessary actions to remediate the cluster and make all non-compliant hosts compliant again.

Run a Remediation Pre-Check or a Draft Pre-Check

You run a remediation pre-check to ensure that the desired configuration is valid and compatible with the current state of the hosts in the cluster, and that the cluster is in good health. You run a draft pre-check to ensure that the draft configuration is valid and can be applied on all hosts in the cluster. A pre-check that finds no issues means that you can successfully remediate the cluster and apply the desired or draft configuration to all hosts.

- Verify that vCenter Server is of version 8.0 Update 1 or later, if you want to run a remediation pre-check of the desired configuration on a cluster that you manage with a single image.
- Verify that vCenter Server is of version 8.0 Update 2 or later, if you want to run a draft pre-check on a cluster that you manage with a single image..
- Verify that you have a draft configuration created, if you want to run a draft pre-check.
- Verify that ESXi is of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image
- Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
- Verify that you have the required license.
- Verify that you have the required privileges for running a remediation or a draft pre-check task. See [Required Privileges for Using vSphere Configuration Profiles](#).

With vSphere 8.0 Update 2, you can create a draft of the cluster configuration in the vSphere Client and edit the configuration settings directly in the user interface. Before applying the draft configuration on a cluster, you can run a draft pre-check.

During a remediation or a draft pre-check, vSphere Configuration Profiles runs various checks on each host and the entire cluster to estimate and validate the impact of the desired or the draft configuration on the current state of the hosts in the cluster. The remediation and draft pre-checks also compute the remediation impact on the hosts.

With vSphere 8.0 Update 1, you can enable vSphere Configuration Profiles on a cluster that has hosts connected to a vSphere Distributed Switch. If you add a host that is connected to a vSphere Standard Switch to such cluster, the remediation pre-check returns an error and you cannot remediate the cluster before you manually add this host to the

vSphere Distributed Switch. For more information about how to add a host to a vSphere Distributed Switch, see the *vSphere Networking* documentation.

Note: In vSphere 8.0 Update 3, vSphere Configuration Profiles can manage the vSphere Distributed Switch configuration of the hosts in a cluster managed with a desired configuration. You can also use the workflows for managing and configuring the vSphere Distributed Switches described in the *vSphere Networking* documentation.

1. In the vSphere Client, navigate to a cluster that you manage with a single image or to a cluster that you manage with baselines.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. Start a pre-check by selecting one of the following options.
 - To run a remediation pre-check of the current desired configuration, on the **Compliance** tab, click **Run pre-check**.
The **Cluster Configuration Remediation Pre-Check** task starts.
 - To run a draft pre-check of the edited cluster configuration, on the **Draft** tab, click **Run pre-check**.
The **Cluster Configuration Draft Pre-check** task starts.
4. Optional: Click the **Operation Details** tab to view details about the remediation or draft pre-check operations.
The **Operation Details** tab gives summarized information or detailed information about the remediation operation, the remediation pre-check, or the draft pre-check operations. Any issues during those operations are also listed on the **Operation Details** tab.

Fix all issues that are reported during the remediation or the draft pre-check and remediate the cluster.

Remediate a Cluster Against the Desired or Draft Configuration

You remediate a cluster against the desired or draft configuration to make all non-compliant hosts compliant with the configuration for the cluster.

- Verify that vCenter Server is of version 8.0 Update 1 or later, if you want to remediate a cluster managed with a single image against the desired configuration.
- Verify that vCenter Server is of version 8.0 Update 2 or later, if you want to remediate a cluster managed with a single image against a draft configuration.
- Verify that ESXi is of version 8.0 Update 1 or later for using vSphere Configuration Profiles on a cluster that you manage with a single image.
- Verify that vCenter Server and ESXi are of version 8.0 Update 3 or later for using vSphere Configuration Profiles on a cluster that you manage with baselines.
- Verify that you have the required license.
- Verify that at least one of the hosts in the cluster is out of compliance with the desired configuration.
- Verify that you have the required privileges for performing remediation. See [Required Privileges for Using vSphere Configuration Profiles](#).

Starting with vSphere 8.0 Update 2, you can create a draft of the cluster configuration in the vSphere Client and edit the configuration settings directly in the user interface. To make all hosts in the cluster compliant with the draft configuration, you must apply it to the cluster. The apply operation first sets the draft configuration as the new desired configuration for the cluster and then remediates all hosts in the cluster against the new desired configuration.

Before the cluster is remediated, an automatic remediation or draft pre-check is triggered. As a result, the remediation impact is calculated for each host in the cluster. The remediation might require a reboot or the host to enter maintenance mode.

Starting with vSphere 8.0 Update 1, you can enable vSphere Configuration Profiles on a cluster that contains hosts connected to a vSphere Distributed Switch. You can't use vSphere Configuration Profiles to manage the vSphere Distributed Switch configuration of the hosts in the cluster. You must manually fix any drifts in the vSphere Distributed Switch networking configuration of each non-compliant host in the cluster before you start the remediation task.

Note: In vSphere 8.0 Update 3, vSphere Configuration Profiles can manage the vSphere Distributed Switch configuration of the hosts in a cluster managed with a desired configuration. You can also use the workflows for managing and configuring the vSphere Distributed Switches described in the *vSphere Networking* documentation.

1. In the vSphere Client, navigate to a cluster that you manage with a single image or to a cluster that you manage with baselines.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. Start the remediation operation.
 - To remediate the cluster against the current desired configuration, on the **Compliance** tab, click **Remediate**.
 - To remediate the cluster against the draft configuration, on the **Draft** tab, click **Apply changes**.

The **Remediate** wizard opens. A remediation pre-check task starts automatically to ensure that remediation can run without issues.

4. On the **Pre-check** page, wait for the remediation pre-check to finish, view the pre-check results and click **Next**.
5. On the **Review Impact** page, review the summary of the remediation impact on the cluster and on each host individually.
6. Click **Remediate** to close the **Remediate** wizard.
The **Remediate Cluster Configuration** task starts.
7. Optional: Click the **Operation Details** tab to view details about the remediation, the remediation pre-check, or the draft pre-check operations.

The **Operation Details** tab gives summarized information or detailed information about the remediation operation or the remediation and draft pre-check operations. Any issues during those operations are listed on the **Operation Details** tab.

All hosts in the cluster are compliant with the desired configuration.

Note:

More than one user with the required privileges can create a draft configuration for a cluster at the same time. If one of the users remediates the cluster with their draft configuration, the draft configuration becomes the new desired configuration for the cluster. In case another user attempts to remediate the cluster with their draft configuration at the same time, the remediation operation fails with an error. That user can either discard his draft configuration and start anew, or export their current draft and import it at a later time.

Backup and Restore Scenarios When Using vSphere Lifecycle Manager

Restoring a vCenter Server instance from a backup might impact the clusters in your environment in a seemingly unexpected way. Whether you use images or baselines to manage your clusters, vSphere Lifecycle Manager behaves in a specific manner during backup and restore operations.

When you back up a vCenter Server instance, you create a backup copy of all clusters in that vCenter Server instance.

Restoring vCenter Server After Switching from Baselines to Images for Cluster Lifecycle Management

Cluster A is a cluster that you manage by using baselines. You back up the vCenter Server instance where the cluster is. After the backup, you switch from using baselines to using images to manage cluster A and you remediate the cluster to apply the image to the hosts in the cluster. You now manage the lifecycle of cluster A by using a single cluster image.

If for some reason you must restore the vCenter Server instance from the backup copy you created, the restored vCenter Server instance contains cluster A. Because cluster A was managed through baselines at the time when you backed up

the vCenter Server system, the restored vCenter Server instance contains cluster A, but you must again use baselines to manage it.

Restoring vCenter Server After Remediating a Cluster Managed by an Image

After remediation, cluster A uses image X with components Y to manage all hosts in the cluster collectively. At a point in time T, you back up the vCenter Server system. Later, you remediate the cluster against a new image X+1 with new components Y+1. Now all hosts in the cluster use image X+1 with components Y+1.

If for some reason you must restore the vCenter Server system from the backup copy that you created at time T, the restored vCenter Server instance contains cluster A, but the compliance check lists the hosts in the cluster as incompatible with the image that cluster A uses. The reason for the incompatibility is that after the restore operation, cluster A reverts back to using image X with components Y, while the hosts in the cluster still run image X+1 with components Y+1. Because you cannot downgrade ESXi, to make the hosts compliant with the cluster image, you must upgrade the cluster to image X+1 with components Y+1.

vCenter Server Configuration

vCenter Server Configuration provides information about configuring VMware vCenter® Server™.

Intended Audience

This information is intended for anyone who wants to configure VMware vCenter Server®. The information is written for experienced system administrators who are familiar with virtual machine technology and data center operations.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

vSphere Client and vCenter Server Management Interface

Instructions in this guide reflect the vSphere Client, an HTML5-based GUI, and the vCenter Server Management Interface. Some additional functions can be performed using the vCenter Server appliance shell and the Direct Console User Interface.

How to Configure vCenter Server

vCenter Server is deployed using a preconfigured virtual machine, which is optimized for running VMware vCenter Server® and the associated services.

During the deployment of vCenter Server, you can create a VMware vCenter® Single Sign-On™ domain or join an existing domain. For information about the vCenter Server deployment, see *vCenter Server Installation and Setup*.

vCenter Server is supported on VMware ESXi™ 6.0 and later. The package contains the following software:

- Project Photon OS® 3.0
- PostgreSQL database
- vCenter Server 8.0 and vCenter Server 8.0 components
- Necessary services for running vCenter Server such as vCenter Single Sign-On, License service, and VMware Certificate Authority

For detailed information about authentication, see *vSphere Authentication*.

Customization of the preconfigured virtual machine is unsupported except for adding memory, CPU, and disk space.

vCenter Server has the following default user names:

- root user with the password that you set during the deployment of the virtual appliance. You use the root user to log in to the vCenter Server Management Interface and to the preconfigured virtual machine operating system.

Important: The password for the root account of vCenter Server expires after 365 days by default. For information about changing the root password and configuring the password expiration settings, see [Change the Password and Password Expiration Settings of the Root User](#).

- administrator@*your_domain_name* which is the vCenter Single Sign-On user with the password and domain name that you set during the deployment of the appliance.

When you install vCenter Server, you can change the vSphere domain. Do not use the same domain name as the domain name of your Microsoft Active Directory or OpenLDAP domain name.

Initially, only the user administrator@*your_domain_name* has the privileges to log in to the vCenter Server system. By default, the administrator@*your_domain_name* user is a member of the SystemConfiguration.Administrators group. This user can add an identity source in which additional users and groups are defined to vCenter Single Sign-On or give permissions to the users and groups. For more information, see *vSphere Security*.

You can configure vCenter Server settings in four ways:

- Use the vCenter Server Management Interface.
You can edit the system settings such as access, network, time synchronization, and the root password settings. This is the preferred way for configuring vCenter Server.
- Use the vSphere Client.
You can navigate to the system configuration settings of vCenter Server and join the deployment to an Active Directory domain. You can manage the services that are running in vCenter Server and modify various settings such as access, network, and firewall settings.
- Use the Bash shell.
You can use TTY1 to log in to the console or can use SSH and run configuration, monitoring, and troubleshooting commands in vCenter Server.
- Use the Direct Console User Interface.
You can use TTY2 to log in to the vCenter Server Direct Console User Interface to change the password of the root user, configure the network settings, or enable access to the Bash shell or SSH.

What Happened to the Platform Services Controller

In vSphere 7.0, all Platform Services Controller services are consolidated into vCenter Server.

Beginning in vSphere 7.0, deploying or upgrading vCenter Server in vSphere 7.0 requires the use of vCenter Server Appliance, a preconfigured virtual machine optimized for running vCenter Server. The new vCenter Server contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, and licensing. It is no longer necessary nor possible to deploy and use an external Platform Services Controller. All Platform Services Controller services are consolidated into vCenter Server, and deployment and administration are simplified.

As these services are now part of vCenter Server, they are no longer described as a part of Platform Services Controller. In vSphere 7.0, the *vSphere Authentication* publication replaces the *Platform Services Controller Administration* publication. The new publication contains complete information about authentication and certificate management. For information about upgrading or migrating from vSphere 6.5 and 6.7 deployments using an existing external Platform Services Controller to vSphere 7.0 using vCenter Server Appliance, see the *vSphere Upgrade* documentation.

Using the vCenter Server Management Interface to Configure vCenter Server

After you deploy vCenter Server, you can log in to the vCenter Server Management Interface and edit the settings.

For information about patching vCenter Server and enabling automatic checks for vCenter Server patches, see the *vSphere Upgrade* documentation.

For information backing up and restoring vCenter Server, see *vCenter Server Installation and Setup*.

Log In to the vCenter Server Management Interface

Log in to the vCenter Server Management Interface to access the vCenter Server configuration settings.

Verify that the vCenter Server is successfully deployed and running.

Note: The login session expires if you leave the vCenter Server Management Interface idle for 10 minutes.

1. In a Web browser, go to the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
2. Log in as root.
The default root password is the password that you set while deploying vCenter Server.

View vCenter ServerHealth Status

You can use the vCenter Server Management Interface to view the overall health status of vCenter Server and health messages.

Log in to the vCenter Server Management Interface as root.

The overall health status of vCenter Server is based on the status of the hardware components such as CPU, memory, database, and storage. It is also based on the update component, which shows whether the software packages are up to date according to the last check for available patches.

Important: If you do not perform regular checks for available patches, the health status of the update component might become out-of-date. For information about how to check for vCenter Server patches and enable automatic checks for vCenter Server patches, see *vSphere Upgrade*.

1. In the vCenter Server Management Interface, click **Summary**.
2. In the Health Status pane, view the Overall Health badge.

Table 130: Health Status

Badge Icon	Description
	Good. All components are healthy.
	Warning. One or more components might become overloaded soon. View the details in the Health Messages pane.
	Alert. One or more components might be degraded. Nonsecurity patches might be available. View the details in the Health Messages pane.
	Critical. One or more components might be in an unusable status and vCenter Server might become unresponsive soon. Security patches might be available. View the details in the Health Messages pane.
	Unknown. No data is available.

Reboot or Shut Down vCenter Server

You can use the vCenter Server Management Interface to restart or power off the virtual machine running.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Summary**.
2. From the top menu pane, click the **Actions** drop-down menu.
3. Click **Reboot** or **Shutdown** to restart or power off the virtual machine.
4. In the confirmation dialog box, click **Yes** to confirm the operation.

Create a Support Bundle

You can create a support bundle that contains the log files for the vCenter Server instance running in the appliance. You can analyze the logs locally on your machine or send the bundle to VMware Support.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Summary**.
2. From the top menu pane, click the **Actions** drop-down menu.
3. Click **Create Support Bundle** and save the bundle on your local machine.

The support bundle is downloaded as a `.tgz` file on your local machine.

Monitor CPU and Memory Use

You can use the vCenter Server Management Interface to monitor the overall CPU and memory use of vCenter Server.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Monitor**.
2. On the Monitor page, click the **CPU & Memory** tab.
3. From the **date range** drop-down menu, select the time period for which you want to generate a CPU utilization trending graph and a memory utilization trending graph.
4. Point to the graphs to see the CPU and memory use for a particular date and time.

Monitor Disk Use

You can use the vCenter Server Management Interface to monitor the disk use of vCenter Server.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Monitor**.
2. On the Monitor page, click the **Disks** tab.

The Monitor Disks pane shows a disk, sortable by name, partition, or utilization.

Monitor Network Use

You can use the vCenter Server Management Interface to monitor the network use of vCenter Server in the last day, week, month, or quarter.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Monitor**.
2. On the Monitor page, click the **Network** tab.
3. From the **date range** drop-down menu, select the time period for generating the network utilization graph.
4. From the table below the graph grid, select a packet or transmit byte rate to monitor.

The options vary depending on your network settings.

The network utilization graph refreshes to display the use of the item you select.

5. Point to the network utilization graph to see the network use data for a particular date and time.

Monitor Database Use

You can use the vCenter Server Management Interface to monitor the use of the embedded database of the vCenter Server by data type. You can also monitor space use trending graphs and filter any of the largest data types.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Monitor**.
2. On the Monitor page, click the **Database** tab to monitor the consumed and free space for the vCenter Server database.
3. From the **date range** drop-down menu, select the time period for which you want to generate the space utilization trending graphs.
4. At the base of the graph, click the title a particular database component to include or exclude that component from the graph.

Option	Description
Seat space utilization trend graph	Allows you to select and view alarm, event, task, and stat trendlines.
Overall space utilization trend graph	Allows you to select and view SEAT, DB Log, and core trendlines.

5. Point to the space utilization graph to see the database use value for a particular date and time.

Enable or Disable SSH and Bash Shell Access

You can use the vCenter Server Management Interface to edit the access settings for the appliance.

Log in to the vCenter Server Management Interface as root.

You can enable or disable an SSH administrator login to the appliance. You can also enable access to the vCenter Server Bash shell for a specific time interval.

1. In the vCenter Server Management Interface, click **Access**, and click **Edit**.
2. Edit the access settings for vCenter Server.

Option	Description
Enable SSH login	Enables SSH access to vCenter Server.
Enable DCUI	Enables DCUI access to vCenter Server.
Enable Console CLI	Enables console CLI access to vCenter Server.
Enable Bash Shell	Enables Bash shell access to vCenter Server for the number of minutes that you enter.

3. Click **OK** to save the settings.

Configure the DNS, IP Address, and Proxy Settings

You can set the combination of static and DHCP, IPv4 and IPv6 addresses, edit the DNS settings, and define the proxy settings for vCenter Server.

- To change the IP address of the appliance, verify that the system name of the appliance is an FQDN. The system name is used as a primary network identifier. If you set an IP address as a system name during the deployment of the appliance, you can later change the PNID to an FQDN.

Note: You can set only IPv4 IP address as system name. The IPv4 IP address must be enabled before this setting.

- To restore a dual stack VC, the base VC after stage 1 deployment should be configured as:
 - If PNID of backed up VC resolves to IPv4 and IPv4 is configured static, the base VC during stage1 should have static or DHCP IPv4 configured.
 - If PNID of backed up VC resolves to IPv4 and IPv4 is configured DHCP, the base VC during stage1 should have DHCP IPv4 configured.
 - If PNID of backed up VC resolves to IPv6 and IPv6 is configured static, the base VC during stage1 should have static or DHCP IPv6 configured.
 - If PNID of backed up VC resolves to IPv6 and IPv6 is configured DHCP, the base VC during stage1 should have DHCP IPv6 configured.
- Log in to the vCenter Server Management Interface as root.

- In the vCenter Server Management Interface, click **Networking**.
- From the Network Settings page, click **Edit**.
- Expand the Hostname and DNS section to configure the DNS settings.

Note: Valid value for Hostname is either FQDN resolving to enabled IP address or IPv4 IP address.

Option	Description
Obtain DNS settings automatically	Obtains the DNS settings automatically from the network.
Enter DNS settings manually	Lets you set the DNS address settings manually. If you select this option, you must provide the following information: - The IP address of the preferred DNS server. - The IP address of the alternate DNS server.

- From the Network Settings page, click **Edit**.
- Expand the NIC 0 section to configure the gateway settings.

Note: PNID and network APIs support only NIC 0 as primary NIC.

Recommended steps to configure network other than NIC 0, when vCenter has multiple NICs configured are:

- NIC 1 is reserved for VCHA and should not be used in multi-NIC configurations.
- As mentioned in the UI specifications, DGW should be set to NIC 0.
- As access to vSphere Client is redirected to PNID, only access to NIC 0 is supported or recommended even in multi-NIC configurations.

- Edit the IPv4 address settings.

Option	Description
Enable or Disable IPv4 settings	Enables or Disables the IPv4 address based on the toggle switch option.
Obtain IPv4 settings automatically	Obtains the IPv4 address for the appliance automatically from the network.
Enter IPv4 settings manually	Uses an IPv4 address that you set manually. You must enter the IP address, subnet prefix length, and the default gateway. Note: For static IPv4 or IPv6 addresses, DNS server must be set manually.

Note: Second and third party solutions need to re-register, when there is a change in the IPv4 and IPv6 IP addressess.

7. Edit the IPv6 address settings.

Option	Description
Enable or Disable IPv6 settings	Enables or disables the IPv6 address based on the toggle switch option.
Obtain IPv6 settings automatically through DHCP	Assigns IPv6 addresses to the appliance automatically from the network by using DHCP.
Obtain IPv6 settings automatically through router advertisement	Assigns IPv6 addresses to the appliance automatically from the network by using router advertisement.
Use static IPv6 addresses	Uses static IPv6 addresses that you set up manually. 1. Click the checkbox. 2. Enter the IPv6 address and the subnet prefix length. 3. Click Add to enter additional IPv6 addresses. 4. Click Save. Note: For static IPv4 or IPv6 addresses, DNS server must be set manually.

You can configure the appliance to obtain the IPv6 settings automatically through both DHCP and router advertisement. You can assign static a IPv6 address at the same time.

Note: Second and third party solutions need to re-register, when there is a change in the IPv4 and IPv6 IP addresses.

8. To configure a proxy server, in the Proxy Settings section, click **Edit**.

9. Select the proxy setting to enable

Option	Description
HTTPS	Enable to configure the HTTPS proxy settings.
FTP	Enable to configure the FTP proxy settings. Note: Ensure that the ICMP is enabled on the proxy server.
HTTP	Enable to configure the HTTP proxy settings.

10. Optional: To add proxy bypass list entries, enter the IPv4 addresses, FQDNs, or domain names that you want to bypass as a comma separated list.

11. Enter the server hostname or IP address.

12. Enter the port.

13. Enter the username (optional).

14. Enter the password (optional).

15. Click **Save**.

Reconfigure the Primary Network Identifier

You can change the FQDN, IP, or PNID of the management network of vCenter Server.

The system name is used as a primary network identifier. If you set an IP address as a system name during the deployment of the appliance, you can later change the PNID to an FQDN. You cannot change an FQDN back to an IP address.

If vCenter High Availability (HA) is enabled, you must deactivate the vCenter HA setup before reconfiguring the PNID.

1. Log in to the vCenter Server Management Interface using your administrator SSO credentials.
2. In the vCenter Server Management Interface, navigate to the **Networking** page and click **Edit**.
3. Select the NIC to be modified and click **Next**.
4. In the **Edit Settings** pane, change the host name and provide the new IP address. Click **Next**.
5. In the **SSO Credentials** pane, provide the administrator SSO credentials. You must use `administrator@<domain_name>` credentials.
6. In the **Ready to Complete** pane, review your new settings and check the backup acknowledgement box. Click **Finish**.
A taskbar shows the status of the network update. To cancel the update, click **Cancel Network Update**. When the network reconfiguration is complete, the UI redirects to the new IP address.
7. To finish the reconfiguration process and restart services, log in using your administrator SSO credentials.
8. On the **Networking** page, verify the new host name and IP address.
 - Re-register all deployed plug-ins.
 - Regenerate all custom certificates.
 - If vCenter HA was enabled, reconfigure vCenter HA.
 - If Active Domain was enabled, reconfigure Active Domain.
 - If Hybrid Link mode was enabled, reconfigure Hybrid Link with the Cloud vCenter Server.

Edit the Firewall Settings

After you deploy vCenter Server, you can edit its firewall settings and create firewall rules using the Management Interface.

Verify that the user who logs in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in vCenter Single Sign-On.

You can set up firewall rules to accept or block traffic between vCenter Server and specific servers, hosts, or virtual machines. You cannot block specific ports, you block all the traffic.

1. In the vCenter Server Management Interface, click **Firewall**.
2. Edit the firewall settings.

Command	Action
Add	1. To create a firewall rule, click Add. 2. Select a network interface of the virtual machine. 3. Enter the IP address of the network to apply this rule to. The IP address can be IPv4 and IPv6 address. 4. Enter a subnet prefix length. If you only block one IP address, use a 32-bit subnet mask. 5. From the Action drop-down menu, select whether to Accept, Ignore, Reject, or Return the connection between vCenter Server and the network that you entered. 6. Click Save.
Edit	1. Select a rule and click Edit. 2. Edit the settings of the rule. 3. Click Save.
Delete	1. Select a rule and click Delete. 2. At the prompt, click Delete again.
Reorder	1. Select a rule and click Reorder.

Command	Action
	2. In the Reorder pane, select the rule to move. 3. Click Move Up or Move Down. 4. Click Save.

Configure the System Time Zone and Time Synchronization Settings

After you deploy vCenter Server, you can change the system time zone and time synchronization settings from the vCenter Server Management Interface.

Log in to the vCenter Server Management Interface as root.

When you deploy vCenter Server, you either use the time settings of the ESXi host on which vCenter Server is running or you configure the time synchronization based on an NTP server. If the time settings in your vSphere network change, you can edit the time zone and time synchronization settings in the appliance.

1. In the vCenter Server Management Interface, click **Time**.
2. Configure the system time zone settings.
 - a) In the Time zone pane, click **Edit**.
 - b) From the **Time zone** drop-down menu, select a location or time zone and click **Save**.
3. Configure the time synchronization settings.
 - a) In the Time synchronization pane, click **Edit**.
 - b) From the **Mode** drop-down menu, configure the time synchronization method.

Option	Description
Disabled	No time synchronization. Uses the system time zone settings.
Host	Enables VMware Tools time synchronization. Uses VMware Tools to synchronize the time of the appliance with the time of the ESXi host.
NTP	Enables NTP synchronization. You must enter the IP address or FQDN of one or more NTP servers.

- c) Click **Save**.

Start, Stop, and Restart Services

You can use the vCenter Server Management Interface to view the status of vCenter Server components and to start, stop, and restart services.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Services**.
The Services page displays a table of the installed services that can be sorted by name, startup type, health, and state.
2. Select a service and click **SET STARTUP TYPE** to configure either Manual or Automatic start of a service.
3. Select a service and click **START** to start a service.
4. Select a service and click **STOP** to stop, or **RESTART** to restart a service, then click **OK**.

Warning: Stopping or Restarting some services may lead to functionality becoming temporarily unavailable.

Configure Update Settings

You can use the vCenter Server Management Interface to configure your update settings and check for new updates.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Update**.
2. To configure your update settings, click **Settings**.
 - a) To check for updates automatically, select the checkbox.
 - b) Select to use a default or custom repository.
 - c) If you selected a custom repository, enter the repository URL, username (optional), and password (optional). Click **Save**.
For the URL, HTTPS and FTPS protocols are supported.
3. To manually check for updates, click the **Check Updates** drop-down menu.
 - a) Select to check a **CD-ROM** or a **CD-ROM + URL** for updates.

The Available Updates table displays available updates sortable by version, type, release date, reboot requirement, and severity.

Change the Password and Password Expiration Settings of the Root User

When you deploy vCenter Server, you set the initial password of the root user, which expires after 90 days by default. You can change the root password and the password expiration settings from the vCenter Server Management Interface.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Administration**.
2. In the Password section, click **Change**.
3. Enter the current password and the new password, then click **Save**.
4. Configure the password expiration settings for the root user.
 - a) In the Password expiration settings section, click **Edit** and select the password expiration policy.

Option	Description
Yes	The password of the root user expires after a specified number of days. You must provide the following information: • Root password validity (days) The number of days after which the password expires. • Email for expiration warning The email address to which vCenter Server sends a warning message before the expiration date.

Option	Description
No	The password of the root user never expires.

- b) In the Password expiration settings pane, click **Save** to apply the new password expiration settings.



vCenter Server Management

Summary

Monitor

Access

Networking

Firewall

Time

Services

Update

Pas

Pas

The Password expiration settings section displays the new expiration date.

Forward vCenter Server Log Files to Remote Syslog Server

You can forward the vCenter Server log files to a remote syslog server to conduct an analysis of your logs.

Log in to the vCenter Server Management Interface as root.

Note:

ESXi can be configured to send log files to a vCenter Server rather than storing them to a local disk. The recommended maximum numbers of supported hosts to collect logs from is 30. See [KB 2003322](#) for information on how to configure ESXi log forwarding. This feature is intended for smaller environments with stateless ESXi hosts. For all other cases, use a dedicated log server. Using vCenter Server to receive ESXi log files might impact vCenter Server performance.

1. In the vCenter Server Appliance Management Interface, select **Syslog**.
2. In the Forwarding Configuration section, click **Configure** if you have not configured any remote syslog hosts. Click **Edit** if you already have configured hosts.
3. In the Create Forwarding Configuration pane, enter the server address of the destination host. The maximum number of supported destination hosts is three.
4. From the **Protocol** drop-down menu, select the protocol to use.

Menu Item	Description
TLS	Transport Layer Security
TCP	Transmission Control Protocol
REL	Reliable Event Logging Protocol
UDP	User Datagram Protocol

5. In the **Port** text box, enter the port number to use for communication with the destination host.
6. In the Create Forwarding Configuration pane, click **Add** to enter another remote syslog server.
7. Click **Save**.
8. Verify that the remote syslog server is receiving messages.
9. In the Forwarding Configuration section, click **Send Test Message**.
10. Verify on the remote syslog server that the test message was received.
The new configuration settings are shown in the Forwarding Configuration section.
11. Optional: To view the logs that are exported, log in to the vCenter Server Appliance Management Interface as root user and view the list of log files in `/etc/vmware-syslog/` directory.

Configure and Schedule Backups

You can use the vCenter Server Management Interface to set a backup location, create a backup schedule, and monitor backup activity.

Log in to the vCenter Server Management Interface as root.

1. In the vCenter Server Management Interface, click **Backup**.
2. To create a backup schedule, click **Configure**. To edit an existing backup schedule, click **Edit**.
 - a) In the **Backup Schedule** pane, enter the backup location using the format `protocol://server-address[:port-number]/folder/subfolder`.
Supported protocols for backup are FTPS, HTTPS, SFTP, FTP, NFS, SMB, and HTTP.
 - b) Enter the username and password for the backup server.
 - c) Enter the time and frequency for the backup to occur.
 - d) (Optional) Enter an encryption password for the backup.
 - e) Indicate the number of backups to retain.
 - f) Indicate the types of data to be backed up.
3. To initiate a manual backup, click **Backup Now**.

Information for scheduled and manual backups is displayed in the **Activity** table.

Using the vSphere Client to Configure vCenter Server

You can perform some configuration operations from the vSphere Client such as joining the appliance to an Active Directory domain, networking, and other settings.

Configuring vCenter Server

You can configure vCenter Server from the vSphere Client and from the vCenter Server Management Interface.

What you can do depends on your deployment.

On-premises vCenter Server

You can change many of the vCenter Server settings, including licensing, statistics collection, logging, and more.

For more information about how to assign a license to your vCenter Server instance, see the [How to Configure License for Your vCenter Server](#) article in the *vCenter Server and Host Management* documentation.

vCenter Server in VMware Cloud on AWS

VMware preconfigures vCenter Server instances when you create an SDDC. You can view configuration settings and advanced settings, and you can set a Message of the Day.

Configuring Statistics Settings

To set up how statistical data is recorded, you configure collection intervals for statistics. You can access the stored statistical information through command-line monitoring utilities or by viewing performance charts in the vSphere Client.

Configure Statistics Collection Intervals in the vSphere Client

Statistic collection intervals determine the frequency at which statistic queries occur, the length of time statistical data is stored in the database, and the type of statistical data that is collected. You can view the collected statistics through the performance charts in the vSphere Client or through command-line monitoring utilities.

Required privilege: **Performance > ModifyIntervals**

Note: Not all interval attributes are configurable.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. To enable or disable a statistics interval, check the box for that interval.
6. To change a statistics interval attribute value, select a value from the drop-down menu.
 - a) In **Interval Duration**, select the time interval in which statistics data is collected.
 - b) In **Save For**, select for how long the archived statistics are kept in the database.
 - c) In **Statistics Level**, select a new level for collecting statistics.
 The lower the level is, the fewer number of statistic counters are used. Level 4 uses all statistics counters. Use it only for debugging purposes.

 The statistics level must be less than or equal to the statistics level that is set for the preceding statistics interval. This requirement is a vCenter Server dependency.
7. Optional: In Database Size, estimate the effect of the statistics settings on the database.
 - a) Enter the number of **Physical Hosts**.
 - b) Enter the number of **Virtual Machines**.
 The estimated space required and number of database rows required are calculated and displayed.
 - c) If necessary, make changes to your statistics collection settings.
8. Click **Save**.

Relationships Between the Default Settings for Statistics Intervals

- Samples that are collected every 5 minutes are stored for 1 day.
- Samples that are collected every 30 minutes are stored for 1 week.
- Samples that are collected every 2 hours are stored for 1 month.
- Samples that are collected on 1 day are stored for 1 year.

For all statistics intervals, the default level is 1. It uses the Cluster Services, CPU, Disk, Memory, Network, System, and Virtual Machine Operations counters.

Data Collection Levels

Each collection interval has a default collection level that determines the amount of data gathered and which counters are available for display in the charts. Collection levels are also referred to as statistics levels.

Table 131: Statistics Levels

Level	Metrics	Best Practice
Level 1	• Cluster Services (VMware Distributed Resource Scheduler) – all metrics • CPU – cpuentitlement, totalmhz, usage (average), usagemhz • Disk – capacity, maxTotalLatency, provisioned, unshared, usage (average), used	Use for long-term performance monitoring when device statistics are not required. Level 1 is the default Collection Level for all Collection Intervals.

Level	Metrics	Best Practice
	- Memory – consumed, mementitlement, overhead, swapinRate, swapoutRate, swapused, totalmb, usage (average), vmmemctl (balloon), totalbandwidth (DRAM or PMem) - Network – usage (average), IPv6 - System – heartbeat, uptime - Virtual Machine Operations – numChangeDS, numChangeHost, numChangeHostDS	
Level 2	- Level 1 metrics - CPU – idle, reservedCapacity - Disk – All metrics, excluding numberRead and numberWrite. - Memory – All metrics, excluding memUsed, maximum and minimum rollup values, read or write latency (DRAM or PMem). - Virtual Machine Operations – All metrics	Use for long-term performance monitoring when device statistics are not required but you want to monitor more than the basic statistics.
Level 3	- Level 1 and Level 2 metrics - Metrics for all counters, excluding minimum and maximum rollup values. - Device metrics	Use for short-term performance monitoring after encountering problems or when device statistics are required.
Level 4	All metrics supported by the vCenter Server, including minimum and maximum rollup values.	Use for short-term performance monitoring after encountering problems or when device statistics are required.

Note:

When the statistics levels, level 3 or level 4 are used beyond the default value, it may cause one particular process, vpxd, to sustain memory growth, if it cannot save the statistics information to the database as quickly as required. If the usage limit of these statistics levels is not monitored closely, it may cause vpxd to grow out of memory and eventually crash.

So, in case the administrator decides to elevate any of these levels, it is necessary for the administrator to monitor the size of the vpxd process to make sure that is not growing boundlessly after the change.

Configure Runtime Settings for vCenter Server

You can change the vCenter Server ID, managed address, and name. You might need to make changes if you run multiple vCenter Server systems in the same environment.

Required privilege: **Global > Settings**

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. In the Edit vCenter Server Settings dialog box, select **Runtime Settings**.
6. In **vCenter Server unique ID**, enter a unique ID.
You can change this value to a number from 0 through 63 to identify each vCenter Server system running in a common environment. By default, an ID value is generated randomly.
7. In **vCenter Server managed address**, enter the vCenter Server system address.
The address can be IPv4, IPv6, a fully qualified domain name, an IP address, or another address format.

8. In **vCenter Server name**, enter the name of the vCenter Server system.
If you change the DNS name of the vCenter Server, you can use this text box to modify the vCenter Server name to match.
9. Click **Save**.

If you made changes to the vCenter Server system unique ID, you must restart the vCenter Server system for these changes to take effect.

Configure User Directory Settings

You can configure some of the ways vCenter Server interacts with the user directory server that is configured as an identity source.

Required privilege: **Global > Settings**

For vCenter Server versions before vCenter Server 5.0, these settings apply to an Active Directory associated with vCenter Server. For vCenter Server 5.0 and later, these settings apply to vCenter Single Sign-On identity sources.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. In the Edit vCenter general settings window, select **User directory**.
6. In **User directory timeout**, type the timeout interval in seconds for connecting to the directory server.
7. Enable the **Query Limit** box to set a query limit size.
8. In **Query Limit Size**, enter the number of users and groups for which you can associate permissions on the child inventory objects of the vCenter Server system.
Note: You can associate permissions with users and groups from the Add Permissions dialog box that displays when you click **Add permissions** in **Manage > Permissions** for a vSphere inventory object.
9. Click **SAVE**.

Configure Mail Sender Settings

You must configure the email address of the sender account to use vCenter Server operations, such as sending email notifications as alarm actions. You can send email alerts and alarms using either anonymous or authentication mode.

Required privilege: **Global > Settings**

SMTP authentication is available for:

- vSphere 7.0 Update 1 and later versions only.

- Office 365 mailbox users only.
- SMTP mail sender should meet the basic requirements for SMTP Authentication as mentioned in Microsoft's document [Requirement for SMTP AUTH client submission](#).

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. Select **Mail** and enter the settings that vCenter Server uses to send email alerts.

6. In **Mail server** text box, enter the SMTP server information.

The SMTP server is the DNS name or IP address of the SMTP gateway to use for sending email messages.

- For sending mails anonymously, you can enter any SMTP server information as the mail server name.
- For SMTP authentication, enter the mail server name as `smtp.office365.com`, unless you have some customized configuration.

Note: Do not use IP address as mail server as IP address is not supported for SMTP authentication.

7. In **Mail sender** text box, enter the sender account information.

The sender account is the email address of the sender.

For SMTP authentication, you must enter a valid SMTP account name in the **Mail sender** text box.

Note: You must enter the full email address, including the domain name.

For example, `mail_server@example.com`

8. Click **SAVE**.
9. This step is applicable only for SMTP authentication.
You must configure the SMTP user settings as below:
 1. Select the **Configure** tab.
 2. Select **Advanced Settings**.
 3. Click **EDIT SETTINGS** and enter the following values for the configuration parameters:

Name	Value
<code>mail.smtp.username</code>	Valid SMTP account name.
<code>mail.smtp.password</code>	Valid SMTP account password. Note: Currently, the account password is not masked and is visible. You must use a dedicated SMTP mail user until the masking is available in upcoming releases.
<code>mail.smtp.port</code>	587

4. Click **SAVE**.

You can perform the following steps to test the mail settings:

1. Create an alarm that is triggered by a user action.
For example, the user action can be powering off a virtual machine.
2. Verify that you receive an email, when the alarm is triggered.

Configure SNMP Settings

You can configure up to four receivers to receive SNMP traps from vCenter Server. For each receiver, specify a host name, port, and community.

Required privilege: **Global > Settings**

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. Select **SNMP receivers**.
6. Select the **Enable receiver 1** box.
7. In **Primary Receiver URL**, enter the host name or IP address of the SNMP receiver.
8. In **Receiver port**, enter the port number of the receiver.
The port number must be a value between 1 and 65535.
9. In **Community string**, enter the community identifier.
10. To send alarms to multiple receivers, select the additional **Enable receiver** boxes and enter the host name, port number, and community identifier for those receivers.
11. Click **Save**.

View Port Settings

You can view the ports used by the Web service to communicate with other applications. You cannot configure these port settings.

The Web service is installed as part of the VMware vCenter Server installation. The Web service is a required component for third-party applications that use the VMware SDK application programming interface (API). For information about installing the Web service, see the *vCenter Server Installation and Setup* documentation.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. Select **Ports**.
The ports used by the Web service are displayed.
6. Click **Save**.

Configure Timeout Settings

You can configure the timeout intervals for vCenter Server operations. These intervals specify the amount of time after which the vSphere Client times out.

Required privilege: **Global > Settings**

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **Advanced Settings**.
4. In **client.timeout.normal**, enter the timeout interval in seconds for normal operations.
The value cannot be less than five (5).
5. In **client.timeout.long**, enter the timeout interval in minutes for long operations.
The value cannot be less than five (5).
6. Click **Save**.
7. Restart the vCenter Server system for the changes to take effect.

Configure Logging Options

You can configure the amount of detail that vCenter Server collects in log files.

Required privilege: **Global > Settings**

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. Select **Logging settings**.
6. Select the logging options.

Option	Description
None	Turns off logging
Error	Displays only error log entries
Warning	Displays warning and error log entries
Info	Displays information, error, and warning log entries
Verbose	Displays information, error, warning, and verbose log entries
Trivia	Displays information, error, warning, verbose, and trivia log entries

7. Click **SAVE**.

Changes to the logging settings take effect immediately. You do not need to restart the vCenter Server system.

Configure Database Settings

You can configure the maximum number of database connections that can occur simultaneously. To limit the growth of the vCenter Server database and save storage space, you can configure the database to discard information about tasks or events periodically.

Note: Do not use the database retention options, if you want to keep a complete history of tasks and events for your vCenter Server.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. In Edit vCenter general settings window, click **Database**.
6. In **Maximum connections** field, enter the required number of connections.

Note: Do not change this value unless one of these issues exists in your system.

- If your vCenter Server system performs many operations frequently and performance is critical, increase the number of connections.
 - If the database is shared and connections to the database are costly, decrease the number of connections.
7. Enable **Task cleanup** option for vCenter Server to delete the retained tasks periodically.
 8. Optional: In **Task retention (days)** field, enter a value (in days).
Information about tasks that are performed on this vCenter Server system is discarded after the specified number of days.
 9. Enable **Event cleanup** option for vCenter Server to clean up the retained events periodically.
 10. Optional: In **Event retention (days)** field, enter a value (in days).

Information about events for this vCenter Server system is discarded after the specified number of days.

Note: Monitor vCenter Server database consumption and disk partition in the vCenter Server Management Interface.

Warning: Increasing the events retention to more than 30 days results in significant increase of vCenter database size and can shutdown the vCenter Server. Ensure that you increase the vCenter database accordingly.

11. Click **SAVE**.
12. Restart the vCenter Server to apply changes manually.

How Do You Verify SSL Certificates for Legacy ESXi Hosts

You can configure vCenter Server to check the SSL certificates of hosts to which it connects. If you configure this setting, vCenter Server and the vSphere Client check for valid SSL certificates before connecting to a host for operations such as adding a host or making a remote console connection to a virtual machine.

Starting with vCenter Server 6.0, SSL certificates are signed by VMware Certificate Authority by default. You can instead use certificates from a third-party CA. Thumbprint mode is supported only for legacy hosts.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. Select **SSL settings**.
6. Determine the host thumbprint for each legacy host that requires validation.
 - a) Log in to the direct console.
 - b) Select **View Support Information** on the **System Customization** menu.

The thumbprint is displayed in the column on the right.
7. Compare the thumbprint you obtained from the host with the thumbprint listed in the vCenter Server SSL settings dialog box.
8. If the thumbprints match, select the check box for the host.

Hosts that are not selected will be disconnected after you click **Save**.
9. Click **Save**.

Configure Advanced Settings

In **Advanced Settings**, you can modify the vCenter Server configuration file, `vpzd.cfg`.

Required privilege: Global > Settings

You can use **Advanced Settings** to add entries to the `vpzd.cfg` file, but not to edit or delete them. VMware recommends that you change these settings only when instructed to do so by VMware technical support or when you are following specific instructions in VMware documentation.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Select **Advanced Settings**.
4. Click **Edit Settings**.
5. In the **Name** field, type a name for the setting. Names must start with 'config.' For example: `config.log`.
6. In the **Value** field, type the value for the specified setting.
7. Click **Add**.
8. Click **Save**.

Newly added advanced settings have `config.` appended to the setting keys in the `vpzd.cfg` file. For example:

```
config.example.setting = exampleValue
```

Many advanced settings changes require you to restart the vCenter Server system. Consult VMware technical support to determine if your changes require a restart.

Send a Message to Other Logged In Users

Administrators can send messages to users who are currently logged in to a vCenter Server system. The message might announce maintenance or ask users to log out temporarily.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Click **Configure**.
3. Select **Settings > Message of the Day** and click **Edit**.
4. Enter a message and click **OK**.

The message appears at the top of the vSphere Client in each active user session.

Join or Leave an Active Directory Domain

You can join vCenter Server to an Active Directory domain. You can attach the users and groups from this Active Directory domain to your vCenter Single Sign-On domain. You can leave the Active Directory domain.

- Verify that the user who logs in to the vCenter Server instance is a member of the `SystemConfiguration.Administrators` group in vCenter Single Sign-On.
- Verify that the system name of the appliance is an FQDN. If, during the deployment of the appliance, you set an IP address as a system name, you cannot join vCenter Server to an Active Directory domain.

Important: Joining vCenter Server to an Active Directory domain with a read-only domain controller (RODC) is not supported. You can join vCenter Server only to an Active Directory domain with a writable domain controller.

If you want to configure permissions so that users and groups from an Active Directory can access the vCenter Server components, you must join the vCenter Server instance to the Active Directory domain.

For example, to enable an Active Directory user to log in to the vCenter Server instance by using the vSphere Client, you must join the vCenter Server instance to the Active Directory domain and assign the Administrator role to this user.

1. Use the vSphere Client to log in as `administrator@your_domain_name` to the vCenter Server instance.
2. From the vSphere Client menu, select **Administration**.
3. Select **Single Sign On > Configuration**.
4. Click on the **Identity Provider** tab, and select **Active Directory Domain** as the identity provide type.
5. Click **JOIN AD**.
6. In the Join Active Directory Domain window, provide the following details.

Option	Description
Domain	Active Directory domain name, for example, mydomain.com. Do not provide an IP address in this text box.
Organization Unit (optional)	The full Organization Unit (OU) LDAP FQDN, for example, OU=Engineering,DC=mydomain,DC=com. Important: Use this text box only if you are familiar with LDAP.
Username	User name in User Principal Name (UPN) format, for example, jchin@mydomain.com. Important: Down-level login name format, for example, DOMAIN\UserName, is unsupported.
Password	Password of the user.

Note: Reboot the node to apply changes.

7. Click **JOIN** to join the vCenter Server to the Active Directory domain.
The operation silently succeeds and you can see the Join AD option turned to Leave AD.
8. Optional: To leave the Active Directory Domain, click **LEAVE AD**.
9. Restart the vCenter Server to apply the changes.
Important: If you do not restart the vCenter Server, you might encounter problems when using the vSphere Client.
10. Select **Identity Sources** tab, and click the **ADD**.
 1. In the Add Identity Source window, select **Active Directory (Integrated Windows Authentication)** as the Identity Source Type.
 2. Enter the identity source settings of the joined Active Directory domain, and click **ADD**.

Table 132: Add Identity Source Settings

Text Box	Description
Domain name	FDQN of the domain. Do not provide an IP address in this text box.
Use machine account	Select this option to use the local machine account as the SPN. When you select this option, you specify only the domain name. Do not select this option if you expect to rename this machine.
Use Service Principal Name (SPN)	Select this option if you expect to rename the local machine. You must specify an SPN, a user who can authenticate with the identity source, and a password for the user.

Text Box	Description
Service principal name	SPN that helps Kerberos to identify the Active Directory service. Include the domain in the name, for example, STS/example.com. You might have to run <code>setspn -S</code> to add the user you want to use. See the Microsoft documentation for information on <code>setspn</code> . The SPN must be unique across the domain. Running <code>setspn -S</code> checks that no duplicate is created.
Username	Name of a user who can authenticate with this identity source. Use the email address format, for example, jchin@mydomain.com. You can verify the User Principal Name with the Active Directory Service Interfaces Editor (ADSI Edit).
Password	Password for the user who is used to authenticate with this identity source, which is the user who is specified in User Principal Name. Include the domain name, for example, jdoe@example.com.

On the **Identity Sources** tab, you can see the joined Active Directory domain.

You can configure permissions so that users and groups from the joined Active Directory domain can access the vCenter Server components. For information about managing permissions, see the *vSphere Security* documentation.

Add a User to the SystemConfiguration.BashShellAdministrators Group

To enable access to the appliance Bash shell by using the vSphere Client, the user you use to log in must be a member of the SystemConfiguration.BashShellAdministrators group. By default, this group is empty and you must add a user to the group manually.

Verify that the user you use to log in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

1. Use the vSphere Client to log in as `administrator@your_domain_name` to the vCenter Server instance.
The address is of the type `http://appliance-IP-address-or-FQDN/ui`.
2. From the vSphere Client menu, select **Administration**.
3. Select **Single Sign On > Users and Groups**.
4. Click on the **Groups** tab, select **SystemConfiguration.BashShellAdministrators** from the options available in the Group Name column.
5. Click **EDIT**.
6. In the **Edit Group** window, to add members, select the domain from the dropdown menu and then search for the required users.
7. Click **SAVE**.

Reboot a Node

In the vSphere Client, you can reboot a node in vCenter Server.

Verify that the user you use to log in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

1. Use the vSphere Client to log in as `administrator@your_domain_name` to the vCenter Server instance.
2. On the vSphere Client main page, click **Administration > Deployment > System Configuration**.
3. Under System Configuration, select a node from the list.
4. Click **REBOOT NODE**.

View the Health Status of Nodes

In the vSphere Client, you can view the health status of vCenter Server nodes.

Verify that the user you use to log in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

vCenter Server instances and machines that run vCenter Server services are considered nodes. Graphical badges represent the health status of nodes.

1. Use the vSphere Client to log in as `administrator@your_domain_name` to the vCenter Server instance.
The address is of the type `http://appliance-IP-address-or-FQDN/ui`.
2. From the vSphere Client menu, select **Administration**.
3. Select **Deployment > System Configuration**.
4. Select a node to view its health status.

Table 133: Health States

Badge Icon	Description
	Good. The health of the object is normal.
	Warning. The object is experiencing some problems.
	Critical. The object is either not functioning properly or will stop functioning soon.
	Unknown. No data is available for this object.

Export a Support Bundle

You can export a support bundle containing the log files for a specific product included in vCenter Server.

Verify that the user who logs in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in vCenter Single Sign-On.

1. Use the vSphere Client to log in as `administrator@your_domain_name` to the vCenter Server instance.
The address is of the type `http://appliance-IP-address-or-FQDN/ui`.

2. On the vSphere Client home page, click **Administration > Deployment > System Configuration**.
3. Select a node from the list and click **Export Support Bundle**.
4. In the **Export Support Bundle** window, expand the trees to view the services running in the appliance and deselect the services for which you do not want to export log files.
All the services are selected by default. If you want to export the support bundle and send it to VMware Support, leave all check boxes selected. The services are separated in two categories: a Cloud infrastructure category, which contains the services of specific products in the appliance, and a Virtual appliance category, which contains the services specific for the appliance and the vCenter Server product.
5. Click **Export Support Bundle** and save the bundle on your local machine.

You saved the support bundle to your machine and can explore it.

Using the Appliance Shell to Configure vCenter Server

You can access all of the vCenter Server API commands and plug-ins that you can use for monitoring, troubleshooting, and configuring the appliance by using the appliance shell.

You can run all commands in the appliance shell with or without the `pi` keyword.

Access the Appliance Shell

To access the plug-ins included in the appliance shell and to be able to see and use the API commands, first access the appliance shell.

1. Access the appliance shell.
 - If you have direct access to the appliance console, press Alt+F1.
 - If you want to connect remotely, use SSH or another remote console connection to start a session to the appliance.
2. Enter a user name and password recognized by the appliance.

You are logged in to the appliance shell and can see the welcome message.

Enable and Access the Bash Shell from the Appliance Shell

If you log in to the appliance shell as a user who has a super administrator role, you can enable access to the Bash shell of the appliance for other users. The root user has access to the appliance Bash shell by default.

The appliance Bash shell is enabled by default for the root

1. Access the appliance shell and log in as a user who has a super administrator role.
The default user with a super administrator role is root.
2. If you want to enable the Bash shell access for other users, run the following command.

```
shell.set --enabled true
```

3. To access the Bash shell run `shell` or `pi shell`.

Keyboard Shortcuts for Editing Commands

You can use various keyboard shortcuts to enter and edit commands in the appliance Bash shell.

Table 134: Keyboard Shortcuts and Function

Keyboard Shortcut	Details
Tab	Completes the current command. If you enter a part of the command name and press the Tab key, the system completes the command name. To view the commands that match a set of characters that you enter, type a character and press the Tab key.
Enter (at the command line)	Runs the command that you entered.
Enter (at the More prompt)	Displays the next page of output.
Delete or Backspace	Deletes the character that is on the left of the cursor.
Left arrow or Ctrl+B	Moves the cursor one character to the left. When you enter a command that extends beyond a single line, you can press the Left Arrow or Ctrl-B keys to go back to the beginning of the command.
Right arrow or Ctrl+F	Moves the cursor one character to the right.
Esc, B	Moves the cursor one word back.
Esc, F	Moves the cursor one word forward.
Ctrl+A	Moves the cursor to the beginning of the command line.
Ctrl+E	Moves the cursor to the end of the command line.
Ctrl+D	Deletes the character selected by the cursor.
Ctrl+W	Deletes the word next to the cursor.
Ctrl+K	Deletes the line forward. When you press Ctrl+K, everything that you entered starting from the cursor location to the end of the command line is deleted.
Ctrl+U or Ctrl+X	Deletes the line backward. When you press Ctrl+U, everything from the beginning of the command line to the cursor location is deleted.
Ctrl+T	Changes the places of the character to the left of the cursor with the character selected by the cursor.
Ctrl+R or Ctrl+L	Displays the system prompt and command line.
Ctrl+V or Esc, Q	Inserts a code to indicate to the system that the following keystroke must be treated as a command entry, not as an editing key.
Up arrow, or Ctrl+P	Recalls commands in the history buffer, beginning with the most recent command.
Down arrow or Ctrl+N	Returns to more recent commands in the history buffer after you use the Up arrow or Ctrl+P to recall commands.
Ctrl+Y	Recalls the most recent entry in the delete buffer. The delete buffer contains the last 10 items you have cut or deleted.
Esc, Y	Recalls the next entry in the delete buffer. The delete buffer contains the last 10 items you have cut or deleted. Press Ctrl+Y first to recall the most recent entry, and then press Esc, Y up to nine times to recall the remaining entries in the buffer.
Esc, C	Capitalizes the character selected by the cursor.
Esc, U	Changes the casing for all characters in the word selected by the cursor, up to the next space, to uppercase.
Esc, L	Changes the capitalized letters in a word from the character selected by the cursor to the end of the word to lowercase.

Get Help on APIs and Utilities

You can access the vCenter Server plug-ins and API commands from the appliance shell. You can use the plug-ins and commands for monitoring, troubleshooting, and configuring the appliance.

You can use the Tab key to autocomplete API commands, plug-in names, and API parameters. Plug-in parameters do not support autocomplete.

1. Access the appliance shell and log in.
2. To get help about the plug-ins, run the `help pi list` or the `? pi list` command.
You receive a list with all the plug-ins in the appliance.
3. To get help about the API commands, run the `help api list` or the `? api list` command.
You receive a list with all the API commands in the appliance.
4. To get help about a particular API command, run the `help api api_name` or the `? api api_name` command.
For example, to receive help about the `com.vmware.appliance.version1.timesync.set` command, run `help api timesync.set` or `? api timesync.set`.

Plug-Ins in vCenter Server Shell

The plug-ins in vCenter Server provide you with access to various administrative tools. The plug-ins reside in the CLI itself. The plug-ins are standalone Linux or VMware utilities, which do not depend on any VMware service.

Table 135: Plug-Ins Available in the vCenter Server

Plug-In	Description
<code>com.vmware.clear</code>	A plug-in that you can use to clear the terminal screen.
<code>com.vmware.cmsso-util</code>	A plug-in that you use for orchestrating changes to PNID, Machine Certificate, unregistering a node from Component Manager, vCenter Single Sign-On, and reconfiguring vCenter Server.
<code>com.vmware.dcli</code>	vAPI based CLI client.
<code>com.vmware.nslookup</code>	A plug-in that you can use to query the Domain Name System (DNS) to obtain domain name or IP address mapping or for any other specific DNS record.
<code>com.vmware.pgrep</code>	A plug-in that you can use to search for all named processes.
<code>com.vmware.pgtop</code>	A plug-in that you can use to monitor the PostgreSQL database.
<code>com.vmware.ping</code>	A plug-in that you can use to ping a remote host. Accepts the same arguments as <code>bin/ping</code> .
<code>com.vmware.ping6</code>	A plug-in that you can use to ping a remote host. Accepts the same arguments as <code>bin/ping6</code> .
<code>com.vmware.portaccess</code>	A plug-in that you can use to troubleshoot the port access of a host.
<code>com.vmware.ps</code>	A plug-in that you can use to see statistics on running processes.
<code>com.vmware.rvc</code>	Ruby vSphere Console
<code>com.vmware.service-control</code>	A plug-in that you can use to manage VMware services.
<code>com.vmware.shell</code>	A plug-in that allows access to the appliance Bash shell.
<code>com.vmware.showlog</code>	A plug-in that you can use to browse the log files.

Plug-In	Description
<code>com.vmware.shutdown</code>	A plug-in that you can use to restart or power off the appliance.
<code>com.vmware.software-packages</code>	A plug-in that you can use to update the software packages in the appliance.
<code>com.vmware.support-bundle</code>	A plug-in that you can use to create a bundle on the local file system and export it to a remote Linux system. If you use the plug-in with the <code>stream</code> command, the support bundle is not created on the local file system, but is directly exported to the remote Linux system.
<code>com.vmware.top</code>	A plug-in that displays process information. Accepts the same arguments as <code>/usr/bin/top/</code> .
<code>com.vmware.tracepath</code>	A plug-in that traces the path to a network host. Accepts the same arguments as <code>/sbin/tracepath</code> .
<code>com.vmware.tracepath6</code>	A plug-in that traces the path to a network host. Accepts the same arguments as <code>/sbin/tracepath6</code> .
<code>com.vmware.updatemgr-util</code>	A plug-in that you can use to configure VMware Update Manager.
<code>com.vmware.vcenter-restore</code>	A plug-in that you can use to restore vCenter Server.
<code>com.vmware.vimtop</code>	A plug-in that you can use to view a list of vSphere services and their resource use.

Browse Log Files Using showlog Plug-In

You can browse the log files in vCenter Server to examine them for errors.

1. Access the appliance shell and log in.
2. Type the `showlog` command, add a space, and press the Tab key to view all the contents of the `/var/log` folder.
3. Run the command for viewing the firstboot log files.

```
showlog /var/log/firstboot/cloudvm.log
```

API Commands in Appliance Shell

The API commands in vCenter Server let you perform various administrative tasks. The API commands are provided by appliance management service. You can edit time synchronization settings, monitor processes and services, set up the SNMP settings, and so on.

Table 136: API Commands Available in vCenter Server

API Command	Description
<code>com.vmware.appliance.health.applmgmt.get</code>	Get the health of the applmgmt service.
<code>com.vmware.appliance.health.databasesstorage.get</code>	Get the health of the database storage.
<code>com.vmware.appliance.health.load.get</code>	Get the CPU load health.
<code>com.vmware.appliance.health.mem.get</code>	Get the memory health.
<code>com.vmware.appliance.health.softwarepackages.get</code>	Get the health of the system update.
<code>com.vmware.appliance.health.storage.get</code>	Get the overall storage health.
<code>com.vmware.appliance.health.swap.get</code>	Get the swap health.

API Command	Description
<code>com.vmware.appliance.health.system.get</code>	Get the system health.
<code>com.vmware.appliance.health.system.lastcheck</code>	Get the time of the last check of the health status.
<code>com.vmware.appliance.monitoring.list</code>	Get a list of monitored items.
<code>com.vmware.appliance.monitoring.get</code>	Get monitored item information.
<code>com.vmware.appliance.monitoring.query</code>	Query a range of values for the monitored items.
<code>com.vmware.appliance.recovery.backup.job.cancel</code>	Cancel a backup job by id.
<code>com.vmware.appliance.recovery.backup.job.create</code>	Start a backup job.
<code>com.vmware.appliance.recovery.backup.job.get</code>	Get a status of the backup job by id.
<code>com.vmware.appliance.recovery.backup.job.list</code>	Get a list of backup jobs.
<code>com.vmware.appliance.recovery.backup.parts.list</code>	Get a list of the vCenter Server components that can be included in a backup job.
<code>com.vmware.appliance.recovery.backup.parts.get</code>	Get detailed info for a backup part.
<code>com.vmware.appliance.recovery.backup.validate</code>	Validate parameters for a backup job without starting the job.
<code>com.vmware.appliance.recovery.restore.job.cancel</code>	Cancel a restore job.
<code>com.vmware.appliance.recovery.restore.job.create</code>	Start a restore job.
<code>com.vmware.appliance.recovery.restore.job.get</code>	Get status of the restore job.
<code>com.vmware.appliance.recovery.restore.validate</code>	Validate restore parameters of a restore job without starting the job.
<code>com.vmware.appliance.system.uptime.get</code>	Gets the system uptime.
<code>com.vmware.appliance.version1.access.consolecli.get</code>	Get information about the state of the console-based controlled CLI (TTY1).
<code>com.vmware.appliance.version1.access.consolecli.set</code>	Set enabled state of console-based controlled CLI (TTY1).
<code>com.vmware.appliance.version1.access.dcu.get</code>	Get information about the state of the Direct Console User Interface (DCUI TTY2).
<code>com.vmware.appliance.version1.access.dcu.set</code>	Set enabled state of the Direct Console User Interface (DCUI TTY2).
<code>com.vmware.appliance.version1.access.shell.get</code>	Get information about the state of Bash shell, that is, access to Bash shell from within the controlled CLI.
<code>com.vmware.appliance.version1.access.shell.set</code>	Set enabled state of Bash shell, that is, access to Bash shell from within the controlled CLI.
<code>com.vmware.appliance.version1.access.ssh.get</code>	Get enabled state of the SSH-based controlled CLI.
<code>com.vmware.appliance.version1.access.ssh.set</code>	Set enabled state of the SSH-based controlled CLI.

API Command	Description
<code>com.vmware.appliance.version1.localaccounts.user.add</code>	Create a new local user account.
<code>com.vmware.appliance.version1.localaccounts.user.delete</code>	Delete a local user account.
<code>com.vmware.appliance.version1.localaccounts.user.get</code>	Get the local user account information.
<code>com.vmware.appliance.version1.localaccounts.user.list</code>	List local user accounts.
<code>com.vmware.appliance.version1.localaccounts.user.password.update</code>	Update the password of a logged in user or of the user that you specify in the <code>username</code> parameter.
<code>com.vmware.appliance.version1.localaccounts.user.set</code>	Update local user account properties, such as role, full name, enabled status, and password.
<code>com.vmware.appliance.version1.monitoring.snmp.disable</code>	Stop an enabled SNMP agent.
<code>com.vmware.appliance.version1.monitoring.snmp.enable</code>	Start a disabled SNMP agent.
<code>com.vmware.appliance.version1.monitoring.snmp.get</code>	Return an SNMP agent configuration.
<code>com.vmware.appliance.version1.monitoring.snmp.hash</code>	Generate localized keys for secure SNMPv3 communications.
<code>com.vmware.appliance.version1.monitoring.snmp.limits</code>	Get SNMP limits information.
<code>com.vmware.appliance.version1.monitoring.snmp.reset</code>	Restore settings to factory defaults.
<code>com.vmware.appliance.version1.monitoring.snmp.set</code>	Set SNMP configuration.
<code>com.vmware.appliance.version1.monitoring.snmp.stats</code>	Generate diagnostics report for SNMP agent.
<code>com.vmware.appliance.version1.networking.dns.domains.add</code>	Add domains to DNS search domains.
<code>com.vmware.appliance.version1.networking.dns.domains.list</code>	Get a list of DNS search domains.
<code>com.vmware.appliance.version1.networking.dns.domains.set</code>	Set DNS search domains.
<code>com.vmware.appliance.version1.networking.dns.hostname.get</code>	Get the Fully Qualified Domain Name.
<code>com.vmware.appliance.version1.networking.dns.hostname.set</code>	Set the Fully Qualified Domain Name.
<code>com.vmware.appliance.version1.networking.dns.servers.add</code>	Add a DNS server. This method fails if you use DHCP.
<code>com.vmware.appliance.version1.networking.dns.servers.get</code>	Get DNS server configuration.
<code>com.vmware.appliance.version1.networking.dns.servers.set</code>	Set the DNS server configuration. If the host is configured to acquire DNS servers and host name by using DHCP, a DHCP refresh is forced.
<code>com.vmware.appliance.version1.networking.firewall.addr.inbound.add</code>	Add a firewall rule to allow or deny access from an incoming IP address.
<code>com.vmware.appliance.version1.networking.firewall.addr.inbound.delete</code>	Delete a specific rule at a given position or delete all rules.
<code>com.vmware.appliance.version1.networking.firewall.addr.inbound.get</code>	Get an ordered list of inbound IP addresses that are allowed or denied by a firewall rule.
<code>com.vmware.appliance.version1.networking.interfaces.get</code>	Get information about a particular network interface.

API Command	Description
<code>com.vmware.appliance.version1.networking.interfaces.list</code>	Get a list of available network interfaces, including interfaces that are not yet configured.
<code>com.vmware.appliance.version1.networking.ipv4.get</code>	Get IPv4 network configuration for interfaces.
<code>com.vmware.appliance.version1.networking.ipv4.list</code>	Get IPv4 network configuration for all configured interfaces.
<code>com.vmware.appliance.version1.networking.ipv4.renew</code>	Renew IPv4 network configuration on interfaces. If the interface is configured to use DHCP for IP address assignment, the lease of the interface is renewed.
<code>com.vmware.appliance.version1.networking.ipv4.set</code>	Set IPv4 network configuration for an interface.
<code>com.vmware.appliance.version1.networking.ipv6.get</code>	Get IPv6 network configuration for interfaces.
<code>com.vmware.appliance.version1.networking.ipv6.list</code>	Get IPv6 network configuration for all configured interfaces.
<code>com.vmware.appliance.version1.networking.ipv6.set</code>	Set IPv6 network configuration for an interface.
<code>com.vmware.appliance.version1.networking.routes.add</code>	Add static routing rules. A destination/prefix of the type 0.0.0.0/0 (for IPv4) or ::/0 (for IPv6) refers to the default gateway.
<code>com.vmware.appliance.version1.networking.routes.delete</code>	Delete static routing rules.
<code>com.vmware.appliance.version1.networking.routes.list</code>	Get routing table. A destination/prefix of the type 0.0.0.0/0 (for IPv4) or ::/0 (for IPv6) refers to the default gateway.
<code>com.vmware.appliance.ntp.get</code>	Get NTP configuration settings. If you run the <code>timesync.get</code> command, you can retrieve the current time synchronization method (by using NTP or VMware Tools). The <code>ntp.get</code> command always returns the NTP server information, even when the time synchronization method is not set to NTP. If time synchronization method is not set by using NTP, the NTP status is displayed as down.
<code>com.vmware.appliance.version1.ntp.server.add</code>	Add NTP servers. This command adds NTP servers to the configuration. If the time synchronization is NTP-based, then NTP daemon is restarted to reload the new NTP servers. Otherwise, this command just adds servers to the NTP configuration.

API Command	Description
<code>com.vmware.appliance.version1.ntp.server.delete</code>	Delete NTP servers. This command deletes NTP servers from the configuration. If the time synchronization mode is NTP-based, the NTP daemon is restarted to reload the new NTP configuration. Otherwise, this command just deletes servers from the NTP configuration.
<code>com.vmware.appliance.version1.ntp.server.set</code>	Set NTP servers. This command deletes old NTP servers from the configuration and sets the input NTP servers in the configuration. If the time synchronization is set by using NTP, the NTP daemon is restarted to reload the new NTP configuration. Otherwise, this command just replaces the servers in NTP configuration with the NTP servers that you provide as input.
<code>com.vmware.appliance.version1.resources.cpu.stats.get</code>	Get CPU statistics.
<code>com.vmware.appliance.version1.resources.load.health.get</code>	Get load health.
<code>com.vmware.appliance.version1.resources.load.stats.get</code>	Get load averages (over 1, 5, and 15-minute intervals).
<code>com.vmware.appliance.version1.resources.mem.health.get</code>	Get memory health.
<code>com.vmware.appliance.version1.resources.mem.stats.get</code>	Get memory statistics.
<code>com.vmware.appliance.version1.resources.net.stats.get</code>	Get network statistics.
<code>com.vmware.appliance.version1.resources.net.stats.list</code>	Get network statistics for all interfaces that are up and running.
<code>com.vmware.appliance.version1.resources.processes.stats.list</code>	Get statistics on all processes.
<code>com.vmware.appliance.version1.resources.softwarepackages.health.get</code>	Get the health of the update component.
<code>com.vmware.appliance.version1.resources.storage.health.get</code>	Get storage health statistics.
<code>com.vmware.appliance.version1.resources.storage.stats.list</code>	Get storage statistics for each logical disk.
<code>com.vmware.appliance.version1.resources.swap.health.get</code>	Get swap health.
<code>com.vmware.appliance.version1.resources.swap.stats.get</code>	Get swap statistics.
<code>com.vmware.appliance.version1.resources.system.health.get</code>	Get the overall health of the system.
<code>com.vmware.appliance.version1.resources.system.stats.get</code>	Get the system status.
<code>com.vmware.appliance.version1.services.list</code>	Get the list of all known services.
<code>com.vmware.appliance.version1.services.restart</code>	Restart a service.
<code>com.vmware.appliance.version1.services.status.get</code>	Get the status of a service.
<code>com.vmware.appliance.version1.services.stop</code>	Stop a service.
<code>com.vmware.appliance.version1.system.storage.list</code>	Gets disk to partition mapping.
<code>com.vmware.appliance.version1.system.storage.resize</code>	Resizes all partitions to 100 percent of disk size.

API Command	Description
<code>com.vmware.appliance.version1.system.time.get</code>	Gets system time.
<code>com.vmware.appliance.version1.system.update.get</code>	Get the URL-based patching configuration.
<code>com.vmware.appliance.version1.system.update.set</code>	Set the URL-based patching configuration.
<code>com.vmware.appliance.version1.system.version.get</code>	Get the version of the appliance.
<code>com.vmware.appliance.version1.timesync.get</code>	Get the time synchronization configuration.
<code>com.vmware.appliance.version1.timesync.set</code>	Set the time synchronization configuration.

Configuring SNMP for vCenter Server

vCenter Server includes an SNMP agent that can send trap notifications and receive `GET`, `GETBULK`, and `GETNEXT` requests.

You can use the appliance shell API commands to enable and configure the vCenter Server SNMP agent. You configure the agent differently depending on whether you want to use SNMP v1/v2c or SNMP v3.

SNMP v3 informs are not supported. vCenter Server supports only notifications such as v1 and v2c traps, and v3 traps with all security levels.

Configure the SNMP Agent for Polling

If you configure the vCenter Server SNMP agent for polling, it can listen for and respond to requests from SNMP management client systems, such as `GET`, `GETNEXT`, and `GETBULK` requests.

By default, the embedded SNMP agent listens on UDP port 161 for polling requests from management systems. You can use the `snmp.set --port` command to configure an alternative port. To avoid conflicts between the port for the SNMP agent and the ports of other services, use a UDP port that is not defined in `/etc/services`.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.

The default user with super administrator role is `root`.

2. Run the `snmp.set --port` command to configure the port.

For example, run the following command:

```
snmp.set --port port
```

Here `port` is the port for the SNMP agent to use for listening for polling requests.

Important: The port you specify must not be already in use by other services. Use IP addresses from the dynamic range, port 49152 and up.

3. Optional: If the SNMP agent is not enabled, enable it by running the `snmp.enable` command.

Configure vCenter Server for SNMP v1 and v2c

When you configure the vCenter Server SNMP agent for SNMP v1 and v2c, the agent supports sending notifications and receiving `GET` requests.

In SNMP v1 and v2c, community strings are namespaces that contain one or more managed objects. Namespaces can act as a form for authentication, but this does not secure the communication. To secure the communication, use SNMP v3.

Configure SNMP Communities

To enable the vCenter Server SNMP agent to send and receive SNMP v1 and v2c messages, you must configure at least one community for the agent.

An SNMP community defines a group of devices and management systems. Only devices and management systems that are members of the same community can exchange SNMP messages. A device or management system can be a member of multiple communities.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.
2. Run the `snmp.set --communities` command to configure an SNMP community.
For example, to configure public, east, and west network operation center communities, run the following command:

```
snmp.set --communities public,eastnoc,westnoc
```


Each time you specify a community with this command, the settings you specify overwrite the previous configuration.
To specify multiple communities, separate the community names with a comma.

Configure SNMP Agent to Send v1 or v2c Notifications

You can use the vCenter ServerSNMP agent to send virtual machine and environmental notifications to management systems.

To send SNMP v1 and v2c notifications with the SNMP agent, you must configure the target, that is the receiver, unicast address, community, and an optional port. If you do not specify a port, the SNMP agent sends notifications to UDP port 162 on the target management system by default.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.
2. Run the `snmp.set --targets` command:

```
snmp.set --targets target_address@port/community
```


Here *target_address*, *port*, and *community* are the address of the target system, the port number to send the notifications to, and the community name, respectively. The port value is optional. If you do not specify a port, the default port, 161, is used.

Each time you specify a target with this command, the settings you specify overwrite all previously specified settings.
To specify multiple targets, separate them with a comma.

For example, run the following command for configuring the targets 192.0.2.1@678/targetcommunity and 2001:db8::1/anothercom:

```
snmp.set --targets 192.0.2.1@678/targetcommunity,2001:db8::1/anothercom
```
3. Optional: If the SNMP agent is not enabled, enable it by running the `snmp.enable` command.
4. Optional: To send a test trap to verify that the agent is configured correctly, run the `snmp.test` command.
The agent sends a `warmStart` trap to the configured target.

Configure vCenter Server for SNMP v3

When you configure the SNMP agent for SNMP v3, the agent supports sending traps. SNMP v3 also provides stronger security than v1 or v2c, including cryptographic authentication and encryption.

SNMP v3 informs are not supported. vCenter Server supports only notifications such as v1/v2c traps and v3 traps with all security levels.

Configure SNMP Engine ID

Every SNMP v3 agent has an engine ID, which serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate localized keys for authentication and encryption of SNMP v3 messages.

If you do not specify an engine ID before you enable the SNMP agent, when you enable the standalone SNMP agent, an engine ID is generated.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.

The default user with super administrator role is root.

2. Run the `snmp.set --engineid` command to configure the target.

For example, run the following command:

```
snmp.set --engineid 80001adc040102030405060708091011
```

where, the value can be broken into hexadecimal digits,

```
80:00:1a:dc:04:24:17:e2:02:b8:61:3f:54:00:00:00:00
```

and decoded as:

80: -- RFC 3411 encoding applies as high bit '10000000' 0x80 and lower 7 bits 0 are part of enterprise id.

00:1a:dc -- rest of enterprise id, the value parsed by vmware is, 1a:dc hexadecimal (6876 in decimal).

04: -- type of engineid as per RFC 3411, page 42. Using here '4' to say remaining octets can be translated to text.

0102030405060708091011 -- operator specified text encoded in hexadecimal characters.

Configure SNMP Authentication and Privacy Protocols

SNMP v3 optionally supports authentication and privacy protocols.

Authentication is used to ensure the identity of users. Privacy allows for encryption of SNMP v3 messages to ensure confidentiality of data. The privacy protocols provide a higher level of security than is available in SNMP v1 and v2c, which use community strings for security.

Both authentication and privacy are optional. However, you must enable authentication if you plan to enable privacy.

The SNMP v3 authentication and privacy protocols are licensed vSphere features and might not be available in some vSphere editions.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.

The default user with super administrator role is root.

2. Optional: Run the `snmp.set --authentication` command to configure authentication.

For example, run the following command:

```
snmp.set --authentication protocol
```

Here, *protocol* must be either *none*, for no authentication, SHA1, or MD5.

3. Optional: Run the `snmp.set --privacy` command to configure privacy protocol.

For example, run the following command:

```
snmp.set --privacy protocol
```

Here, *protocol* must be either *none*, for no privacy, or AES128.

Configure SNMP Users

You can configure up to five users who can access SNMP v3 information. User names must be no more than 32 characters long.

- Verify that you have configured the authentication and privacy protocols before configuring users.

- Verify that you know the authentication and privacy passwords for each user that you plan to configure. Passwords must be at least eight characters long. Store these passwords in files on the host system.

While configuring a user, you generate authentication and privacy hash values based on the user's authentication and privacy passwords and on the SNMP agent's engine ID. After configuring users, if you change the engine ID, the authentication protocol, or the privacy protocol, the users are no longer valid and must be reconfigured.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.
2. If you are using authentication or privacy, get the authentication and privacy hash values for the user by the running `snmp.hash --auth_hash --priv_hash` command.

For example, run the following command:

```
snmp.hash --auth_hash secret1 --priv_hash secret2
```

Here, *secret1* is the path to the file containing the user's authentication password and *secret2* is the path to the file containing the user's privacy password. Alternatively, you can specify the flag `--raw_secret` and set the boolean parameter to *true*.

The authentication and privacy hash values are displayed.

3. Configure the user by running `snmp.set --user`.

For example, run the following command:

```
snmp.set --user userid/authhash/privhash/security
```

The parameters in the command are as follows.

Parameter	Description
<i>userid</i>	Replace with the user name.
<i>authhash</i>	Replace with the authentication hash value.
<i>privhash</i>	Replace with the privacy hash value.
<i>security</i>	Replace with the level of security enabled for that user, which can be <i>auth</i> , for authentication only, <i>priv</i> , for authentication and privacy, or <i>none</i> , for no authentication or privacy.

Configure SNMP v3 Targets

Configure SNMP v3 targets to allow the SNMP agent to send SNMP v3 traps.

You can configure a maximum of three SNMP v3 targets, in addition to a maximum of three SNMP v1 or v2c targets.

To configure a target, you must specify a host name or IP address of the system that receives the traps, a user name, a security level, and whether to send traps. The security level can be either *none*, for no security, *auth*, for authentication only, or *priv*, for authentication and privacy.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.

2. Run the `snmp.set --v3targets` command to set up the SNMP v3 target.

For example, run the following command:

```
snmp.set --v3targets hostname@port/userid/secLevel/trap
```

The parameters in the command are as follows.

Parameter	Description
<i>hostname</i>	Replace with the host name or IP address of the management system that receives the traps.
<i>port</i>	Replace with the port on the management system that receives the traps. If you do not specify a port, the default port, 161, is used.
<i>userid</i>	Replace with the user name.
<i>secLevel</i>	Replace with either <code>none</code> , <code>auth</code> , or <code>priv</code> to indicate the level of authentication and privacy you have configured. Use <code>auth</code> if you have configured authentication only, <code>priv</code> if you have configured both authentication and privacy, and <code>none</code> if you have configured neither.

3. Optional: If the SNMP agent is not enabled, enable it by running the `snmp.enable` command.
4. Optional: To send a test trap to verify that the agent is configured correctly, run the `snmp.test` command.
The agent sends a `warmStart` trap to the configured target.

Configure SNMP Agent to Filter Notifications

You can configure the vCenter Server SNMP agent to filter out notifications if you do not want your SNMP management software to receive those notifications.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is `root`.
2. Run the `snmp.set --notraps` command to filter traps.
 - To filter specific traps, run the following command:

```
snmp.set --notraps oid_list
```

 Here, *oid_list* is a list of object IDs for the traps to filter, separated by commas. This list replaces any object IDs that were previously specified using this command.
 - To clear all trap filters, run the following command:

```
snmp.set --notraps reset
```
3. Optional: If the SNMP agent is not enabled, enable it by running the `snmp.enable` command.

The traps identified by the specified object IDs are filtered out of the output of the SNMP agent, and are not sent to SNMP management software.

Configure SNMP Management Client Software

After you have configured vCenter Server to send traps, you must configure your management client software to receive and interpret those traps.

Download the VMware MIB files from <https://kb.vmware.com/s/article/1013445>.

To configure your management client software, specify the communities for the managed device, configure the port settings, and load the VMware MIB files. See the documentation for your management system for specific instructions for these steps.

1. In your management software, specify the vCenter Server instance as an SNMP-based managed device.
2. If you are using SNMP v1 or v2c, set up appropriate community names in the management software.
These names must correspond to the communities set for the SNMP agent on vCenter Server.
3. If you are using SNMP v3, configure users and authentication and privacy protocols to match the protocols configured on vCenter Server.
4. If you configured the SNMP agent to send traps to a port on the management system other than the default UDP port 162, configure the management client software to listen on the port you configured.
5. Load the VMware MIBs into the management software to view the symbolic names for the vCenter Server variables.
To prevent lookup errors, load these MIB files in the following order before loading other MIB files:
 1. VMWARE-ROOT-MIB.mib
 2. VMWARE-TC-MIB.mib
 3. VMWARE-PRODUCTS-MIB.mib

The management software can now receive and interpret traps from vCenter Server.

Reset SNMP Settings to Factory Defaults

You can reset SNMP settings to factory defaults. You can also reset the value of a specific argument to the factory default.

You can reset a specific argument, such as the communities or targets. You can also reset the SNMP configuration to the factory defaults.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.
2. To reset specific arguments, run the command `snmp.set --arguments reset`.
For example, to reset the communities that you configured, run the following command:

```
snmp.set --communities reset
```
3. To reset the whole SNMP configuration to the factory defaults, run the command `snmp.reset`.

Configuring Time Synchronization Settings in vCenter Server

You can change the time synchronization settings in vCenter Server after deployment.

When you deploy vCenter Server, you can choose the time synchronization method to be either by using an NTP server or by using VMware Tools. In case the time settings in your vSphere network change, you can edit the vCenter Server and configure the time synchronization settings by using the commands in the appliance shell.

When you enable periodic time synchronization, VMware Tools sets the time of the guest operating system to be the same as the time of the host.

After time synchronization occurs, VMware Tools checks once every minute to determine whether the clocks on the guest operating system and the host still match. If not, the clock on the guest operating system is synchronized to match the clock on the host.

Native time synchronization software, such as Network Time Protocol (NTP), is typically more accurate than VMware Tools periodic time synchronization and is therefore preferred. You can use only one form of periodic time synchronization

in vCenter Server. If you decide to use native time synchronization software, vCenter Server VMware Tools periodic time synchronization is deactivated.

Related Links

[Synchronize ESXi Clocks with a Network Time Server on page 1347](#)

Before you install vCenter Server, make sure all machines on your vSphere network have their clocks synchronized.

Use VMware Tools Time Synchronization

You can set up vCenter Server to use VMware Tools time synchronization.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.
2. Run the command to enable VMware Tools time synchronization.

```
timesync.set --mode host
```
3. Optional: Run the command to verify that you successfully applied the VMware Tools time synchronization.

```
timesync.get
```


The command returns that the time synchronization is in host mode.

The time of the appliance is synchronized with the time of the ESXi host.

Add or Replace NTP Servers in the vCenter Server Configuration

To set up the vCenter Server to use NTP-based time synchronization, you must add the NTP servers to the vCenter Server configuration.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.
2. Add NTP servers to the vCenter Server configuration by running the following `ntp.set` command.

```
ntp.set --servers IP-addresses-or-host-names
```


In this command, *IP-addresses-or-host-names* is a comma-separated list of IP addresses or host names of the NTP servers.

This command removes the current NTP servers (if any) and adds the new NTP servers to the configuration. If the time synchronization is based on an NTP server, then the NTP daemon is restarted to reload the new NTP servers. Otherwise, this command replaces the current NTP servers in the NTP configuration with the new NTP servers you specify.
3. Optional: To verify that you successfully applied the new NTP configuration settings, run the following command.

```
ntp.get
```


The command returns a space-separated list of the servers configured for NTP synchronization. If the NTP synchronization is activated, the command returns that the NTP configuration is in Up status. If the NTP synchronization is deactivated, the command returns that the NTP configuration is in Down status.
4. Optional: To verify if the NTP server is reachable, run the following command.

```
ntp.test --servers IP-addresses-or-host-names
```


The command returns the status of the NTP servers.

If the NTP synchronization is deactivated, you can configure the time synchronization settings in the vCenter Server to be based on an NTP server. See [Synchronize the Time in vCenter Server with an NTP Server](#).

Synchronize the Time in vCenter Server with an NTP Server

You can configure the time synchronization settings in the vCenter Server to be based on an NTP server.

Set up one or more Network Time Protocol (NTP) servers in the vCenter Server configuration. See [Add or Replace NTP Servers in the vCenter Server Configuration](#).

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.

The default user with super administrator role is root.

2. Run the command to enable NTP-based time synchronization.

```
timesync.set --mode NTP
```

3. Optional: Run the command to verify that you successfully applied the NTP synchronization.

```
timesync.get
```

The command returns that the time synchronization is in NTP mode.

Managing Local User Accounts in vCenter Server

If you log in to the appliance shell as a super administrator, you can manage the local user accounts in vCenter Server by running commands in the appliance shell. The default user with a super administrator role is root.

User Roles in vCenter Server

There are three main user roles in vCenter Server.

The local users of vCenter Server have the rights to perform various tasks. Three user roles are available in vCenter Server:

Operator	Local users with the operator user role can read vCenter Server configuration.
Administrator	Local users with the administrator user role can configure vCenter Server.
Super Administrator	Local users with the super administrator user role can configure vCenter Server, manage the local accounts, and use the Bash shell.

Get a List of the Local User Accounts in vCenter Server

You can see the list of the local user accounts so that you can decide which user account to manage from the appliance shell.

1. Access the appliance shell and log in as a user who has a super administrator role.

The default user with a super administrator role is root.

2. Run the `localaccounts.user.list` command.

You can see a list of the local users. The information about a user includes the user name, status, role, status of the password, full name, and email.

Note: The list of local users includes only the local users who have their default shell as appliance shell.

Create a Local User Account in vCenter Server

You can create a new local user account.

For information about the user roles, see [User Roles in vCenter Server](#).

1. Access the appliance shell and log in as a user who has a super administrator role.

The default user with a super administrator role is root.

2. Run the `localaccounts.user.add --role --username --password` command.

For example, to add the local user account test with the operator user role, run the following command:

```
localaccounts.user.add --role operator --username test --password
```

The role can be operator, admin, or superAdmin.

You can also set up a new local user account and specify an email and the full name of the user. For example, to add the local user account test1 with the operator user role, full name TestName and the email address test1@mymail.com, run the following command:

```
localaccounts.user.add --role operator --username test1 --password --fullname TestName --email test1@mymail.com
```

You cannot use spaces in full names.

3. Enter and confirm the password of the new local user when prompted.

You created a new local user in the appliance.

Update Password of a Local User in vCenter Server

You can update the password of a local user in vCenter Server for security reasons.

1. Access the appliance shell and log in as a user who has a super administrator role.

The default user with a super administrator role is root.

2. Run the `localaccounts.user.password.update --username user name --password` command.

For example, to change the password of a user with user name test, run the following command:

```
localaccounts.user.password.update --username test --password
```

3. Enter and confirm the new password when prompted.

The password expiration period for local accounts is 90 days by default.

Update a Local User Account in vCenter Server

You can update an existing local user account in vCenter Server.

For information about the user roles, see [User Roles in vCenter Server](#).

1. Access the appliance shell and log in as a user who has a super administrator role.

The default user with a super administrator role is root.

2. Run the `localaccounts.user.set --username` command to update an existing local user.

- To update the role of the local user, run the following command:

```
localaccounts.user.set --username user name --role new role
```

Here, *user name* is the name of the user that you want to edit and *new role* is the new role. The role can be operator, admin, or superAdmin.

- To update the email of the local user, run the following command:

```
localaccounts.user.set --username user name --email new email address
```

Here, *user name* is the name of the user that you want to edit and *new email address* is the new email address.

- To update the full name of the local user, run the following command:

```
localaccounts.user.set --username user name --fullname new full name
```

Here, *user name* is the name of the user that you want to edit and *new full name* is the new full name of the user.

- To update the status of the local user, run the following command:

```
localaccounts.user.set --username user name --status new status
```

Here, *user name* is the name of the user that you want to edit and *status* is the new status of the local user. The status can be either disabled or enabled.

Delete a Local User Account in vCenter Server

You can delete a local user account in vCenter Server.

1. Access the appliance shell and log in as a user who has a super administrator role.

The default user with a super administrator role is root.

2. Run the `localaccounts.user.delete --username` command.

For example, to delete the user with user name test, run the following command:

```
localaccounts.user.delete --username test
```

The user is deleted.

Monitor Health Status and Statistics in vCenter Server

You can monitor the hardware health status of vCenter Server by using the API commands in the appliance shell. You can also monitor the health status of the update component for information about available patches.

You can view the status of the hardware components such as memory, CPU, storage, and network, as well as the update component that shows if the software packages are up to date according to the last check for available patches.

A particular health status can be green, yellow, orange, red, or gray. For more information, see [View vCenter Server Health Status](#).

For a complete list of the API commands that you can use for monitoring statistics and health of the vCenter Server system, see [API Commands in Appliance Shell](#).

1. Access the appliance shell and log in.

The user name that you use to log in can be of a user with an operator, administrator, or super administrator user role.

2. View the health status of a particular component.

- To view the health of the memory in vCenter Server, run the `mem.health.get` command.
- To view the health of the storage in vCenter Server, run the `storage.health.get` command.
- To view the health of the swap in vCenter Server, run the `swap.health.get` command.
- To view the health of the update component in vCenter Server, run the `softwarepackages.health.get` command.

Important: If you do not perform regular checks for available patches, the health status of the update component might become out-of-date. For information about checking for vCenter Server patches and enabling automatic checks for vCenter Server patches, see *vSphere Upgrade*.

- To view the overall health of the vCenter Server system, run the `health.system.get` command.

3. To view statistics about a particular hardware component, run the respective command.

For example, to view storage statistics for each logical disk, run the `storage.stats.list` command.

Using the vimtop Plug-In to Monitor the Resource Use of Services

You can use the `vimtop` utility plug-in to monitor vSphere services that run in the vCenter Server.

`vimtop` is a tool similar to `esxtop`, which runs in the environment of the vCenter Server. By using the text-based interface of `vimtop` in the appliance shell, you can view overall information about the vCenter Server, and a list of vSphere services and their resource use.

Monitor Services By Using vimtop in Interactive Mode

You can use the `vimtop` plug-in to monitor services in real time.

The default view of the `vimtop` interactive mode consists of the overview tables and the main table. You can use single-key commands in interactive mode to switch the view from processes to disks or network.

1. From an SSH client application, log in to the vCenter Server shell.
2. Run the `vimtop` command to access the plug-in in interactive mode.

Interactive Mode Command-Line Options

You can use various command-line options when you run the `vimtop` command to enter the plug-in interactive mode.

Table 137: Interactive Mode Command-Line Options

Option	Description
<code>-h</code>	Prints help for the <code>vimtop</code> command-line options.
<code>-v</code>	Prints the <code>vimtop</code> version number.
<code>-c filename</code>	Loads a user-defined <code>vimtop</code> configuration file. If the <code>-c</code> option is not used, the default configuration file is <code>/root/vimtop/vimtop.xml</code> . You can create your own configuration file, specifying a different filename and path by using the <code>w</code> single-key interactive command.
<code>-n number</code>	Sets the number of performed iterations before the <code>vimtop</code> exits interactive mode. <code>vimtop</code> updates the display <code>number</code> number of times and exits. The default value is 10000.
<code>-p / -dseconds</code>	Sets the update period in seconds.

Interactive Mode Single-Key Commands for vimtop

When running in interactive mode, `vimtop` recognizes several single-key commands.

All interactive mode panels recognize the commands listed in the following table.

Table 138: Interactive Mode Single-Key Commands

Key Names	Description
h	Show a help menu for the current panel, giving a brief summary of commands, and the status of secure mode.
i	Show or hide the top line view of the overview panel of the <code>vimtop</code> plug-in.
t	Show or hide the Tasks section, which displays information in the overview panel about the tasks currently running on the vCenter Server instance.
m	Show or hide the Memory section in the overview panel.
f	Show or hide the CPU section which displays information in the overview panel about all available CPUs.
g	Show or hide the CPUs section which displays information in the overview panel about the top 4 physical CPUs.
spacebar	Immediately refreshes the current pane.
p	Pause the displayed information about the services resource use in the current panels.
r	Refresh the displayed information about the services resource use in the current panels.
s	Set refresh period.
q	Exit the interactive mode of the <code>vimtop</code> plug-in.
k	Displays the Disks view of the main panel.
o	Switch the main panel to Network view.
Esc	Clear selection or return to the Processes view of the main panel.
Enter	Select a service to view additional details.
n	Show or hide names of the headers in the main panel.
u	Show or hide the measurement units in the headers in the main panel.
left, right arrows	Select columns.
up, down arrows	Select rows.
<,>	Move a selected column.
Delete	Remove selected column.
c	Add a column to the current view of the main panel. Use spacebar to add or remove columns from the displayed list.
a	Sort the selected column in ascending order.
d	Sort the selected column in descending order.
z	Clear the sort order for all columns.
l	Set width for the selected column.
x	Return the column widths to their default values.
+	Expand selected item.
-	Collapse selected item.
w	Write the current setup to a <code>vimtop</code> configuration file. The default file name is the one specified by <code>-c</code> option, or <code>/root/vimtop/vimtop.xml</code> if the <code>-c</code> option is not used. You can also specify a different file name on the prompt generated by the <code>w</code> command.

Using the Direct Console User Interface to Configure vCenter Server

After you deploy vCenter Server, you can reconfigure the network settings and enable access to the Bash shell for troubleshooting. To access the Direct Console User Interface, you must log in as root.

The home page of the Direct Console User Interface contains a link to the support bundle of the vCenter Server. The link to the support bundle is of the type `https://appliance-host-name:443/appliance/support-bundle`.

Log In to the Direct Console User Interface

The Direct Console User Interface lets you interact with vCenter Server locally by using text-based menus.

1. Log in to the vSphere Client or host UI client on which the vCenter Server VM is deployed.
2. Select the vCenter Server VM.
3. Launch the remote console or web console from the VM summary page.
4. Click inside the console window and press F2 to customize the system.
5. Type the password for the root user and press Enter.

Important: If you enter invalid credentials thrice, the root account is locked for five minutes.

You logged in to the Direct Console User Interface. You can change the password of the root user, edit the network settings, and enable access to the vCenter Server appliance Bash shell.

Change the Password of the Root User

To prevent unauthorized access to the vCenter Server Direct Console User Interface, you can change the password of the root user.

The default root password for the vCenter Server instance is the password you enter during deployment.

Important: The password for the root account of vCenter Server expires after 90 days. You can change the expiry time for an account by logging as root to the vCenter Server Bash shell, and running `chage -M number_of_days -W warning_until_expiration user_name`. To increase the expiration time of the root password to infinity, run the `chage -M -1 -E -1 root` command.

1. Log in to the Direct Console User Interface.
2. Select **Configure Password** and press Enter.
3. Type the old password of the root user, and press Enter.
4. Set up the new password and press Enter.
5. Press Esc until you return to the main menu of the Direct Console User Interface.

You changed the password of the root user of the appliance.

Configure the Management Network of vCenter Server

The vCenter Server instance can obtain networking settings from a DHCP server, or use static IP addresses. You can change the networking settings of vCenter Server from the Direct Console User Interface. You can change the IPv4, IPv6, and DNS configuration.

To change the IP address of the vCenter Server instance, verify that the system name is an FQDN. If, during deployment you set an IP address as a system name, you cannot change the IP address after the deployment. The system name is always used as a primary network identifier.

1. Log in to the Direct Console User Interface of vCenter Server
2. Select **Configure Management Network** and press Enter.
3. Change the IPv4 settings from **IP Configuration**.

Option	Description
Use dynamic IP address and network configuration	Obtains networking settings from a DHCP server if one is available on your network
Set static IP address and network configuration	Sets static networking configuration

4. Change the IPv6 settings from **IPv6 Configuration**.

Option	Description
Enable IPv6	Enables or disables IPv6
Use DHCP stateful configuration	Uses a DHCP server to obtain IPv6 addresses and networking settings
Use ICMP stateless configuration	Uses a Stateless Address Autoconfiguration (SLAAC) to obtain IPv6 addresses and network settings

5. Change the DNS settings from **DNS Configuration**.

Option	Description
Obtain DNS server address and hostname automatically	Obtains the DNS server address and host name automatically. Use this option if the IP settings are obtained automatically from a DHCP server.
Use the following DNS server address and hostname	Sets the static IP address and host name for the DNS server.

6. Set custom DNS suffixes from **Custom DNS Suffixes**.
If you do not specify any suffixes, a default suffix list is derived from the local domain name.
7. Press Esc until you return to the main menu of the Direct Console User Interface.

Restart the Management Network of vCenter Server

Restart the management network of vCenter Server to restore the network connection.

1. Log in to the Direct Console User Interface of vCenter Server.
2. Select **Restart Management Network** and press Enter.
3. Press F11.

Enable Access to the Bash Shell

You can use the Direct Console User Interface to enable local and remote access to the Bash shell. Bash shell access enabled through Direct Console User Interface remains enabled for 3600 seconds.

1. Log in to the Direct Console User Interface of vCenter Server.
2. Select **Troubleshooting Options** and press Enter.
3. From the Troubleshooting Mode Options menu, select to enable either Bash shell or SSH.
4. Press Enter to enable the service.
5. Press Esc until you return to the main menu of the Direct Console User Interface.

Access the vCenter Server Bash shell for troubleshooting.

Access the Bash Shell for Troubleshooting

Log in to the Bash shell for troubleshooting purposes only.

1. Access the shell using one of the following methods.
 - If you have direct access to the vCenter Server instance, press Alt+F1.
 - If you want to connect remotely, use SSH or another remote console connection to start a session.
2. Enter a user name and password.
3. In the shell, enter the command `pi shell` or `shell` to access the Bash shell.

Export a vCenter Server Support Bundle for Troubleshooting

You can export the support bundle of the vCenter Server instance in the appliance for troubleshooting using the URL displayed on the DCUI home screen.

You can also collect the support bundle from the vCenter Server appliance Bash shell by running the `vc-support.sh` script.

The support bundle is exported in `.tgz` format.

1. Log in to the Windows host machine on which you want to download the bundle.
2. Open a Web browser and enter the URL to the support bundle displayed in the DCUI.
`https://appliance-fully-qualified-domain-name:443/appliance/support-bundle`
3. Enter the user name and password of the root user.
4. Click **Enter**.

The support bundle is downloaded as `.tgz` file on your Windows machine.

vCenter Server and Host Management

Learn how to use the VMware®vSphere Client components, configure and manage hosts, migrate virtual machines, and manage licenses in your vCenter Server environment.

vCenter Server and Host Management also provides brief introductions to the various tasks you can perform within the system, and it cross-references to the documentation that describes the tasks in detail.

vCenter and Host Management covers ESXi and vCenter Server.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

vCenter Server and Host Management is intended for system administrators who are experienced Windows or Linux system administrators and who are familiar with virtual machine technology and data center operations.

What Are the Main vSphere Concepts and Features

VMware vSphere® uses the power of virtualization to transform data centers into simplified cloud computing infrastructures, enabling IT organizations to deliver flexible and reliable IT services.

The two core components of vSphere are VMware ESXi™ and VMware vCenter Server®. ESXi is the hypervisor on which you create and run virtual machines. vCenter Server is a service that acts as a central administrator for ESXi hosts that are connected on a network. With vCenter Server, you can pool and manage the resources of multiple hosts. vCenter Server allows you to monitor and manage your physical and virtual infrastructure.

Some vSphere components are available as plug-ins that extend the functionality of the vSphere product.

What Are the Basics of Virtualization

A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. ESXi is the hypervisor in a vSphere environment. The hypervisor serves as a platform for running virtual machines and allows for the consolidation of computing resources.

Each virtual machine contains its own virtual, or software-based, hardware, including a virtual CPU, memory, hard disk, and network interface card.

ESXi is installed on physical or virtual hardware in a virtualized data center, and acts as a platform for virtual machines. The hypervisor provides physical hardware resources dynamically to virtual machines to support the operation of the virtual machines. The hypervisor allows virtual machines to operate with a degree of independence from the underlying physical hardware. For example, a virtual machine can be moved from one physical host to another, or its virtual disks can be moved from one type of storage to another, without affecting the functioning of the virtual machine.

Because virtual machines are decoupled from the underlying physical hardware, virtualization allows you to consolidate physical computing resources such as CPUs, memory, storage, and networking into pools of resources. These resources can be dynamically and flexibly made available to virtual machines. With the vCenter Server management platform, you can increase the availability and security of your virtual infrastructure.

What Is the Physical Topology of a vSphere Data Center

A typical VMware vSphere data center consists of physical building blocks such as x86 virtualization servers, storage networks and arrays, IP networks, a management server, and desktop clients.

The vSphere data center includes the following components.

ESXi Hosts

Industry standard x86 servers that run ESXi on the bare metal. ESXi software provides resources for and runs the virtual machines. You can group similarly configured x86 servers with connections to the same network and storage subsystems. This grouping creates an aggregate set of resources in the virtual environment, called a cluster.

Storage networks and arrays

VMware vSphere uses Fibre Channel SAN arrays, iSCSI SAN arrays, and NAS arrays to meet different data center storage needs. With storage area networks, you can connect and share storage arrays between groups of servers. This arrangement allows aggregation of the storage resources and provides more flexibility in provisioning them to virtual machines. For detailed information, see *vSphere Storage*.

IP networks

Each compute server can have multiple physical network adapters to provide high bandwidth and reliable networking to the entire VMware vSphere data center. For detailed information, see *vSphere Networking*.

vCenter Server

vCenter Server provides a single point of control to the data center. It provides essential data center services such as access control, performance monitoring, and configuration. It unifies the resources from the individual computing servers to be shared among virtual machines in the entire data center. It manages the assignment of virtual machines to the ESXi hosts and the assignment of resources to the virtual machines within a given computing server. These assignments are based on the policies that the system administrator sets.

Compute servers continue to function even if vCenter Server becomes unreachable (for example, if the network is severed). The ESXi hosts can be managed separately and continue to run the virtual machines assigned to them based on the resource assignment that was last set. After connection to vCenter Server is restored, it can manage the data center as a whole again.

Management clients

VMware vSphere provides several interfaces for data center management and virtual machine access. These interfaces include vSphere Client for access through a web browser or vSphere Command-Line Interface (vSphere CLI).

What Are the Software Components of vSphere

VMware vSphere is a suite of software components for virtualization. These include ESXi, vCenter Server, and other software components such as vCenter Single Sign-On, and vCenter Server plug-ins that fulfill several different functions in the vSphere environment.

vSphere includes the following software components:

ESXi

The hypervisor runs virtual machines. Each virtual machine has a set of configuration and disk files that together perform all the functions of a physical machine.

Through ESXi, you run the virtual machines, install operating systems, run applications, and configure the virtual machines. Configuration includes identifying the virtual machine's resources, such as storage devices.

The server provides bootstrapping, management, and other services that manage your virtual machines.

vCenter Server

A service that acts as a central administrator for VMware ESXi hosts that are connected on a network. vCenter Server directs actions on the virtual machines and the ESXi hosts.

vCenter Server is installed to run automatically on a preconfigured virtual machine. The vCenter Server service runs continuously in the background. It performs its monitoring and managing activities even when no vSphere Clients are connected and when no one is logged on to the computer where it resides. It must have network access to all the hosts it manages.

vCenter Server is deployed as a preconfigured virtual machine optimized for running vCenter Server and the vCenter Server components. You can deploy vCenter Server on ESXi hosts 6.5 or later.

All prerequisite services for running vCenter Server and the vCenter Server components are bundled in the vCenter Server installation. All vCenter Server services run as child processes of the VMware Service Library Lifecycle Manager service. See the *vCenter Server Installation and Setup* documentation for details about setting up this configuration.

vCenter Single Sign-On

A service that is part of the vCenter Server management infrastructure. The vCenter Single Sign-On authentication service makes the VMware Cloud infrastructure platform more secure by allowing the various vSphere software components to communicate with each other. The vCenter Single Sign-On authentication service uses a secure token exchange mechanism instead of requiring each component to authenticate a user separately with a directory service like Active Directory.

When you install vCenter Single Sign-On, the following components are deployed.

- STS (Security Token Service)** STS certificates enable a user who has logged on through vCenter Single Sign-On to authenticate to any vCenter service that vCenter Single Sign-On supports. The STS service issues Security Assertion Markup Language (SAML) tokens. These security tokens represent the identity of a user in each of the vCenter Single Sign-On identity sources.
- Administration server** The administration server allows users with vCenter Single Sign-On administrator privileges to configure the vCenter Single Sign-On service and manage users and groups from the vSphere Client. Initially, only the user `administrator@vsphere.local` has these privileges.
- vCenter Lookup Service** vCenter Lookup Service contains topology information about the vSphere infrastructure, enabling vSphere

components to connect to each other securely. Unless you are using Simple Install, you are prompted for the Lookup Service URL when you install other vSphere components. For example, the Inventory Service and the vCenter Server installers ask for the Lookup Service URL and then contact the Lookup Service to find vCenter Single Sign-On. After installation, the Inventory Service and vCenter Server system are registered with the vCenter Lookup Service so other vSphere components, like the vSphere Client, can find them.

VMware Directory Service

Directory service associated with the vsphere.local domain. This service is a multi-tenanted, peer-replicating directory service that makes an LDAP directory available on port 389. In multisite mode, an update of VMware Directory Service content in one VMware Directory Service instance results in the automatic update of the VMware Directory Service instances associated with all other vCenter Single Sign-On nodes.

vCenter Server plug-ins

Applications that provide additional features and functionality to vCenter Server. Typically, plug-ins consist of a server component and a client component. After the plug-in server is installed, it is registered with vCenter Server and the plug-in client is available to the for download. After a plug-in is installed on the vSphere Client, it might alter the interface by adding views, tabs, toolbar buttons, or menu items related to the added functionality.

Plug-ins leverage core vCenter Server capabilities, such as authentication and permission management, but can have their own types of events, tasks, metadata, and privileges.

Some vCenter Server features are implemented as plug-ins, and can be managed using the vSphere Client Plug-in Manager. These features include vCenter Storage Monitoring, vCenter Hardware Status, and vCenter Service Status.

vCenter Server database

Persistent storage for maintaining the status of each virtual machine, host, and user managed in the vCenter Server environment. The vCenter Server database can be remote or local to the vCenter Server system.

The database is installed and configured during vCenter Server installation.

tcServer

If you are accessing your ESXi host directly through the VMware Host Client, and not through a vCenter Server system and associated , you do not use a vCenter Server database.

Many vCenter Server functions are implemented as web services that require the tcServer. The tcServer is installed on the vCenter Server machine as part of the vCenter Server installation.

Features that require the tcServer to be running include: ICIM/ Hardware Status tab, Performance charts, WebAccess, Storage Policy-Based services, and vCenter Service status.

vCenter Server agent

On each managed host, the software that collects, communicates, and runs the actions received from vCenter Server. The vCenter Server agent is installed the first time any host is added to the vCenter Server inventory.

Host agent

On each managed host, the software that collects, communicates, and runs the actions received through the vSphere Client. It is installed as part of the ESXi installation.

vSphere Cluster Services

vSphere Cluster Services (vCLS) is activated by default and runs in all vSphere clusters. vCLS ensures that if vCenter Server becomes unavailable, cluster services remain available to maintain the resources and health of the workloads that run in the clusters. vCenter Server is still required to run DRS and HA.

vCLS is activated when you upgrade to vSphere 7.0 Update 3 or when you have a new vSphere 7.0 Update 3 or later deployment. vCLS is upgraded as part of vCenter Server upgrade.

vCLS uses agent virtual machines to maintain cluster services health. The vCLS agent virtual machines (vCLS VMs) are created when you add hosts to clusters. Up to three vCLS VMs are required to run in each vSphere cluster, distributed within a cluster. vCLS is also activated on clusters which contain only one or two hosts. In these clusters the number of vCLS VMs is one and two, respectively.

New anti-affinity rules are applied automatically. Every three minutes a check is performed, if multiple vCLS VMs are located on a single host they will be automatically redistributed to different hosts.

Table 139: Number of vCLS Agent VMs in Clusters

Number of Hosts in a Cluster	Number of vCLS Agent VMs
1	1
2	2
3 or more	3

vCLS VMs run in every cluster even if cluster services like vSphere DRS or vSphere HA are not activated on the cluster. The life cycle operations of vCLS VMs are managed by vCenter Server services like ESX Agent Manager and Workload Control Plane. vCLS VMs do not support NICs.

A cluster activated with vCLS can contain ESXi hosts of different versions if the ESXi versions are compatible with vCenter Server. vCLS works with vSphere Lifecycle Manager clusters.

Monitoring vSphere Cluster Services

You can monitor the resources consumed by vCLS VMs and their health status.

vCLS VMs are not displayed in the inventory tree in the **Hosts and Clusters** tab. vCLS VMs from all clusters within a data center are placed inside a separate VMs and templates folder named **vCLS**. This folder and the vCLS VMs are visible

only in the **VMs and Templates** tab of the vSphere Client. These VMs are identified by a different icon than regular workload VMs. You can view information about the purpose of the vCLS VMs in the **Summary** tab of the vCLS VMs.

You can monitor the resources consumed by vCLS VMs in the **Monitor** tab.

Table 140: vCLS VM Resource Allocation

Property	Size
VMDK size	245 MB (thin disk)
Memory	128 MB
CPU	1 vCPU
Hard disk	2 GB
Storage on datastore	480 MB (thin disk)

Note: Each vCLS VM has 100MHz and 100MB capacity reserved in the cluster. Depending on the number of vCLS VMs running in the cluster, a max of 400 MHz and 400 MB of capacity can be reserved for these VMs.

You can monitor the health status of vCLS in the **Cluster Services** portlet displayed in the **Summary** tab of the cluster.

Table 141: Health status of vCLS

Status	Color Coding	Summary
Healthy	Green	If there is at least one vCLS VM running, the status remains healthy, regardless of the number of hosts in the cluster.
Degraded	Yellow	If there is no vCLS VM running for less than 3 minutes (180 seconds), the status is degraded.
Unhealthy	Red	If there is no vCLS VM running for 3 minutes or more, the status is unhealthy in a DRS enabled cluster.

Maintaining Health of vSphere Cluster Services

vCLS VMs are always powered-on because vSphere DRS depends on the availability of these VMs. These VMs should be treated as system VMs. Only administrators can perform selective operations on vCLS VMs. To avoid failure of cluster services, avoid performing any configuration or operations on the vCLS VMs.

vCLS VMs are protected from accidental deletion. Cluster VMs and folders are protected from modification by users, including administrators.

Only users which are part of the Administrators SSO group can perform the following operations:

- ReadOnly access for vCLS VMs
- Console access to vCLS VMs
- Relocate vCLS VMs to either new storage, compute resource or both using cold or hot migration
- Use tags and custom attributes for vCLS VMs

Operations that might disrupt the healthy functioning of vCLS VMs:

- Changing the power state of the vCLS VMs

- Resource reconfiguration of the vCLS VMs such as changing CPU, Memory, Disk size, Disk placement
- VM encryption
- Triggering vMotion of the vCLS VMs
- Changing the BIOS
- Removing the vCLS VMs from the inventory
- Deleting the vCLS VMs from disk
- Enabling FT of vCLS VMs
- Cloning vCLS VMs
- Configuring PMem
- Moving vCLS VM to a different folder
- Renaming the vCLS VMs
- Renaming the vCLS folders
- Enabling DRS rules and overrides on vCLS VMs
- Enabling HA admission control policy on vCLS VMs
- Enabling HA overrides on vCLS VMs
- Moving vCLS VMs to a resource pool
- Recovering vCLS VMs from a snapshot

When you perform any disruptive operation on the vCLS VMs, a warning dialog box appears.

Troubleshooting:

The health of vCLS VMs, including power state, is managed by VMware ESX Agent Manager and Workload Control Plane services. In case of power on failure of vCLS VMs, or if the first instance of DRS for a cluster is skipped due to lack of quorum of vCLS VMs, a banner appears in the cluster summary page along with a link to a Knowledge Base article to help troubleshoot the error state.

Because vCLS VMs are treated as system VMs, you do not need to backup or snapshot these VMs. The health state of these VMs is managed by vCenter Server services.

What Are the Client Interfaces for vSphere

You can access vSphere components through the vSphere Client, the VMware Host Client, and the vSphere Command-Line Interfaces.

vSphere Client

The vSphere Client, introduced in vSphere 6.5, is an HTML5-based client and is included with vCenter Server. As of vSphere 7.0, the vSphere Web Client has been deprecated. The vSphere Client is the primary interface for connecting to and managing vCenter Server instances.

VMware Host Client

The VMware Host Client is a Web-based application that you can use to manage individual ESXi hosts that are not connected to a vCenter Server system.

For more information about the VMware Host Client, see *vSphere Single Host Management - VMware Host Client*.

vSphere Command-Line Interfaces

vSphere supports multiple command-line interfaces for configuring virtual machines, ESXi hosts, and vCenter Server.

What Are the Managed Inventory Objects in vSphere

In vSphere, the inventory is a collection of virtual and physical objects, such as data centers, networks, and datastores, on which you can place permissions, monitor tasks and events, and set alarms. You can rename inventory objects, and you can group most inventory objects by using folders to more easily manage them.

All inventory objects, with the exception of hosts, can be renamed to represent their purposes. For example, they can be named after company departments or locations or functions.

Note: Managed object names cannot exceed 214 bytes (UTF-8 encoded).

vCenter Server monitors and manages the following inventory objects:

Data Centers

Unlike folders, which are used to organize specific object types, a data center is an aggregation of all the different types of objects used to work in virtual infrastructure.

Within each data center, there are four separate hierarchies.

- Virtual machines (and templates)
- Hosts (and clusters)
- Networks
- Datastores

Clusters

A collection of ESXi hosts and associated virtual machines intended to work together as a unit. When you add a host to a cluster, the host's resources become part of the cluster's resources. vCenter Server manages the resources of all hosts in a cluster as one unit.

Datastores

A virtual representation of physical storage resources in the data center. A datastore is the storage location for virtual machine files. In an on-premises SDDC, these physical storage resources can come from the local SCSI disk of the ESXi host, the Fibre Channel SAN disk arrays, the iSCSI SAN disk arrays, or Network Attached Storage (NAS) arrays. For both on-premises and cloud SDDCs, vSAN datastores hide the idiosyncrasies of the underlying physical storage and present a uniform model for the storage resources required by virtual machines.

Folders

Folders allow you to group objects of the same type so you can easily manage them. For example, you can use folders to set permissions across objects, to set alarms across objects, and to organize objects in a meaningful way.

A folder can contain other folders, or a group of objects of the same type: data centers, clusters, datastores, networks, virtual machines, templates, or hosts. For example, one folder can contain hosts and a folder containing hosts, but it cannot contain hosts and a folder containing virtual machines.

Hosts

The physical computer on which ESXi is installed. All virtual machines run on hosts or clusters.

Networks

A set of virtual network interface cards (virtual NICs), distributed switches or vSphere Distributed Switches, and port groups or distributed port groups that connect virtual machines to each other or to the physical network outside of the virtual data center. You can monitor networks and set permissions and alarms on port groups and distributed port groups.

Resource pools

Resource pools are used to compartmentalize the CPU and memory resources of a host or cluster. Virtual machines run in, and draw their resources from, resource pools. You can create multiple resource pools as direct children of a standalone host or

Templates

cluster and then delegate control over each resource pool to other individuals or organizations.

You can monitor resources and set alarms on them.

A template is a primary copy of a virtual machine that can be used to create and provision new virtual machines. Templates can have a guest operating system and application software installed. They can be customized during deployment to ensure that the new virtual machine has a unique name and network settings.

Virtual machines

A virtualized computer environment in which a guest operating system and associated application software can run. Multiple virtual machines can operate on the same managed host machine concurrently.

vApps

vSphere vApp is a format for packaging and managing applications. A vApp can contain multiple virtual machines.

Optional vCenter Server Components

Optional vCenter Server components are packaged and installed with the base product, but might require a separate license.

Optional vCenter Server features include:

VMware vSphere® vMotion®

Enables you to move running virtual machines from one ESXi host to another ESXi host without service interruption. It requires licensing on both the source and target host. vCenter Server centrally coordinates all vSphere vMotion activities.

vSphere® Storage vMotion®

Allows you to move the disks and configuration file of a running virtual machine from one datastore to another without service interruption. It requires licensing on the virtual machine's host.

VMware vSphere® High Availability

Enables a cluster with High Availability. If a host fails, all virtual machines that were running on the host are promptly restarted on different hosts in the same cluster.

When you enable the cluster for vSphere HA, you specify the number of hosts you want to be able to recover. If you specify the number of host failures allowed as **1**, vSphere HA maintains enough capacity across the cluster to tolerate the failure of one host. All running virtual machines on that host can be restarted on remaining hosts. By default, you cannot turn on a virtual machine if doing so violates required failover capacity.

VMware vSphere® Distributed Resource Scheduler™

Helps improve resource allocation and power consumption across all hosts and resource pools. vSphere DRS collects resource use information for all hosts and virtual machines in the cluster and gives recommendations (or migrates virtual machines) in one of two situations:

- Initial placement – When you power on a virtual machine in the cluster for the first time, DRS either places the virtual machine or makes a recommendation.
- Load balancing – DRS attempts to improve resource use across the cluster by performing automatic migrations of virtual machines (vSphere vMotion) or by providing a recommendation for virtual machine migrations.

vSphere DRS includes VMware vSphere® Distributed Power Management™ (DPM) capabilities. When DPM is enabled, the system compares cluster-level and host-level capacity to the demands of virtual machines that are running in the cluster. Based on the results of the comparison, DPM recommends (or

VMware vSphere® Storage DRS™

implements) actions that can reduce the power consumption of the cluster.

Allows you to manage multiple datastores as a single resource, called a datastore cluster. A datastore cluster is an aggregation of multiple datastores into a single logical, load-balanced pool. You can treat the datastore cluster as a single flexible storage resource for resource management purposes. You can assign a virtual disk to a datastore cluster, and vSphere Storage DRS finds an appropriate datastore for it. The load balancer takes care of initial placement and future migrations based on workload measurements. Storage space balancing and I/O balancing minimize the risk of running out of space and the risk of I/O bottlenecks slowing the performance of virtual machines.

VMware vSphere® Fault Tolerance

vSphere Fault Tolerance provides continuous availability for virtual machines by creating and maintaining a Secondary VM that is identical to the Primary VM. This Secondary VM is continuously available to replace the Primary VM in a failover situation.

What is vCenter Enhanced Linked Mode

vCenter Enhanced Linked Mode links multiple vCenter Server systems. With Enhanced Linked Mode, you can view and search across all linked vCenter Server systems. This mode replicates roles, permissions, licenses, and other key data across systems.

Enhanced Linked Mode provides the following features for vCenter Server:

- You can log in to all linked vCenter Server systems simultaneously with a single user name and password.
- You can view and search the inventories of all linked vCenter Server systems within the vSphere Client.
- Roles, permission, licenses, tags, and policies are replicated across linked vCenter Server systems.

To join vCenter Server systems in Enhanced Linked Mode, connect them to the same vCenter Single Sign-On domain.

For more information about vCenter Enhanced Linked Mode, see the *vCenter Server Installation and Setup* guide.

vCenter Server Plug-Ins

vCenter Server plug-ins extend the capabilities of vCenter Server by providing additional features and functions.

Some plug-ins are installed as part of the base vCenter Server product.

vCenter Server Storage Monitoring

Allows you to review information on storage use and to map relationships visually between all storage entities available in vCenter Server.

vCenter Server Hardware Status

Uses CIM monitoring to display the hardware status of hosts that vCenter Server manages.

vCenter Server Service Status

Displays the status of vCenter Server services.

Some plug-ins are packaged separately from the base product and require separate installation. You can update plug-ins and the base product independently of each other. VMware modules include:

vSphere Lifecycle Manager

Enables administrators to apply updates and patches across ESXi hosts and all managed virtual machines. Administrators can create user-defined security baselines that represent a set of security standards. Security administrators can compare hosts and virtual machines against these baselines to identify and remediate systems that are not in compliance.

VMware Aria Automation Orchestrator

A workflow engine that enables you to create and run automated workflows in your vSphere environment. VMware Aria Automation

Orchestrator coordinates workflow tasks across multiple VMware products and third-party management and administration solutions through its open plug-in architecture. VMware Aria Automation Orchestrator provides a library of workflows that are extensible. You can use any operation available in the vCenter Server API to customize VMware Aria Automation Orchestrator workflows.

Using the vSphere Client

Use the vSphere Client to connect to vCenter Server systems and manage vSphere inventory objects.

Use of the vSphere Client requires a supported Web browser.

The home screen of the vSphere Client is a system dashboard that aggregates data from different sources in the environment together in a single, unified view.

VMware has tested and supports the following guest operating systems and browser versions for the vSphere Client.

Table 142: Supported Guest Operating Systems and Browser Versions for the vSphere Client.

Operating system	Browser
Windows 32-bit and 64-bit	Microsoft Edge: 90 or later. Mozilla Firefox: 80 or later. Google Chrome: 89 or later.
Mac OS	Microsoft Edge: 90 or later. Mozilla Firefox: 80 or later. Google Chrome: 89 or later.

Later versions of these browsers are likely to work, but have not been tested.

How to Use and Configure Your vSphere Client

Learn how to log in to a vCenter Server instance, configure the vSphere Client timeout settings, and edit the login banner.

How to Log in to Your vCenter Server by Using the vSphere Client

Learn how to use the vSphere Client to log in to vCenter Server and manage your vSphere inventory.

The vSphere Client is automatically installed as part of the vCenter Server appliance deployment. This way, the vSphere Client always points to the same vCenter Single Sign-On instance.

1. Open a Web browser and enter the URL for your vCenter Server instance:

`https://vcenter_server_ip_address_or_fqdn`

As an alternative, you can open a Web browser and enter the URL for the vSphere Client:

`https://vcenter_server_ip_address_or_fqdn/ui.`

2. If a warning message about a potential security risk appears, select to continue to the website.

Browser	Action
Microsoft Edge	1. Click Details . 2. Under the additional message that appears, click Go on to the webpage .
Mozilla Firefox	1. Click Advanced .

Browser	Action
	2. Under the additional message that appears, click Accept the Risk and Continue .
Google Chrome	1. Click Advanced . 2. Under the additional message that appears, click Proceed to vcenter_server_ip_address_or_fqdn .

3. On the vSphere Welcome page, select **Launch vSphere Client (HTML5)**.
4. If the warning message about a potential security risk appears again, repeat Step 2.
5. Enter the credentials of a user who has permissions on vCenter Server and click **Login**.

The vSphere Client connects to all the vCenter Server systems on which the specified user has permissions, and you can view and manage the vSphere inventory.

Manage the Login Message to the vSphere Client Login Page

You can add a login message to the vSphere Client login page. You can also configure a custom login message and provide a check box for user consent.

1. Log in with the vSphere Client to the vCenter Server.
2. Specify the user name and password for administrator@vsphere.local or another member of the vCenter Single Sign-On Administrators group.
If you specified a different domain during installation, log in as administrator@mydomain.
3. Navigate to the Configuration UI.
 - a) From the **Home** menu, select **Administration**.
 - b) Under **Single Sign On**, click **Configuration**.
4. Click the **Login Message** tab.
5. Click **Edit** and configure the login message.

Option	Description
Show login message	Toggle on Show login message to enable the login message. You cannot make changes to the login message unless you toggle on this switch.
Login message	Title of the message. By default, when Consent checkbox is toggled on, the login message text is <code>I agree to Terms and Conditions</code> . You must replace <code>Terms and Conditions</code> with your own text. If the Consent checkbox is toggled off, then <code>Login message</code> appears, over which you enter your message.
Consent checkbox	Toggle on Consent checkbox to require that the user clicks a check box before logging in. You can also display a message without a check box.
Details of login message	Message that the user sees when clicking the login message, for example, the text of the terms and conditions. You must enter some details in this text box.

6. Click **Save**.

How to Configure Your vSphere Client Timeout Value

By default, the vSphere Client sessions close after 120 minutes of idle time, and require the user to log in again to resume using the client. You can change the timeout value by editing the `webclient.properties` file.

You can configure the timeout of the vSphere Client by changing the value of the `session.timeout` property in the `webclient.properties` file.

Edit the Properties File of the vSphere Client

1. On the computer where the vSphere Client is installed, navigate to the `/etc/vmware/vsphere-ui` file path where the `webclient.properties` file is located.
2. Edit the file to include the line `session.timeout = <custom_value>`, where `<custom_value>` is the timeout value in minutes. Uncomment the line, if necessary.
To set the client to never timeout, specify a negative or 0 value for the timeout.
For example, to set the timeout value to 60 minutes, include the line `session.timeout = 60`.
3. Restart the VMware vSphere Client service.

Change the Client Configuration in the vSphere Client

1. In the vSphere Client, navigate to **Home > Administration**.
2. Under **Deployment**, select **Client Configuration**.
3. In the **Client Configuration** pane, click **Edit**.
The **Edit client configuration** dialog box appears.
4. Enter a new value in minutes in the session timeout text box.
5. Click **Save**.
6. Restart the VMware vSphere Client service.

How Do You Refresh Data in the vSphere Client

You must manually refresh the data in the vSphere Client to see changes made to objects by other users during your session.

For performance reasons, the vSphere Client does not continuously refresh data on all objects in the inventory. All changes that you make during your current session are immediately reflected in the client user interface. Changes made by other users or in other sessions are not reflected until you manually refresh the data.

To update all data in the current vSphere Client view, click the refresh icon () .

The client view is updated.

How to Use Your vSphere Client Inventory

Learn how to navigate, search, and sort the vSphere Client inventory.

How to Use Your Global Inventory Lists Navigator

You can use the vSphere **Global Inventory Lists** navigator to browse and select objects in the vSphere inventory as an alternative to the hierarchical inventory tree.

Unlike the inventory tree, which presents hierarchical arrangements of parent and child objects arranged in the **Hosts and Clusters**, **VMs and Templates**, **Storage**, and **Networking** inventory lists, the Global Inventory Lists navigator presents a list of all related objects in the inventory. You can navigate from an object to its related objects, regardless of their type.

1. From the vSphere Client home page, click **Global Inventory Lists**.
2. In the left pane, select one of the object or resource categories to view objects of that type.
For example, click **Hosts** to view the hosts in the vSphere inventory.

3. In the left navigation pane, click an object from the list once to display information about the object.
4. Optional: Click the object again to open it.
Opening an object brings it to the top of the navigator and displays related object categories beneath it.
For example, opening a host allows you to see the child resource pools, virtual machines, datastores, networks, and distributed switches associated with this host.
5. To access additional information or manage the selected object, click one of the tabs in the center pane.

Option	Description
Summary	You can view basic status and configuration for an object.
Monitor	You can view alarms, performance data, resource allocation, events, and other status information for an object.
Configure	Depending on the selected object, you can edit settings, alarm definitions, tags, permissions, and so on.
Permissions	You can view, add, change, and delete permissions. This tab is only available for logged in users with administration privileges.
Related Objects tabs	You can view and manage the objects related to the object that you selected. For example, if you select a host, the tabs that you see are VMs , Resource Pools , Datastores , and Networks .
Updates	Depending on the selected object, you can check the hardware compatibility of a host against the VMware Compatibility Guide, check the status of virtual machines, upgrade the VMware Tools version or the virtual hardware version of the virtual machines, and perform vSphere Lifecycle Manager operations. Note: The Updates tab is available only for certain types of inventory objects, for example, hosts, virtual machines, and clusters.

Inventory Keyboard Shortcuts in the vSphere Client

Discover some of the inventory keyboard shortcuts that you can use to quickly navigate to different inventories in the vSphere Client.

Table 143: Inventory Keyboard Shortcuts

Keyboard Combination on a Windows machine	Keyboard Combination on a Mac machine	Action
Ctrl+Alt+Home	Control+Option+Home	Home
Ctrl+Alt+1	Control+Option+1	Shortcuts
Ctrl+Alt+2	Control+Option+2	Hosts and Clusters
Ctrl+Alt+3	Control+Option+3	VMs and Templates
Ctrl+Alt+4	Control+Option+4	Storage
Ctrl+Alt+5	Control+Option+5	Networking
Ctrl+Alt+6	Control+Option+6	Content Libraries
Ctrl+Alt+7	Control+Option+7	Workload Management
Ctrl+Alt+8	Control+Option+8	Global Inventory Lists
Ctrl+Alt+R	Control+Option+R	Refresh
Ctrl+Alt+S repeat twice	Control+Option+S repeat twice	Search

How to Sort Your vSphere Client Inventory Objects

You can use sortable columns and input filters to sort and find objects in your vSphere Client inventory that meet certain criteria.

You can sort list views by column. These list views can be selected in the **Global Inventory Lists** pane or found in search results.

For example, you can sort virtual machines by name, provisioned space, used space, and so on.

1. In the vSphere Client home page, select **Global Inventory Lists**.
2. To open an object list view, select an object type from the list.
If objects from the selected type are not present in the inventory, a blank page appears.
3. Optional: Click the column you want to sort objects by. You can click again on that column to reverse the sorting order.
4. Add or remove columns from the object list view:
 - a) Click **Manage Columns** at the bottom left corner of the object list view.
The **Show Columns** pop-up appears.
 - b) Optional: To show a column in the object list view, select the check box for the column from the list.
 - c) Optional: To hide a column from the object list view, deselect the check box for the column from the list.

How to Use Quick Filters to Find Objects in Your vSphere Client Inventory

From a list view, you can use the quick filter field to filter objects by entering specific values.

You can use input filters to sort and find objects in your vSphere Client inventory that meet certain criteria. You can filter inventory objects by name, state, status, and so on.

1. In the vSphere Client home page, select **Global Inventory Lists**.
2. To open an object list view, select an object type from the list.
If objects from the selected type are not present in the inventory, a blank page appears.
3. Click the drop-down menu in the right upper corner of the inventory object list table, and select **Quick Filter**.
4. Enter your filtering parameter in the empty filter field. For example, in the virtual machine list view, to filter for all virtual machines that are in a powered on state, you must enter `powered on`.
A list of inventory objects that meet your filtering criteria is displayed in your selected sorting order.

To clear the filtered list of inventory objects, delete the filtering parameter from the filter field.

How to Use Advanced Filters to Find Objects in Your vSphere Client Inventory

You can use the advanced filter options to create a set of filtering criteria and find specific objects in your vSphere Client.

1. In the vSphere Client home page, select **Global Inventory Lists**.
2. To open an object list view, select an object type from the list.
If objects from the selected type are not present in the inventory, a blank page appears.
3. Click the drop-down menu in the right upper corner of the inventory object list table, and select **Advanced Filter**.
4. Click **Add New**.
A pop-up window appears.

5. From the drop-down menu under Filter, select the filtering parameter.
6. Depending on the parameter you select, you must select values from a pre-defined list, or enter data in a filter operator.
 - From the Value list, select all values that you want to apply as filtering criteria.
 - From the filter operator, select how the filtering criteria must relate to each other and enter the filter value.

Note: If you want to specify the filtering criteria further, you can click **Add Condition**, and enter the required details for the second filter operator.
7. Click **Apply**.
A list of inventory objects that meet your filtering criteria is displayed in your selected sorting order. The filters that you apply are displayed on top of the inventory list.
Note: You can apply multiple advanced filters to the inventory list. They are combined with the logical operator *and*.
8. Optional: To edit a filter that you applied, click on the filter displayed on top of the inventory list, edit the filtering criteria, and click **Apply**.
9. Optional: To hide all filters applied to the inventory list, click **Hide All**.
The filters continue to be applied.

To delete a filter that you applied, click the remove icon () next to the filtering criterion. If the filter consists of more than one filtering criteria, you should remove each one separately.

To delete all filters applied to the inventory list, click **Clear All**.

How to Search Your vSphere Client Inventory

With the vSphere Client, you can search the inventory for objects that match specified criteria. You can search the inventories of all vCenter Server systems.

You can only view and search for inventory objects that you have permission to view.

Note: If your permissions change while you are logged in, the search service might not immediately recognize these changes. To ensure that your search is performed with up-to-date permissions, log out of all your open sessions and log in again before you perform the search.

Perform a Quick Search

A quick search checks all types of objects for the specified search term within the name or other properties of the object.

1. Enter the search term in the search box at the top of the vSphere Client window.
Multiple search terms in a quick or simple search are treated as if they are connected by ORs. For example, searching for `example machine` finds all objects with names containing either "example" or "machine".
The search results appear below the search box as you type. The number of items displayed is limited to 10.
2. To run a saved search, click the saved search icon and select a saved search query.
3. To display an item in the inventory, click the item from the drop-down list of the search results.
4. To see more search results or more details about the search results, click the summary of the search results from the drop-down list of search results.
 - a. To display an object in the inventory, click the object in the search results page.
The search results are listed in tables arranged by object type. For example, if a search finds hosts and virtual machines, the following tables appear: **Hosts**, which shows only host results and **Virtual Machines**, which shows only virtual machine results.

Save a Search

You can save search queries in the vSphere Client so you can rerun, rename, or delete them later.

1. In a vSphere Client window, enter a query for a simple search in the search box.
2. Click the summary of the search results from the drop-down list with search results.
3. On the search results page, click **Save Search**.
4. Enter a name for the search query. Names must be lowercase with no spaces.
5. Click **Save**.
The search query you entered is saved. You can reload the query later and repeat the search.

Manage a Saved Search

In the vSphere Client, you can rename, delete, or create a duplicate of a saved search query.

1. At the top of the vSphere Client window, click inside the search box.
2. Click the saved search icon and select a saved search query.
3. On the Search Results page, click **Actions** and select one of the options:

Option	Description
Save as	Creates a duplicate of the saved search query.
Rename	Renames the saved search query.
Delete	Deletes the saved search query.

4. Confirm the changes.

How to Export Your vSphere Inventory Lists

In the vSphere Client, you can export the contents of an inventory list view to a CSV file.

1. In the vSphere Client, click **Menu** and select **Global Inventory Lists**.
2. In the left pane, select one of the object or resource categories to view objects of that type.
For example, click **Hosts** to view the hosts in the vSphere inventory.
3. Select whether you want all columns or your current selection of columns to be listed in the CSV file.
 - a) Click **Manage Columns** at the bottom left corner of the object list view.
The **Show Columns** pop-up appears.
 - b) Optional: To show a column in the object list view, select the check box for the column from the list.
 - c) Optional: To hide a column from the object list view, deselect the check box for the column from the list.
 - d) Optional: To show all columns, click **SELECT ALL**.
4. Click **Export** at the bottom left corner of the list view.

The CSV file is generated and available for download.

How to Manage Your vCenter Server Client Plug-Ins

Learn how to monitor, add, or remove vCenter Server client plug-ins. You can monitor client plug-in activities including installation progress, failures, deployments, upgrades, and uninstallations.

From the vSphere Client, you can monitor vCenter Server client plug-ins using the **Recent Tasks** pane, the **Task Console**, or the **Administration > Client Plugins** view. In addition to monitoring installation progress, you can troubleshoot installation failures and plug-in incompatibilities. You can also add and remove client plug-ins.

For more information about how to develop client plug-ins, see the *Developing Remote Plug-ins with the vSphere Client SDK* documentation on <https://developer.vmware.com>.

View Details about a Client Plug-In

1. In the vSphere Client home page, click **Administration > Solutions > Client Plugins**.
The **Client Plugins** main view shows a list of all deployed client plug-ins. You can view information about the name, version, deployment status, type, certification status, vendor, and plug-in ID of each client plug-in.
2. In the Name column, click a client plug-in name to open the client plug-in details view.
You can use the client plug-in details view to view more information about the client plug-in instances, control the visibility of the client plug-ins on specific vCenter Server systems, and remove client plug-ins from the vSphere Client.

Hide and Show a Client Plug-In on a vCenter Server Instance

If your environment includes several vCenter Server systems joined in Enhanced Linked Mode, the vSphere Client can load views from client plug-ins added to any of the linked vCenter Server instances. You configure what is visible in a specific vCenter Server system, by using the **Hide** and **Show** options from the details view of a selected client plug-in.

1. From the vSphere Client Home menu, select **Administration**.
2. Under **Solutions**, click **Client Plugins**.
3. From the Name column, click the client plug-in name.
The detailed client plug-in view opens.
4. Change the visibility of a client plug-in instance on a vCenter Server instance.
 - a. Select the check box on the client plug-in instance.
 - To hide a client plug-in in the vSphere Client UI for a vCenter Server instance, select the check box on the instance version and click **Hide**.
 - To make a client plug-in instance visible in the vSphere Client UI for a vCenter Server instance, select the client plug-in version and click **Show**.
5. Confirm your choice in the pop-up dialog box that appears.

Add a Client Plug-In with the vSphere Client

Starting with vSphere 8.0, you can deploy remote client plug-ins using the vSphere Client. In an environment with several vCenter Server instances connected with Enhanced Linked Mode, you can select on which vCenter Server systems to deploy the client plug-in.

1. From the vSphere Client Home menu, select **Administration**.
2. Under **Solutions**, click **Client Plugins**.
3. In the **Client Plugins** pane, click **Add**.
The **Install new solution** wizard opens.
4. On the **Select an OVF template** page, specify the location of the client plug-in OVF descriptor file and click **Next**.

Option	Action
URL	Enter the URL to the OVF descriptor file located on the Internet. The only supported URL sources are HTTPS. Note: Make sure that all elements of the client plug-in are at the same location as the OVF descriptor file. Missing files might lead to failure in plug-in deployment.
Local file	Click Upload files and select all files associated with the client plug-in.

5. On the **Select a name and folder** page, enter a unique name for the virtual machine, select a deployment location, and click **Next**.
The default name for the virtual machine is the same as the human-readable name of the plug-in defined in the OVF file. If you change the default name, choose a name that is unique within each vCenter Server virtual machine folder.
6. On the **Select a compute resource** page, select a resource where to run the deployed client plug-in VM, and click **Next**.
7. On the **Review details** page, verify the client plug-in details and click **Next**.
8. On the **License agreement** page, accept the end-user license agreements and click **Next**.
9. On the **Select storage** page, define where and how the files for the deployed client plug-in are stored.
 - Select a storage policy for the virtual machine.
This option is available only if storage policies are enabled on the destination resource.
 - Select the **Disable Storage DRS for this virtual machine** check box to deactivate Storage DRS for the virtual machine.
 - Select a datastore to store the deployed client plug-in.
The plug-in manifest file and virtual disk files are stored on the datastore. Select a datastore large enough to accommodate the virtual machine and all associated virtual disk files.
10. On the **Select networks** page, select a source network and map it to a destination network. Define the IP allocation settings of the client plug-in server. Click **Next**.
The Source Network column lists all networks as defined in the OVF descriptor file.
11. On the **Customize template** page, view the read-only settings of the client plug-in and click **Next**.
12. On the **Associate vCenter Servers** page, select the vCenter Server instances on which the client plug-in is deployed and click **Next**.
13. On the **Ready to complete** page, review your selections and click **Finish**.

As a result new tasks for downloading and installing the client plug-in appear in the **Recent Tasks** pane. After successful completion of the tasks, refresh your browser to see the newly added client plug-in.

Remove a Client Plug-In from a vCenter Server Instance

You can remove a client plug-in from the vSphere Client UI for a vCenter Server instance. In an environment with several vCenter Server instances connected with Enhanced Linked Mode, you can select from which vCenter Server systems to remove the client plug-in.

Note:

When you remove a client plug-in from the vSphere Client UI, the virtual machine on which the client plug-in instance is running is not removed from the vCenter Server instance. You must delete the virtual machine from the datastore to complete the client plug-in removal.

1. From the vSphere Client Home menu, select **Administration**.
2. Under **Solutions**, click **Client Plugins**.
3. From the Name column in the **Client Plugins** pane, select the client plug-in.
The detailed client plug-in view opens.
4. Select the check box on one of the client plug-in instances and click **Remove**.
5. In the confirmation dialog box, click **Yes** to remove the client plug-in instance.

As a result, the client plug-in is undeployed from the vCenter Server instance. Make sure that you refresh your browser to remove the client plug-in from the vSphere Client UI.

How to Start, Stop, and Restart Services in Your vSphere Client

In the vSphere Client, you can start, stop, and restart services that are running on vCenter Server. You can restart services upon a configuration change or for suspected functional or performance issues.

Verify that the user you use to log in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

1. Log in to the vCenter Server Management Interface.
 - Log in to the vCenter Server Management Interface directly.
 - In the vSphere Client home page, select **Administration > System Configuration**. Click a node from the list.
2. On the vCenter Server Management Interface home page, click **Services**.
3. Select a service from the service list.
4. From the top menu, click **Restart**, **Start**, or **Stop**.
Unavailable options are dimmed.

Note: Restarting the Content Library Service also restarts the Transfer Service and the OVF Service. The Content Library Service, the Transfer Service, and the OVF Service run on the same Tomcat server.

How to Attach a File to a VMware Service Request

Learn how to attach files, such as log files or screenshots, to VMware Service Requests directly from the vSphere Client.

If you already have a Service Request ID with VMware, you can use the vSphere Client to upload the system log bundles directly to your service request.

1. From the vSphere Client sidebar, click **Administration**.
2. In the Administration panel, under Support, click **Upload File to Service Request**.
3. Click the **Upload File to Service Request** button.
4. Enter the Service Request ID.
5. Click **Browse** and select the file to attach.
6. Click **Upload**.

How to Provide Feedback for vSphere

Learn how to use the vSphere Client to provide feedback and submit feature requests for improvements to vSphere.

How to Provide Feedback with Your vSphere Client

Use the improved feedback tool in the vSphere Client to provide timely feedback to our developers.

1. From the vSphere Client home screen, click the feedback icon in the top right corner.
2. Select the type of feedback you want to give and enter your feedback in the **Description** window.
3. Optional: Provide an email address and screenshots or other images.
4. Click **Send**.

How to Use the vSphere Ideas Portal to Submit Suggestions for New Features and Feature Requests

Learn how to provide suggestions for new vSphere features or share feature requests through the vSphere Ideas portal. All users with valid Broadcom Support Portal accounts can access the portal.

Verify that you have a valid Broadcom Support Portal account.

The vSphere Ideas portal is public. When you publish new ideas and feature requests, consider whether you want to share personal data.

You can access the vSphere Ideas portal at <https://vsphere.ideas.aha.io/> by entering a valid Broadcom Support Portal user name and password.

Alternatively, you can access the vSphere Ideas portal by using the vSphere Client.

1. On the vSphere Client home page, click the feedback icon in the top-right corner.
2. In the **Send Feedback** dialog box, click **Idea**.
3. Click **Visit ideas portal**.
4. On the Broadcom Support Portal login page, enter a valid user name and password.
You are redirected to <https://vsphere.ideas.aha.io/>.
5. On the vSphere Ideas portal home page, click **Add a new idea**.
6. On the **Add a new idea** page, enter the details about your idea and click **Share idea**.

Your idea appears in the **Recent** tab. All users of the vSphere Ideas portal can view, subscribe, vote, and comment on your idea.

Configuring Customer Experience Improvement Program

You can participate in the Customer Experience Improvement Program (CEIP) to provide anonymous feedback or information to VMware for quality, reliability, and functionality improvements of VMware products and services.

VMware Customer Experience Improvement Program

vCenter Server participates in VMware's Customer Experience Improvement Program (CEIP).

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

Join or Leave Customer Experience Improvement Program in vSphere Client

You can join your vCenter Server to the Customer Experience Improvement Program (CEIP), or decide to leave the CEIP at any time. To leave and rejoin your host to the CEIP, see *vSphere Single Host Management - VMware Host Client* documentation.

Obtain the user name and password of the administrator account.

1. From the vSphere Client login page, log in to vCenter Server by using the credentials of the administrator account.
2. On the vSphere Client home page, click **Administration**.
3. Under Deployment, click **Customer Experience Improvement Program**.
4. Click **Join** to join the CEIP or **Leave CEIP** to leave the Program.

Organizing Your vSphere Inventory

Plan how to set up your virtual inventory and consider how the virtual machines that it will support are going to be used and administered.

A large vSphere implementation might contain several virtual data centers with a complex arrangement of hosts, clusters, resource pools, and networks. It might involve multiple vCenter Server systems connected using Enhanced Linked Mode. Smaller implementations might require a single virtual data center with a much less complex topology.

Here are the questions to answer as you create and organize an inventory of virtual objects:

- Will some virtual machines require dedicated resources?

- Will some virtual machines experience periodic spikes in workload?
- Will some virtual machines need to be administered as a group?
- Will some virtual objects require one set of system permissions, while other objects will require a different set of permissions?
- Do you want to use multiple vSphere Standard Switches, or you want to have a single vSphere Distributed Switch per data center?
- Do you want to use vMotion and Distributed Resource Management with certain virtual machines but not others?

The left pane of the vSphere Client displays your vSphere inventory. You can add and arrange objects in any way with the following considerations:

- The name of an inventory object must be unique with its parent.
- vApp names must be unique within the Virtual Machines and Templates view.
- System permissions are inherited and cascade. When you assign a system permission to an object in the inventory, the same permission propagates down the object hierarchy.

Tasks for Organizing Your Inventory

Populating and organizing your inventory involves the following activities:

Note: vCenter Server supports an 80 character limit for all its entities. When you create a new inventory object, if the name exceeds this limit, the task fails and a warning appears.

- Creating data centers.
- Adding hosts to the data centers.
- Organizing inventory objects in folders.
- Setting up networking by using vSphere Standard Switches or vSphere Distributed Switches. To use services such as vMotion, TCP/IP storage, VMware vSAN™, and Fault Tolerance, setup VMkernel networking for these services. For more information, see *vSphere Networking*.
- Configuring storage systems and creating datastore inventory objects to provide logical containers for storage devices in your inventory. See *vSphere Storage*.
- Creating clusters to consolidate the resources of multiple hosts and virtual machines. You can enable vSphere HA and vSphere DRS for increased availability and more flexible resource management. See *vSphere Availability* for information about configuring vSphere HA, and *vSphere Resource Management* for information about configuring vSphere DRS.
- Creating resource pools to provide logical abstraction and flexible management of the resources in vSphere. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources. See *vSphere Resource Management* for details.

How Do You Create a vSphere Data Center

A virtual data center is a container for all the inventory objects required to complete a fully functional environment for operating virtual machines. You can create multiple data centers to organize groups of environments to meet different user needs.

Required privileges:

• **Datacenter > Create datacenter**

For example, you can create a data center for each organizational unit in your enterprise or create some data centers for high-performance environments and other data centers for less demanding environments.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Right-click the vCenter Server object and select **New Datacenter**.
3. Optional: Enter a name for the data center and click **OK**.

Add hosts, clusters, resource pools, vApps, networking, datastores, and virtual machines to the data center.

How to Create a Folder in Your vSphere Inventory

In the vSphere Client, you can create folders and use them to group objects of the same type for easier management. For example, you can apply a common set of permissions to the folder and these permissions apply to all objects grouped in the folder.

A folder can contain other folders, or a group of objects of the same type. For example, one folder can contain both virtual machines and another folder that contains virtual machines, but it cannot contain both hosts and a folder that contains virtual machines.

1. In the vSphere Client, select either a data center or another folder as a parent object for the folder that you want to create.
2. Right-click the parent object and click **New Folder**.
 - If the parent object is a folder, the new folder is of the same type as the parent folder - it can contain only objects of the same type that the parent folder contains.
 - If the parent object is a data center, you can create one of four types of folders: **Host and Cluster** folders, **Network** folders, **Storage** folders, and **VM and Template** folders.
3. Enter a name for the folder and click **OK**.

Move objects into the folder by right-clicking the object and selecting **Move To**. Select the folder as the destination. You can also move an object by dragging it to the destination folder.

How to Add a Host to Your vSphere Data Center or Folder

In the vSphere Client, you can add hosts under a data center object, a folder object, or a cluster object. If a host contains virtual machines, those virtual machines are also added to the inventory.

- Verify that a data center or a folder exists in the inventory.
- Obtain the user name and password of the root user account for the host.
- Verify that hosts behind a firewall are able to communicate with the vCenter Server system and all other hosts through port 902 or another custom-configured port.
- Verify that all NFS mounts on the host are active.
- Verify that you have the proper privileges. Different sets of privileges apply when you add multiple hosts to a cluster and a single host to a cluster or a data center. For more information, see [Required Privileges for Common Tasks](#) in the *vSphere Security* documentation.

- If you want to add a host with more than 512 LUNs and 2,048 paths to the vCenter Server inventory, verify that the vCenter Server instance is suitable for a large or an x-large environment.
1. In the vSphere Client, navigate to a data center or folder within a data center.
 2. Right-click the data center or folder and select **Add Host**.
 3. On the **Name and location** page, enter the IP address or the name of the host and click **Next**.
 4. On the **Connection settings** page, enter administrator credentials and click **Next**.
 5. On the **Host summary** page, review the summary of the host details and click **Next**.
 6. On the **Host lifecycle** page, the check box **Manage host with an image** check is selected by default.
 - If you want to manage the host with an image, leave the check box selected and click **Next**.
 For more information about using an image to manage the lifecycle of a standalone ESXi host that is part of your vCenter Server inventory, see [Managing Standalone ESXi Hosts with vSphere Lifecycle Manager Images](#) in the *Managing Host and Cluster Lifecycle* documentation.
 - If you do not want to manage the host with an image, deselect the check box and click **Next**.
 7. If you manage the host with an image, on the **Image** page, set up the desired image and click **Next**.
 - a) Select an ESXi version.
 - b) Optional: Select a vendor add-on and the version of the add-on.
 8. On the **Assign license** page, add a license to the host.
 - Assign an already existing license and click **Next**.
 - Assign a new license.
 - a. Click **Create New Licenses**. The **Add Host** wizard minimizes in Work in Progress and the **New Licenses** wizard appears.
 - b. Enter or copy and paste the new license key from My VMware and click **Next**.
 - c. Enter a new name for the license and click **Next**.
 - d. Review the new license and click **Finish**.
 - e. On the **Assign license** page, click **Next**.
 9. Optional: Select a lockdown mode option to deactivate the remote access for the administrator account after vCenter Server takes control of this host and click **Next**.
 10. Optional: On the **VM location** page, select a location for the virtual machines that reside on the host and click **Next**.
 11. Review the summary and click **Finish**.

A new task for adding the host appears in the Recent Tasks pane. It might take a few minutes for the task to complete.

How Do You Create and Configure Clusters in the vSphere Client

You can create clusters in the vSphere Client and then configure them manually or by using the Quickstart workflow.

A cluster is a group of hosts. When a host is added to a cluster, the resources of the host become part of the resources of the cluster. The cluster manages the resources of all hosts that it contains.

When you create clusters, you can enable vSphere High Availability (HA), vSphere Distributed Resource Scheduler (DRS), and the VMware vSAN features.

Starting with vSphere 7.0, you can create a cluster that you manage with a single image. By using vSphere Lifecycle Manager images, you can easily update and upgrade the software and firmware on the hosts in the cluster. Starting with vSphere 7.0 Update 2, during cluster creation, you can select a reference host and use the image on that host as the

image for the newly created cluster. For more information about using images to manage ESXi hosts and clusters, see the *Managing Host and Cluster Lifecycle* documentation.

Starting with vSphere 7.0 Update 1, vSphere Cluster Services (vCLS) is enabled by default and runs in all vSphere clusters. vCLS ensures that if vCenter Server becomes unavailable, cluster services remain available to maintain the resources and health of the workloads that run in the clusters. For more information about vCLS, see [vSphere Cluster Services](#).

How Do You Create a vSphere Cluster with the vSphere Client

You can create a new vSphere cluster object by using the vSphere Client.

- Verify that a data center, or a folder within a data center, exists in the inventory.
- Verify that hosts have the same ESXi version and patch level.
- Obtain the user name and password of the root user account for the host.
- If you want to extend a cluster with initially configured networking, verify that hosts do not have a manual vSAN configuration or a manual networking configuration.
- To create a cluster that you manage with a single image, review the requirements and limitations information in the [Managing Host and Cluster Lifecycle](#) documentation.

Required privileges:

- **Host > Inventory > Create cluster**

Starting with vSphere 7.0, the clusters that you create can use vSphere Lifecycle Manager images for host updates and upgrades.

A vSphere Lifecycle Manager image is a combination of vSphere software, driver software, and desired firmware with regard to the underlying host hardware. The image that a cluster uses defines the full software set that you want to run on all ESXi hosts in the cluster: the ESXi version, additional VMware-provided software, and vendor software, such as firmware and drivers.

The image that you define during cluster creation is not immediately applied to the hosts. If you do not set up an image for the cluster, the cluster uses baselines and baseline groups. Starting with vSphere 7.0 Update 2, during cluster creation, you can select a reference host and use the image on that host as the image for the newly created cluster. For more information about using images and baselines to manage hosts in clusters, see the [Managing Host and Cluster Lifecycle](#) documentation.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select a data center.
3. Right-click the data center and select **New Cluster**.
4. Enter a name for the cluster.
5. Select DRS, vSphere HA, or vSAN cluster features.

Option	Description		
To use DRS with this cluster	1. Slide the switch to the right to enable the DRS service. 2. Click the info icon on the left to see the Default Settings for the DRS service. The default values are: — Automation Level: Fully Automated Migration — Threshold: 3		
To use vSphere HA with this cluster	1. Slide the switch to the right to enable the vSphere HA service. 2. Click the info icon on the left to see the Default Settings for the vSphere HA service. You are presented with the following default values:		
	<table> <tr> <td>Host Monitoring:</td><td>Enabled</td></tr> </table>	Host Monitoring:	Enabled
Host Monitoring:	Enabled		

Option	Description	
	Admission Control:	Enabled
	VM Monitoring:	Disabled
To use vSAN with this cluster	Slide the switch to the right to enable the vSAN service. For more information on vSAN, see Creating a vSAN Cluster in the <i>vSAN Planning and Deployment</i> documentation.	

You can override the default values later on in the workflow.

- Optional: To create a cluster that you manage with a single image, select the **Manage all hosts in the cluster with a single image** check box.

For information about creating a cluster that you manage with a single image, see the [Managing Host and Cluster Lifecycle](#) documentation.

- Click **Next**.

- Review the cluster details and click **Finish**.

The cluster appears in the vCenter Server inventory. The **Quickstart** workflow appears under **Configure > Configuration**.

You have created an empty cluster in the vCenter Server inventory.

You can use the Quickstart workflow to easily configure and expand the cluster. You can also skip the Quickstart workflow and continue configuring the cluster and its hosts manually.

How to Extend and Configure Your vSphere Cluster by Using the Quickstart Workflow

You can extend and configure a cluster by using the simplified Quickstart workflow in the vSphere Client. Quickstart groups common tasks and offers configuration wizards that guide you through the process of configuring and extending a cluster.

After you provide the required information on each wizard, your cluster is configured based on your input. When you add hosts using the Quickstart workflow, hosts are automatically configured to match the cluster configuration.

Once the cluster creation operation completes, you are redirected to the **Cluster quickstart** page located under **Configure > Configuration**. The Quickstart workflow consists of three cards for quickly extending and configuring your clusters.

Table 144: Quickstart workflow cards for extending and configuring a cluster

Cluster Quickstart Workflow Card	Description
Cluster basics	You can edit the cluster name and activate or deactivate cluster services. The card lists the services you enabled.
Add hosts	You can add new ESXi hosts. After the hosts are added, the card shows the total number of hosts present in the cluster and displays health check validation for those hosts.
Configure cluster	You can configure network settings for vMotion and vSAN traffic, review and customize cluster services, and set up a vSAN datastore. After the cluster is configured, the card provides details on configuration mismatch and reports cluster health results through the vSAN Health service.

Characteristics of a Cluster Configured Using the Quickstart Workflows

A vSphere cluster that is configured using the Quickstart workflow has the following characteristics.

- Hosts are of version ESXi 6.7 Update 2 and later.
- Hosts end up having similar configuration, including network settings.
- Cluster configuration is based on the recommended default settings for networking and different services.
- Licenses cannot be assigned using the Quickstart workflow. You must manually assign the required licenses. See [vSphere License Management](#).

Skipping Quickstart

The **Skip Quickstart** button prompts you to continue configuring the cluster and its hosts manually. To confirm exiting the simplified configuration workflow, click **Continue**. After you dismiss the **Cluster quickstart** workflow, you cannot restore it for the current cluster.

How to Add an ESXi Host to Your vSphere Cluster by Using the Quickstart Workflow

You can add new and existing ESXi hosts to a cluster in the vCenter Server inventory.

- Verify that hosts have the same ESXi version and patch level.
- Obtain the user name and password of the root user account for the host.
- If you want to extend a cluster with initially configured networking, verify that hosts do not have a manual vSAN configuration or a manual networking configuration.
- Verify that you have the proper privileges. Different sets of privileges apply when you add multiple hosts to a cluster and a single host to a cluster or a data center. For more information, see [Required Privileges for Common Tasks](#) in the *vSphere Security* documentation.
- To add a host to a cluster that you manage with a single image, see the [Managing Host and Cluster Lifecycle](#) documentation.

You can also add hosts to a DRS cluster. For more information, see the *vSphere Resource Management* documentation.

When you add the first three hosts to the cluster, vSphere Cluster Services (vCLS) agent virtual machines are added by default to the cluster. A quorum of up to three vCLS agent virtual machines are required to run in a cluster, one agent virtual machine per host. For more information about vCLS, see [vSphere Cluster Services](#).

Note: After you configure the cluster by using Quickstart, if you modify any cluster networking settings outside of Quickstart, you cannot use the Quickstart workflow to add and configure new hosts to the cluster.

1. In the vSphere Client, navigate to a cluster within a data center.
2. On the **Configure** tab, select **Configuration > Quickstart**.
3. Click **Add** in the **Add hosts** card.
4. On the **Add hosts** wizard, add new or existing hosts to the cluster.
 - Add hosts that are not part of the vCenter Server inventory.
 - a. Click the **New hosts** tab.
 - b. Populate the IP Address and credentials text boxes for those hosts.
 - c. To add more new hosts, click the **Add Host** button.
 - d. To reuse the credentials for all added hosts, select the **Use the same credentials for all hosts** check box.
 - Add hosts that are managed by your vCenter Server instance and are in the same data center as your cluster. The hosts must not be part of another cluster on the vCenter Server instance.
 - a. Click the **Existing hosts** tab.
 - b. From the list, select the hosts that you want to add to the cluster.

5. Click **Next**.

The **Host summary** page lists all hosts that will be added to the cluster and related warnings.

Note:

If a host cannot be validated automatically by the system, you are prompted to manually validate its certificate and accept its thumbprint in the **Security Alert** pop-up.

6. On the **Host summary** page, review the details of the added hosts and click **Next**.

7. Optional: On the **Import Image** page, select the host whose image to use as the image for the cluster.

The **Import Image** page appears when you add hosts to a cluster managed with a single image. For information about adding a host to a cluster that you manage with a single image, see the [Managing Host and Cluster Lifecycle](#) documentation.

8. On the **Ready to complete** page, review the IP addresses or FQDN of the added hosts and click **Finish**.

Review the number of added hosts and the health check validation, performed by the vSAN Health service, in the **Add hosts** card.

9. Optional: Click **Re-validate** to trigger the validation of the hosts.

Note:

If an error occurs, it is visible in the **Recent Tasks** pane only.

All hosts are placed in maintenance mode and added to your cluster. You can manually exit the maintenance mode.

Configure your cluster default settings through the Quickstart workflow.

How to Configure Your vSphere Cluster by Using the Quickstart Workflow

Learn how to use the Quickstart workflow to quickly configure the settings of a vSphere cluster.

To configure the host networking settings on your host and to customize the cluster settings, start the **Configure cluster** wizard, part of the cluster Quickstart workflow. Learn how to use Quickstart to configure a vSAN cluster, see *Use Quickstart to Configure a vSAN Cluster* in the *vSAN Planning and Deployment* documentation.

1. In the vSphere Client, navigate to a cluster.
2. On the **Configure** tab, select **Configuration > Quickstart**.

The **Cluster Quickstart** page appears.

Note:

To configure your cluster host networking and services manually by referring to different parts of the vSphere software, click the **Skip quickstart** button. If you dismiss the **Cluster quickstart** workflow, you cannot restore it, and you have to configure manually any hosts that you add to this cluster in the future.

3. In the **Configure cluster** card, select **Configure**.

The **Configure cluster** wizard appears.

4. On the **Distributed switches** page, configure the cluster networking.

Alternatively, you can select the **Configure networking settings later** check box to hide all options that are related to cluster networking and to change the recommended host networking manually at a later time.

CAUTION:

After you select the **Configure networking settings later** check box, and complete the **Configure cluster** workflow, you cannot perform the cluster networking configuration in the future by using the **Configure cluster** wizard.

- a) In the **Distributed switches** section, specify the number of distributed switches to create from the drop-down menu.

Note: You can select up to three distributed switches.

The selected distributed switches are configured as part of this workflow and all hosts in the cluster connect to them.

- b) Enter a unique name for each of the distributed switches you are about to create.
- c) Optional: Click **Use Existing** to select an existing compatible distributed switch and an existing compatible distributed port group.
- d) In the **Port groups** section, set up the vMotion network by selecting a distributed switch from the drop-down menu and assigning a new default port group to it.
- e) In the **Physical adapters** section, for each physical network adapter (NIC), select the distributed switch name from the drop-down menu.

The new distributed switch must be assigned to at least one physical adapter.

Note: If you are using an existing distributed switch, the physical adapter selection must match the current mapping of the distributed switch. Any variation results in an error.

This mapping of physical NICs to the distributed switches is applied to all hosts in this cluster.

- f) Click **Next**.

5. If the vSphere DRS feature is enabled on the cluster, configure the networking options in the **vMotion traffic** page.

- a) Optional: Select the **Use VLAN** check box and enter an ID for the vMotion distributed port group.
- b) Optional: Select a protocol type from the drop-down menu.
- c) Select an IP configuration type from the drop-down menu.
- d) Optional: Populate the text boxes for each host in the cluster depending on the IP address type you need for setting up the networking.

If the IP address type is set to **DHCP**, DHCP is used to obtain the IP settings for the vMotion traffic.

- e) Click **Next**.

The **Advanced options** page appears.

6. On the **Advanced options** page, customize the cluster settings depending on the services you enabled.
 - a) If you enabled the vSphere HA feature during cluster creation, use the options in the **vSphere HA** section to activate or deactivate host failure monitoring, virtual machine monitoring, and admission control.
 If you enable admission control, you can specify the failover capacity by number of hosts.
 - b) If you enabled the vSphere DRS feature during cluster creation, the **Distributed Resource Scheduler** section is visible.
 1. Set the Automation level to Fully Automated, Partially Automated or Manual.
 2. Select one of the five migration settings from the Migration threshold drop-down menu.
 - c)
 1. Select a lockdown mode from the drop-down menu.
 2. Enter an NTP server address or FQDN.
 If you enter more than one NTP server addresses, separate them by using commas.
 3. Select a host update option from the drop-down menu.
 The settings are applied across all hosts in this cluster.
 - d) In the **Enhanced vMotion Compatibility** section, enable Enhanced vMotion™ Compatibility (EVC) and select the CPU model from the EVC mode drop-down menu.
 - e) Click **Next**.

The **Review** page appears.

7. Review the settings and select **Finish**.

The card closes, and the progress of the operation appears in the **Recent Tasks** pane.

You have created a fully configured cluster in the vCenter Server inventory.

Expand your cluster manually or by using the Quickstart workflow and the **Add hosts** card.

How to Extend Your vSphere Cluster with Host Networking Configuration

By using the Quickstart Workflow, you can extend a hyper-converged cluster by adding hosts and configuring their networking to match the cluster configuration.

- Verify that you have an existing cluster and hosts added to it.
- In the initial cluster configuration, you configured the host networking.
- Verify that hosts have the same ESXi version and patch level.
- Obtain the user name and password of the root user account for the host.
- If you want to extend a cluster with initially configured networking, verify that hosts do not have a manual vSAN configuration or a manual networking configuration.
- To add a host to a cluster that you manage with a single image, see the *Managing Host and Cluster Lifecycle* documentation.

Note: If you modify any cluster networking settings outside of Quickstart, you cannot use the Quickstart workflow to add and configure new hosts to the cluster.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters** and select a configured cluster.
2. Right-click the cluster and select **Add Hosts**.

The **Add hosts** wizard appears.

3. On the **Add hosts** page, add new or existing hosts to the cluster.
 - Add hosts that are not part of the vCenter Server inventory.
 - a. Click the **New hosts** tab.
 - b. Populate the IP Address and credentials text boxes for those hosts.
 - c. To add more new hosts, click the **Add Host** button.
 - d. To reuse the credentials for all added hosts, select the **Use the same credentials for all hosts** check box.
 - Add hosts that are managed by your vCenter Server instance and are in the same data center as your cluster.
 - a. Click the **Existing hosts** tab.
 - b. From the list, select the hosts that you want to add to the cluster.
4. Click **Next**.
The **Host summary** page lists all hosts that will be added to the cluster and related warnings.
Note: If a host cannot be validated automatically by the system, you are prompted to manually validate its certificate and accept its thumbprint in the **Security Alert** pop-up.
5. On the **Host summary** page, review the details of the added hosts and click **Next**.
6. Optional: On the **Import Image** page, select the host whose image to use as the image for the cluster.
The **Import Image** page appears when you add hosts to a cluster managed with a single image. For information about adding a host to a cluster that you manage with a single image, see the *Managing Host and Cluster Lifecycle* documentation.
7. On the **Ready to complete** page, review the IP addresses or FQDN of the added hosts and click **Finish**.
The **Extend Cluster Guide** page appears.
8. In the **Configure cluster** card, select **Configure**.
The **Configure cluster** wizard appears.
9. If the vSphere DRS feature is enabled on the cluster, configure the networking options in the **vMotion traffic** page.
 - a) Optional: Select the **Use VLAN** check box and enter an ID for the vMotion distributed port group.
 - b) Optional: Select a protocol type from the drop-down menu.
 - c) Select an IP type from the drop-down menu.
 - d) Optional: Populate the text boxes for each host in the cluster depending on the IP address type you need for setting up the networking.
If the IP address type is set to `DHCP`, DHCP is used to obtain the IP settings for the vMotion traffic.
10. Click **Next**.
The **Review** page appears.
11. Review the settings and select **Finish**.
The card closes, and the progress of the operation appears in the **Recent Tasks** pane.

After successful validation, your newly added hosts are configured as the existing hosts in your cluster and the **Configure** button in the **Configure hosts** card becomes inactive. You can only click **Re-validate** to verify the cluster configuration.

Add more hosts to the cluster.

How to Extend and Configure Your vSphere Cluster Manually

You can extend a cluster by manually adding hosts to that cluster. If you previously configured the cluster without setting up the host networking, the configuration of the existing hosts in the cluster is applied to the newly added hosts.

- Verify that you have an existing cluster and hosts added to it.

- Verify that hosts have the same ESXi version and patch level.
- Obtain the user name and password of the root user account for the host.
- During the initial cluster configuration, select the Configure networking settings later check box. For more information, see [How to Configure Your vSphere Cluster by Using the Quickstart Workflow](#).
- To add a host to a cluster that you manage with a single image, see the *Managing Host and Cluster Lifecycle* documentation.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters** and select a configured cluster.
2. Right-click the cluster and select **Add Hosts**.
The **Add hosts** wizard appears.
3. On the **Add hosts** page, add new or existing hosts to the cluster.
 - Add hosts that are not part of the vCenter Server inventory.
 - a. Click the **New hosts** tab.
 - b. Populate the IP Address and credentials text boxes for those hosts.
 - c. To add more new hosts, click the **Add Host** button.
 - d. To reuse the credentials for all added hosts, select the **Use the same credentials for all hosts** check box.
 - Add hosts that are managed by your vCenter Server instance and are in the same data center as your cluster. The hosts must not be part of another cluster on the vCenter Server instance.
 - a. Click the **Existing hosts** tab.
 - b. From the list, select the hosts that you want to add to the cluster.

4. Click **Next**.

The **Host summary** page lists all hosts that will be added to the cluster and related warnings.

Note:

If a host cannot be validated automatically by the system, you are prompted to manually validate its certificate and accept its thumbprint in the **Security Alert** pop-up.

5. On the **Host summary** page, review the details of the added hosts and click **Next**.
6. Optional: On the **Import Image** page, select the host from which an image must be imported to the cluster.

The **Import Image** page appears when you add hosts to a cluster managed with a single image. For information about adding a host to a cluster managed with a single image, see the *Managing Host and Cluster Lifecycle* documentation.

7. On the **Review** page, click **Finish**.
The **Cluster Quickstart** page appears.
8. Click **Skip Quickstart** to continue configuring the cluster manually.

CAUTION: Once you dismiss the Quickstart workflow, you cannot enable it again on the cluster.

Note: You must manually configure all hosts added to the cluster after dismissing the Quickstart workflow.

Manually configure all cluster settings and add more hosts to the cluster.

vSphere Tags and Custom Attributes

Tags and attributes allow you to attach metadata to objects in the vSphere inventory to make it easier to sort and search for these objects.

For vSphere tags and attributes, VMware Cloud on AWS supports the same set of tasks as an on-premises SDDC.

What Are vSphere Tags and How to Use Them

You can create, edit, and delete tag categories and vSphere tags. You can also assign vSphere tags to objects in the vCenter Server inventory and manage the user privileges for working with tags and categories.

A tag is a label that you can apply to objects in the vSphere inventory. When you create a tag, you assign that tag to a category. Categories allow you to group related tags together. When you define a category, you can specify the object types for its tags, and whether more than one tag in the category can be applied to an object.

For example, if you wanted to tag your virtual machines by guest operating system type, you can create a category called operating system. You can specify that it applies to virtual machines only and that only a single tag can be applied to a virtual machine at any time. The tags in this category might be Windows, Linux, and Mac OS.

Tags and categories can span multiple vCenter Server instances:

- If multiple on-premises vCenter Server instances are configured to use Enhanced Linked Mode, tags and tag categories are replicated across all these vCenter Server instances.
- When you use Hybrid Linked Mode, tags and tag categories are maintained across your linked domain. That means the on-premises SDDC and the VMware Cloud on AWS SDDC share tags and tag attributes.

How Do You Create and Edit a vSphere Tag Category

You use categories in your vSphere environment to group tags together and define how tags can be applied to objects. You create and edit a tag category from the vSphere Client.

The required privilege depends on the task that you want to perform.

Task	Privilege
Create a tag category	vSphere Tagging > Create vSphere Tag Category on any vCenter Server inventory object.
Edit a tag category	vSphere Tagging > Edit vSphere Tag Category on any vCenter Server inventory object.

You can create a tag category explicitly, as explained here, or as part of the tag creation process. Each tag has to belong to at least one tag category.

1. From the vSphere Client Home menu, click **Tags & Custom Attributes**.
2. Click the **Tags** tab and click **Categories**
3. Start the task that you want to perform.

Task	Steps to perform
Create a tag category	Click New .
Edit a tag category	Select a category and click Edit .

4. Edit the category settings.

Setting	Description
Category Name	The category name must be unique to the currently selected vCenter Server system.
Description	You can provide text in the description to describe the purpose or use of the category.
Tags Per Object	• If you select One Tag, you can apply only one tag from this category to an object. Use this option for categories whose tags are mutually exclusive. For example, if you have a category called Priority

Setting	Description
	with tags High, Medium, and Low, then each object should have only one tag because an object can have only one priority. If you select Many tags, you can apply more than one tag from the category to an object. Use this option for categories whose tags are not mutually exclusive. After you have set Tags Per Object, you can change from One Tag to Many Tags, but not from Many Tags to One Tag.
Associable Object Types	Select whether tags in this category can be assigned to all objects or only to a specific type of object, such as a virtual machine or a datastore. Changes to the associable object type are limited. - If you initially selected a single object type, you can later change the category to work for all object types. - If you initially selected All Objects, you cannot restrict the category later.

- Click **Create** or **Save**.

How Do You Delete a vSphere Tag Category

Learn how to remove a vSphere tag category from your vSphere environment.

Required privilege: **vSphere Tagging > Delete vSphere Tag Category** on any vCenter Server inventory object.

Deleting a category also deletes all tags associated with the category.

- From the vSphere Client Home menu, click **Tags & Custom Attributes**.
- Click the **Tags** tab and click **Categories**
- Select a category from the list and click **Delete**.
The **Delete Category** dialog opens.
- To confirm the operation, click **Delete**.

The category and all associated tags are deleted.

How Do You Create and Edit a vSphere Tag

Read how to create or edit a vSphere tag. After you create a tag, you can use it to add metadata to inventory objects.

The required privilege depends on the task that you want to perform.

Task	Privilege
Create a tag	vSphere Tagging > Create vSphere Tag on any vCenter Server inventory object.

Task	Privilege
Edit a tag	vSphere Tagging > Edit vSphere Tag on any vCenter Server inventory object.

You can use tags to add metadata to inventory objects. You can also record information about your inventory objects in tags, and you can use the tags in searches. For information about creating and editing a tag category, see [How Do You Create and Edit a vSphere Tag Category](#).

1. From the vSphere Client Home menu, click **Tags & Custom Attributes**.
2. Select the **Tags** tab and click **Tags**.
3. Start the task that you want to perform.

Task	Steps to perform
Create a tag	Click New .
Edit a tag	Select a tag and click Edit .

4. In the dialog box, enter or edit the name for the tag.
You can use up to 256 characters for the name of the tag.
5. Enter or edit the description of the tag.
6. If you create a tag, specify the tag category.
 - Select an existing category from the **Category** drop-down menu.
 - Click **Create New Category** and select the newly created category.
7. Click **Create** or **Save**.

How Do You Delete a vSphere Tag

You can delete an existing vSphere tag when you no longer need it by using the vSphere Client.

1. From the vSphere Client Home menu, click **Tags & Custom Attributes**.
2. Select the **Tags** tab and click **Tags**.
3. Select one or more tags from the list and click **Delete**.
4. To confirm the operation, click **Delete**.

The selected tags are deleted.

How Do You Assign or Remove a vSphere Tag

After you create vSphere tags, you can apply or remove them as metadata to objects in the vCenter Server inventory.

1. In the vSphere Client, navigate to an inventory tree.
2. Select an object from the list and click **Actions**.
3. From the drop-down menu, select **Tags and Custom Attributes**.
4. Select the required action.
 - To assign a tag, select **Assign Tag**.
 - To remove a tag, select **Remove tag**.

5. Select a tag from the list and click **Assign**.

In the **Assign Tag** dialog box, you can also create new tags with **Add tag**. For information about creating a tag, see [How Do You Create and Edit a vSphere Tag](#).

How Do You Add Permissions for vSphere Tags and Tag Categories

You can manage the user privileges for working with tags and categories. The procedure for assigning permissions to tags is the same as the procedure for assigning permissions to tag categories.

When you create a tag, you can specify which users and groups can operate with that tag. For example, you can grant administrative rights only to administrators and set read-only permissions for all other users or groups. You must have vSphere administrator credentials to set and manage permissions for tags.

Permissions for tags work similar to permissions for vCenter Server inventory objects. For more information, see the *vSphere Security* documentation.

1. From the vSphere Client Home menu, click **Tags & Custom Attributes**.
2. On the **Tags** tab, click the **Tags** or **Categories** button.
Depending on the button that you click, you see the list of tags or the list of tag categories.
3. Select an item from the list and click **Add Permission**.
4. In the **Add Permission** dialog box, select a domain from the drop-down menu.
5. Search for a user or a group to add.
6. Select a role to add from the drop-down menu.
7. To enable permission inheritance, select the **Propagate to children** check box.
8. Click **OK**.

vSphere Tagging Best Practices

Incorrect tagging can lead to replication errors. To avoid these errors, diligently follow best practices when tagging objects in your vSphere environment.

When working with tags in multiple node situations, expect replication delays between the nodes (generally 30 seconds to 2 minutes depending on your setup).

Follow these best practices to avoid replication errors:

- After creating a tag, if you immediately assign that tag to a local object, assign it from the management node where you created the tag.
- After creating a tag, if you immediately assign that tag to a remote object, assign it from the management node to which the object is local. Depending on your environment setup, allow for replication time to propagate the new tag before you use the tag.
- Avoid simultaneously creating categories and tags from different management nodes before categories and tags across nodes can finish the replication process. If duplicate categories or tags are created from different nodes at the same time, the duplicates might not be detected and will appear. If you see these results, manually delete duplicates from one management node.
- Before you remove an object from the vCenter Server inventory that is not a virtual machine or a host, first remove the tags from the object.

For more information about tagging performance best practices, see the blog post titled *vSphere 7.0 U1 Tagging Performance Best Practices* at <https://blogs.vmware.com/performance>.

What Are Custom Attributes in the vSphere Client

You can use custom attributes in the vSphere Client to assign user-specific values for each object of the custom attribute type.

After you create the attributes, set the value for the attribute on each virtual machine or managed host, as appropriate. This value is stored with vCenter Server and not with the virtual machine or managed host. Use the new attribute to filter information about your virtual machines and managed hosts. If you no longer need the custom attribute, remove it. A custom attribute is always a string.

For example, suppose that you have a set of products and you want to sort them by sales representative. Create a custom attribute for the sales person's name, Name. Add the custom attribute, Name, column to one of the list views. Add the appropriate name to each product entry. Click the column title Name to sort alphabetically.

The custom attributes feature is available only when you are connected to a vCenter Server system.

How Do You Add and Edit Custom Attributes in the vSphere Client

You can create custom attributes in the vSphere Client and associate these custom keys with an object, such as a host, virtual machine, cluster, or network. You can then assign a user-specific value using the custom attributes to specific objects in your vCenter Server environment.

After you create the attributes, set an appropriate value for the attribute on each inventory object. The attribute value is stored with vCenter Server and not with the inventory object. Use the new attribute to filter your inventory objects. If you no longer need the custom attribute, remove it. A custom attribute is always a string.

For example, suppose that you have lots of virtual machines in your inventory and you want to sort them by owner.

1. Create a `VM-owner` custom attribute for the virtual machine type of inventory objects.
2. Add the name of the owner to each VM entry you want and use the **Global Inventory Lists** pane to add the `VM-owner` custom attribute column to one of the list views.
3. You can now click the **VM-owner** column to sort alphabetically by owner.

Tip: If you cannot use the custom attribute column to sort the list of inventory objects, you can export the content of the list view and then use the sorting options on the saved CSV file.

Note: Tags and tag categories support a finer-grained mechanism for tagging your object. Consider using tags and tag categories instead of custom attributes.

1. From the vSphere Client Home menu, click **Tags & Custom Attributes**.
2. Click **Custom Attributes**.

All currently defined custom attributes for vCenter Server are displayed.

3. Click **New**.

Tags & Custom Attributes

Tags Custom Attributes

SC2

NEW EDIT DELETE

☐	Attribute
☐	
☐	
☐	
☐	
☐	
☐	

Manage Columns


4. Enter the values for the custom attribute.

- Enter the name of the attributes in the **Attribute** text box.
- Select the attribute type from the **Type** drop-down menu.
- Click **OK**.

After you have defined an attribute on an object, it is available to all objects of that type in the inventory. However, the value you specify is applied only to the currently selected object.

5. You can later add a value to the custom attribute for each inventory object that you want.

- a) Navigate to an inventory object in your vCenter Server inventory.
- b) On the **Summary** tab, navigate to the **Custom Attributes** card.
- c) Click **Edit**.



- If you want to add a value to an existing custom attribute, enter or edit the attribute value in the Value column.
- If you want to create a new custom attribute, enter a name and value to the attribute and select the attribute type from the **Type** drop-down menu. Click **Add**.

d) Click **Save**.

Working with vSphere Tasks

vSphere tasks are activities and actions that occur on an object within the vSphere inventory.

vSphere tasks represent system activities that do not complete immediately, such as migrating a virtual machine. For example, powering off a virtual machine is a task. You can perform this task manually every evening, or you can set up a scheduled task to power off the virtual machine every evening.

How Do You View vSphere Tasks

Learn how to view vSphere tasks that are associated with a single object or all objects in a vSphere environment.

By default, the task list for an object also includes tasks performed on its child objects. You can filter the task list by entering the object that you want in the **Target** column filter.

If you are logged in to a vCenter Server instance that is part of a vCenter Enhanced Linked Mode group, the **Server** column in the task list displays the name of the vCenter Server system on which the task is performed.

1. In the vSphere Client, navigate to an object in the inventory.
2. Click the **Monitor** tab.
3. Under **Tasks and Events**, select **Tasks**.
The task list contains tasks performed on the object and detailed information, such as target, task status, initiator, and start/completion time of the task.
You can limit the number of tasks visible in the **Tasks** pane by selecting an option from the **Tasks per page** drop-down menu.

4. Optional: Click the filter (⌵) icon on a column header and filter the task list by entering your filtering criteria.

5. Optional: Click the expand row (⌵) icon to view related events for a task. You can view task information about multiple tasks at the same time.

6. Optional: Click **Open in new tab** to open only the **Task Console** in a separate tab on your browser.

How Do You Cancel a vSphere Task

Canceling a vSphere task stops a running task from occurring. Canceling a scheduled task does not cancel subsequent runs. To cancel a scheduled task that has not run, reschedule it.

Required privileges:

- Manual tasks: **Tasks > Update task**
- Scheduled tasks: **Scheduled task > Remove task**

1. Locate the task in the **Recent Tasks** pane.
By default, the **Recent Tasks** pane is displayed at the bottom of the vSphere Client.

2. Click the cancel button (⊗) next to the progress bar in the Status column for the task you want.
If the cancel option is unavailable, the selected task cannot be canceled.

3. In the **Cancel Task** pop-up dialog box, click **OK**.

The vCenter Server system or ESXi host stops the progress of the task and returns the object to its previous state. The vSphere Client displays the task with a *The task was canceled by a user.* status.

Cancelled tasks and the corresponding events are retained in the vCenter Server database. For information about the retention of tasks and events in the vCenter Server database, and how to configure the database settings, see the *vSphere Monitoring and Performance* documentation.

How Do You Filter vSphere Tasks

To lower the number of vSphere tasks displayed for a single object or all objects in your environment, filter the task history in the database using several criteria, such as status, initiator, and time range.

1. In the vSphere Client, navigate to an object in the inventory.
2. Click the **Monitor** tab.
3. Under **Tasks and Events**, select **Tasks**.
4. Click **Filter**.

The **Filter Tasks** dialog opens.

5. Filter the tasks you want to view.
 - a) Select the respective check boxes in the **Status** options list to filter the tasks by their status.
 - b) From the **Type** drop-down menu, select the task type.
 - **User** - Displays the tasks initiated by all users, such as `vsphere.local` users, Active Directory domain users, and system users.
 - **System** - Displays the tasks initiated only by system users. System users are user accounts associated with an application or service.
 - c) To view the tasks initiated by all users or by specific users, select the **All users** or **Specific users** radio button.

If you select to view the tasks performed by specific users, create a list of users.

 1. Select the domain for the users from the drop-down menu.
 2. Select the users who initiated the tasks.
 - a. Enter a name in the Search box.

The system searches from the users in the selected domain.
 - b. Select a user and click **Add**.
 - d) From the **Select time set** drop-down menu, select the time when the task was started, completed or queued.
 - e) From the **Select time range** options, select the time range when the tasks were performed.
 - **Relative** - If you want to filter the tasks for a relative time range, specify the relative time range from the **Last** and **Range** drop-down menus.
 - **Date range** - If you want to filter the tasks for a specific time period, specify the date and time from the drop-down menus.

For example, to filter the tasks that got queued for the last 2 weeks, from the **Select time set** drop down menu, select **Queued time** and select the **Relative** radio button. Then, from the **Last** drop-down menu, select **2** and from the **Range** drop-down menu, select **Weeks**.

6. Click **Filter**.

The **Tasks** page displays the tasks that match your criteria. If you want to clear the task filter, click **Clear all filters** in the upper-right corner.

How Do You Export vSphere Tasks

You can export the detailed information for selected vSphere tasks or all tasks listed in the **Tasks** page to a CSV file.

To export the information about specific tasks, you can filter the task list and then select only the tasks that you want. You can also export the information about all tasks listed for an object or all objects in the vSphere environment.

1. In the vSphere Client, navigate to an object in the inventory.
2. Click the **Monitor** tab.
3. Under **Tasks and Events**, select **Tasks**.
4. Optional: To export one or more tasks, click the check boxes in front of the respective tasks.
5. Click **Export** at the top left corner of the **Tasks** pane.

Note: By default, if you did not select any tasks, the detailed information about all displayed tasks is exported.

- a) Optional: If you selected one or more tasks but you want to export the information about all displayed tasks, click **All Rows**.
- b) Optional: If you want to export only the information for the selected tasks, click **Selected Rows**.

The CSV file is generated and available for download.

How Do You Export vSphere Tasks with Advanced Export

You can use the advanced export option to export vSphere tasks that meet your specific criteria, such as status, initiator, and time range.

With the advanced export option for tasks, you can export up to 100 tasks that meet your specific criteria. You can use the exported information for monitoring and troubleshooting purposes. For example, you can export the tasks performed by specific users during the time range when some suspicious activity was detected.

1. In the vSphere Client, navigate to an object in the inventory.
2. Click the **Monitor** tab.
3. Under **Tasks and Events**, select **Tasks**.
4. From the **Export** drop-down menu, select **Advanced export**.

The **Advanced Export Tasks** dialog opens.

5. Filter the tasks for which you want to export information.
 - a) To export only tasks with specific status, select the respective check boxes in the **Status** options list.
 - b) To export tasks based on the initiator of the tasks, select one of the options from the **Type** drop-down menu.
 - **User** - Exports tasks initiated by all users, such as `vsphere.local` users, Active Directory domain users, and system users.
 - **System** - Exports tasks initiated only by system users. System users are user accounts associated with an application or service.
 - c) From the **Users** radio buttons, select whether to export the tasks initiated by specific users or by all users.

If you select the **Specify users** radio button, you must create a list of users.

1. From the **Domain** drop-down menu, select the domain for the users.
2. Create a list of users who initiated the tasks.
 - a. Enter a name in the Search text box.

The system searches from the users in the selected domain.

b. Select a user from the list and click **Add**.

d) From the **Select time set** drop-down menu, select a time set that defines the time range of the exported tasks.

You can select between start time, completion time, and queued time.

e) From the **Select time range** options, select the time range for the exported tasks.

- **Relative** - If you want to export the tasks for a relative time range, specify the relative time range from the **Last** and **Range** drop-down menus.
- **Data range** - If you want to export the tasks for a specific time period, specify the date and time from the drop-down menus.

For example, to export the tasks that got queued between March 13, 2022 10:30 AM and June 13, 2022 10:30 PM from the **Select time set** drop down menu, select **Queued time**. Then select the **Date range** radio button and specify the time period using the respective text fields.

f) From the **Include columns** radio buttons, select whether you want to export the information from all columns or from specific columns.

6. Click **Export**.

A .csv file is generated and available for download.

Scheduling vSphere Tasks

You can schedule vSphere tasks to run once in the future or multiple times, at a recurring interval.

The vSphere tasks that you can schedule are listed in the following table.

Table 145: Scheduled Tasks

Scheduled Task	Description
Add a host	Adds the host to the specified data center or cluster.
Change the power state of a virtual machine	Powers on, powers off, suspends, or resets the state of the virtual machine.
Change cluster power settings	Activate or deactivate DPM for hosts in a cluster.
Change resource settings of a resource pool or virtual machine	Changes the following resource settings: • CPU – Shares, Reservation, Limit. • Memory – Shares, Reservation, Limit.
Check compliance of a profile	Checks that a host's configuration matches the configuration specified in a host profile.
Clone a virtual machine	Makes a clone of the virtual machine and places it on the specified host or cluster.
Create a virtual machine	Creates a new virtual machine on the specified host.
Deploy a virtual machine	Creates a new virtual machine from a template on the specified host or cluster.
Migrate a virtual machine	Migrate a virtual machine to the specified host or datastore by using migration or migration with vMotion.
Make a snapshot of a virtual machine	Captures the entire state of the virtual machine at the time the snapshot is taken.
Scan for Updates	Scans templates, virtual machines, and hosts for available updates. This task is available only when is installed.
Remediate	Installs missing patches from the baselines selected for remediation on the hosts discovered during the scan operation and applies the newly configured settings. This task is available only when is installed.

You create scheduled tasks by using the **Scheduled Task** wizard. For some scheduled tasks, this wizard opens the wizard used specifically for that task. For example, if you create a scheduled task that migrates a virtual machine, the **Scheduled Task** wizard opens the **Migrate Virtual Machine** wizard, which you use to set up the migration details.

Scheduling one task to run on multiple objects is not possible. For example, you cannot create one scheduled task on a host that powers on all virtual machines on that host. You must create a separate scheduled task for each virtual machine.

After a scheduled task runs, you can reschedule it to run again at another time.

Note: At the time when a scheduled task is about to run, the user who created or last edited the task must be authenticated and authorized to run the task. If, for example, the user permissions to create or edit a scheduled task are removed, or the user account is deleted, the scheduled task fails with an error message. An administrator or a user with the required privileges can reschedule the task for subsequent runs.

How Do You Create a Scheduled vSphere Task

You can create scheduled tasks for operations that you want to run automatically once or at a recurring interval.

Required privilege: **Schedule Task > Create tasks**

Note: At the time when a scheduled task is about to run, the user who created the task must still have the **Schedule Task > Create tasks** privileges. If the user is no longer authenticated or authorized to run the task, for example, their permissions to create a scheduled task are removed, or the user account is deleted, the scheduled task fails with an error message. An administrator or another user with the required privileges can reschedule the task for subsequent runs.

If the task to schedule is not available in the vSphere Client, use the vSphere API. See the *vSphere SDK Programming Guide*.

CAUTION: Do not schedule multiple tasks simultaneously on the same object. The results are unpredictable.

1. In the vSphere Client, navigate to the object for which you want to schedule a task.
2. Select **Configure > Scheduled Tasks > New Scheduled Task**.
3. From the **New Scheduled Task** drop-down menu, select the task to schedule.
The **Schedule New Task** wizard opens for the task with the task name appended next to the name of the wizard. The wizard contains a page where you configure the scheduling options for the task.
4. Enter the required scheduling options for the task.
 - a) Enter a name and description of the task.
 - b) Choose a frequency for the task.

Option	Description
Once	Runs the scheduled task at the time selected.
After vCenter startup	Runs the task a specified number of minutes after vCenter Server startup.
Hourly	1. Enter the repeat frequency. 2. Enter the start date and time. 3. Enter the end date and time. For example, to start a task at the half-hour mark of every fifth hour, enter 5 hours and 30 minutes.
Daily	1. Enter the repeat frequency. 2. Enter the start date and time. 3. Enter the end date and time. For example, to run the task at 2:30 pm every four days, enter 4 and 14 : 30.
Weekly	1. Enter the repeat frequency.

Option	Description
	2. Select the day of the week. 3. Enter the start date and time. 4. Enter the end date and time. For example, to run the task at 6 am every Tuesday and Thursday, enter 1 week, 6 am, and select Tuesday and Thursday.
Monthly	1. Enter the repeat frequency. 2. Select the days by using one of the following methods. — Enter a specific day of the month and the number of months. For example, the tenth day every five months. — Select first, second, third, fourth, or last, and select the day of the week and the number of months. last runs the task on the last week in the month that the day occurs. For example, if you select the last Monday of the month and the month ends on a Sunday, the task runs six days before the end of the month.

c) Optional: Enter one or more email addresses to receive an email notification when the task is complete.

To schedule some tasks, you must enter additional details for the task. For example, to schedule taking a virtual machine snapshot, the **Schedule New Task (Take Snapshot)** wizard opens. In the **Scheduling Options** page, you set the scheduling options for the task, and in the **Snapshot Settings** page, you enter the properties for the snapshot.

5. Click **Schedule the task**.

How Do You Change or Reschedule a Scheduled vSphere Task

After a scheduled vSphere task is created, you can change the schedule, frequency, and other attributes of the task. You can edit and reschedule tasks before or after they run.

Required privilege: **Schedule Task > Modify**

Note: At the time when a scheduled task is about to run, the user who last edited the scheduled task must still have the **Schedule Task > Modify** privileges. If the user is no longer authenticated or authorized to run the task, for example, their permissions to edit a scheduled task are removed, or the user account is deleted, the scheduled task fails with a warning. An administrator or another user with the required privileges can reschedule the task for subsequent runs.

1. In the vSphere Client, navigate to the object for which you want to edit a scheduled task.
To view all scheduled tasks for a vCenter Server instance, navigate to that vCenter Server instance.
2. Select **Configure**, and select **Scheduled Tasks**.
3. Select a task from the list on the left and click **Edit**.
4. Right-click the task and select **Edit**.
5. Change the task attributes as necessary.
6. Click **Save**.

How Do You Remove a Scheduled vSphere Task

Removing a scheduled task removes all future occurrences of the task. The history associated with all completed occurrences of the task remains in the vCenter Server database.

Required privilege: **Scheduled task > Remove**

1. In the vSphere Client, navigate to the object for which you want to remove a scheduled task.
To view all scheduled tasks for a vCenter Server instance, navigate to that vCenter Server instance.

2. Select **Configure**, and select **Scheduled Tasks**.
3. Select a task from the list on the left and click **Remove**.

vSphere Host Configuration

Before you set up your virtual environment and consider how the virtual machines that it will support are going to be used and administered, you should configure ESXi hosts in vCenter Server.

The configuration of ESXi hosts involves several tasks. Learn how to configure the boot device of the ESXi host, the agent VM settings, advanced host attributes, and time configuration settings.

For information on configuring network and storage access, and security settings, see the configuration information for the specific vSphere component in the *vSphere Security* documentation, the *vSphere Storage* documentation, and the *vSphere Networking* documentation.

How to Configure the Boot Device on Your ESXi Host

On servers running ESXi, you can select the device that the server boots from.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select a host in the inventory and click the **Configure** tab.
3. Under Hardware, select **Overview** and click the **Boot Options** button.
4. Select a boot device from the drop-down menu.
5. Optional: To reboot immediately from the device you selected, select **Apply and Reboot on OK**.
If you do not select **Apply and Reboot on OK**, the new setting takes effect when the host reboots next.
6. Click **OK**.

How to Configure Your Agent VM Settings in the vSphere Client

In your vSphere environment, you can configure the datastore and network settings for the ESX agent virtual machines that you deploy on a host.

An ESX agent is a virtual machine, or a virtual machine and a vSphere Installation Bundle (VIB), that extend the functions of an ESXi host to provide additional services that a vSphere solution requires.

For example, a solution might require a particular network filter or firewall configuration to function. A solution can use an ESX agent to connect to the vSphere Hypervisor and extend the host with functions specific to that solution. For example, the ESX agent can filter network traffic, act as a firewall, or gather other information about the virtual machines on the host.

When you configure the datastore and network settings for ESX agents on a host, all of the ESX agents that you deploy on the host use that datastore and network configuration.

Important: ESX agents are deployed only if you configure the network and datastore settings.

1. Select a host in the vSphere Client inventory.
2. Click the **Configure** tab.
3. Under **Virtual Machines**, select **Agent VM Settings**.
The current settings for the ESX agents on the host, if any, appear.

4. Click **Edit**.
5. From the **Datastore** drop-down menu, select a datastore in which to deploy the ESX agent virtual machines.
6. From the **Network** drop-down menu, select a network to connect the ESX agents.
7. Click **OK**.

For information about ESX agents and ESX Agent Manager, see *Developing and Deploying vSphere Solutions, vServices, and ESX Agents*.

How Do You Set Advanced Host Attributes in the vSphere Client

Learn how to set advanced attributes for an ESXi host.

CAUTION: Changing advanced options is considered unsupported. Typically, the default settings produce the optimum result. Change the advanced options only when you get specific instructions from VMware technical support or a knowledge base article.

1. Browse to the host in the vSphere Client.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. Click the **Edit** button.
5. Find the appropriate item and change the value.
6. Click **OK**.

Synchronizing Clocks on the vSphere Network

Verify that all components on the vSphere network have their clocks synchronized. If the clocks on the physical machines in your vSphere network are not synchronized, SSL certificates and SAML tokens, which are time-sensitive, might not be recognized as valid in communications between network machines.

Unsynchronized clocks can result in authentication problems, which can cause the installation to fail or prevent the vCenter Server `vmware-vpxd` service from starting.

Time inconsistencies in vSphere can cause the first boot of a component in your environment to fail at different services depending on where in the environment time is not accurate and when the time is synchronized. Problems most commonly occur when the target ESXi host for the destination vCenter Server is not synchronized with NTP or PTP. Similarly, issues can arise if the destination vCenter Server migrates to an ESXi host set to a different time due to fully automated DRS.

To avoid time synchronization issues, ensure that the following is correct before installing, migrating, or upgrading a vCenter Server instance.

- The target ESXi host where the destination vCenter Server is to be deployed is synchronized to NTP or PTP.
- The ESXi host running the source vCenter Server is synchronized to NTP or PTP.
- When upgrading or migrating from vSphere 6.7 to vSphere 8.0, if the vCenter Server appliance is connected to an external Platform Services Controller, ensure the ESXi host running the external Platform Services Controller is synchronized to NTP or PTP.
- If you are upgrading or migrating from vSphere 6.7 to vSphere 8.0, verify that the source vCenter Server or vCenter Server appliance and external Platform Services Controller have the correct time.

Verify that any Windows host machine on which vCenter Server runs is synchronized with the Network Time Server (NTP) server.

To synchronize ESXi clocks with an NTP or a PTP server, you can use the VMware Host Client. For information about editing the time configuration of an ESXi host, see topic *Edit the Time Configuration of an ESXi Host in the VMware Host Client* in the *vSphere Single Host Management - VMware Host Client* documentation.

To learn how to change time synchronization settings for vCenter Server, see topic *Configure the System Time Zone and Time Synchronization Settings* in the *vCenter Server Configuration* documentation.

To learn how to edit the time configuration for a host by using the vSphere Client, see topic *Editing the Time Configuration Settings of a Host* in the *vCenter Server and Host Management* documentation.

Editing the Time Configuration Settings of Your ESXi Host

To ensure precise timestamping of events and synchronization of the time between an ESXi host and the other components in the vSphere network, configure the time settings of the ESXi host manually or synchronize the time and date of the host with an NTP or PTP server.

Exact timestamping shows the precise sequence of events that occur in the vSphere network. Time synchronization between the components of the vSphere network can prevent authentication problems, backup issues, and incorrect logs.

To ensure that all components in your environment have the correct time, as a best practice you must use the same time synchronization mechanism across the board.

Configure the Date and Time on Your ESXi Host Manually

Learn how to configure the date and time of the ESXi host manually in the vSphere Client. You can use the manual time configuration only after you stop synchronizing the time and date with NTP or PTP.

- Verify that the NTP and the PTP configurations are not used on the host.

If you notice some significant clock skew between a host in your environment and the rest of the vSphere components, you might need to first manually set the time and date on the host before synchronizing it to an NTP server or PTP.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select a host.
3. On the **Configure** tab, select **System > Time Configuration**.
4. Click **Manual Set-Up**.

The **Manual time configuration** dialog box appears.

5. Enter a date and time and click **OK**.

Note:

ESXi hosts use UTC (Coordinated Universal Time) and do not support changing time zones. In the vSphere Client, you see your local time as the current host time.

Use NTP Servers for Time and Date Synchronization of Your ESXi Host

To avoid time synchronization problems between an ESXi host and other components in the vSphere network, you can synchronize the time and date of the host to an NTP server.

- Verify that the PTP service is not running. For information about how to change its status, see *Manage the NTP and PTP Service Status on Your ESXi Host*.

The NTP and the PTP services cannot run simultaneously. Deactivate the PTP service and then enable the NTP server synchronization. Additionally, when you enable the NTP service, the manual time configuration becomes inactive.

Note:

You can set a startup policy to control the start and stop of the NTP service. You can also change the NTP service status manually. For more information about services, see *Manage the NTP and PTP Service Status on Your ESXi Host*.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select a host.
3. On the **Configure** tab, select **System > Time Configuration**.
4. Click **Add Service** and select **Network Time Protocol** from the drop-down menu.
5. In the **Network Time Protocol** dialog box, edit the network time protocol settings.
 - a) To monitor all events in the vSphere environment, select **Enable monitoring events**.
 - b) In the **NTP Servers** text box, enter the IP addresses or host names of the NTP servers that you want to use. A best practice is to use at least three NTP servers for the time synchronization of the host.
 - c) Click **OK**.

The NTP configuration is triggered immediately.

Use PTP for Time and Date Synchronization of Your ESXi Host

To ensure that the time of an ESXi host is synchronized with the time of other components of the vSphere network, you can synchronize the time and date of the host to PTP.

- Verify that the NTP service is not running. For information about how to change its status, see *Manage the NTP and PTP Service Status on Your ESXi Host*.
- For a hardware-based timestamping, verify that you enabled PCI passthrough for a network device on the host. See the *vSphere Networking* documentation.

The Precision Time Protocol allows both software- and hardware-based timestamping on an ESXi hosts and provides a highly accurate time synchronization. Starting with vSphere 7.0 Update 3, PTP provisions hardware-based timestamping for the virtual machines and the hosts within a network.

The PTP and the NTP services cannot run simultaneously. Stop the NTP service and then enable the PTP service. Additionally, when you enable the PTP service, the manual time configuration becomes inactive.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select a host.
3. On the **Configure** tab, select **System > Time Configuration**.
4. Click **Add Service** and select **Precision Time Protocol** from the drop-down menu.
5. In the **Precision Time Protocol** dialog box, edit the PTP settings.
 - a) Select the PTP-enabled network device for time synchronization from the **Network adapter type** drop-down menu.

PCI passthrough	For configuring a hardware timestamping with a PCI passthrough network device, follow these steps: 1. From the Device name drop-down menu, select the PCI passthrough network device. For vSphere 7.0 U3, the supported network adapters are IntelX710 and E810.
-----------------	--

	2. Enter the IP configuration settings for the network device. – To use DHCP to obtain the IP address configuration of the network device, select Obtain IPv4 settings automatically. – To set a static IP configuration, select the Use static IPv4 settings option and enter the IPv4 configuration settings. – To use the IEEE 802.3 network transport, select No IP configuration (use IEEE 802.3 transport).
VMkernel Adapter	For configuring software timestamping PTP, from the Device name drop-down menu, select the VMkernel adapter configured to use PTP.

6. Optional: Create a fallback mechanism in case the PTP synchronization fails.

- To monitor the events in vSphere, select the **Enable monitoring events** check box.
- To enable a fallback mechanism in case PTP become inaccessible, click **Enable fallback**.

You can create a fallback mechanism only if you enabled tracking the events in vSphere.

- In the **NTP Servers** text box, enter the IP addresses or host names of the NTP servers that you want to use. A best practice is to use at least three NTP servers for the time synchronization of the host.

7. Click **OK**.

Finishing the configuration immediately triggers the PTP synchronization.

8. Optional: To test whether the service used for time synchronization of the host is running correctly, click **Test Services**.

The **Time Synchronization Services test** dialog box appears. You can see information about how the activated time configuration works by collecting information from various systems.

Manage the NTP and PTP Service Status on Your ESXi Host

In the vSphere Client, you can configure a startup policy for the NTP or PTP service that runs on the ESXi host. You can also manually start, stop, or restart the respective service and in this way override the configured startup policy.

- In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
- Select a host.
- On the **Configure** tab, select **System > Services**.
- Change the status of the NTP or PTP service manually.

Option	Description
Option	Action
Change the NTP service status	1. Select NTP Daemon . 2. Click Start , Stop , or Restart .
Change the PTP service status	1. Select PTP Daemon . 2. Click Start , Stop , or Restart .

5. Optional: Select the NTP or PTP service and click **Edit Startup Policy**.

Option	Description
Start and stop with port usage	Starts or stops the service when a client port is activated or deactivated for access to the security profile of the host.
Start and stop with host	Starts and stops the service when the host powers on and shuts down.
Start and stop manually	You must manually control the status of the service.

Delete a Time Synchronization Service on Your ESXi Host

Read how to delete a time synchronization service that is no longer needed by using the vSphere Client.

Verify that the service you want to delete is not used to synchronize the time of the host.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select a host.
3. On the **Configure** tab, select **System > Time Configuration**.
4. Select the **Precision Time Protocol** or **Network Time Protocol** check box and click **Delete**.
5. Click **OK**.

Managing Hosts in vCenter Server

To monitor all hosts in your vSphere environment from one place and to simplify host configuration, connect the hosts to a vCenter Server system.

For information about configuration management of ESXi hosts, see the *vSphere Networking* documentation, the *vSphere Storage* documentation, and the *vSphere Security* documentation.

Disconnecting and Reconnecting an ESXi Host

You can disconnect and reconnect a host that a vCenter Server system manages. Disconnecting a managed host does not remove it from vCenter Server, but temporarily suspends all monitoring activities that vCenter Server performs.

The managed host and its associated virtual machines remain in the vCenter Server inventory. By contrast, removing a managed host from vCenter Server deletes the managed host and all its associated virtual machines from the vCenter Server inventory.

If an ESXi host becomes disconnected due to a loss of network connectivity to vCenter Server, the ESXi host reconnects automatically to vCenter Server after network connectivity is restored. If you manually disconnect an ESXi host, see [How to Reconnect Your Managed ESXi Host](#) for more information.

Reconnecting an ESXi host automatically or manually does not impact the running virtual machines, unless the host is part of a cluster and there are resource pool privileges configured. For more information about creating and managing resource pools, see the *vSphere Resource Management* documentation.

Tip: You can use the vSphere Automation API to automate the management of managed ESXi hosts. For more information about the available options, see the [vCenter Host APIs](#) reference documentation.

How to Disconnect Your Managed ESXi Host

Use the vSphere Client to disconnect a managed host from vCenter Server.

1. Navigate to **Home > Hosts and Clusters** and select a host.
2. Right-click the host and select **Connection > Disconnect** from the pop-up menu.
3. In the confirmation dialog box that appears, click **Ok**.

If the managed host is disconnected, the word “disconnected” is appended to the object name in parentheses, and the object is dimmed. All associated virtual machines are similarly dimmed and labeled.

How to Reconnect Your Managed ESXi Host

Use the vSphere Client to reconnect a managed host to a vCenter Server system.

1. Navigate to **Home > Hosts and Clusters** and select a host.
2. Right-click the host and select **Connection > Connect** from the pop-up menu.

When the managed host's connection status to vCenter Server is changed, the statuses of the virtual machines on that managed host are updated to reflect the change.

How Do You Reconnect ESXi Hosts After Changes to the vCenter Server SSL Certificate

vCenter Server uses an SSL certificate to encrypt and decrypt host passwords stored in the vCenter Server database. If the certificate is replaced or changed, vCenter Server cannot decrypt host passwords, and therefore cannot connect to managed hosts. If vCenter Server fails to decrypt a host password, the host is disconnected from vCenter Server.

You must reconnect the host and supply the login credentials, which are encrypted and stored in the database using the new certificate.

How to Relocate Your ESXi Host

You can move a host to another location within the vSphere inventory by dragging the host to the new location. The new location can be a folder, a cluster, or you can place the host as a standalone object in the data center.

Power off all virtual machines that are running on the host, or migrate the virtual machines to a new host by using vMotion.

When a host is moved from a cluster to another destination in the vSphere inventory, the resources which the host provides are deducted from the total cluster resources. You can either keep the virtual machines in the same cluster and migrate them to other hosts, or keep them on the host and remove them from the cluster. For information about removing a host from a cluster, see the *vSphere Resource Management* documentation.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters** and select a host.
2. If the host is part of a cluster, put it in maintenance mode.
 - a) Right-click the host and select **Maintenance Mode > Enter Maintenance Mode**.
 - b) Optional: If the host is part of a DRS cluster, evacuate the powered off or suspended virtual machines to other hosts within the cluster by selecting the check box **Move powered-off and suspended virtual machines to other hosts in the cluster**.
 - c) In the confirmation dialog box, click **OK**.
The host enters maintenance mode.
3. Select the host in the vSphere inventory panel and drag it to the new location within the inventory.
4. Right-click the host and select **Maintenance Mode > Exit Maintenance Mode**.
5. Optional: Power on the virtual machines that you powered off before you put the host in maintenance mode.

How to Remove an ESXi Host from Your vCenter Server Instance

Remove a managed host from vCenter Server to stop vCenter Server from monitoring and managing that host.

- Make sure that NFS mounts are active. If NFS mounts are unresponsive, the operation fails.

- If the host you want to remove from the cluster is connected to a distributed switch, remove the host from the switch. For more information, see *Remove Hosts from a vSphere Distributed Switch* in the *vSphere Networking* documentation.

If possible, remove managed hosts while they are connected. Removing a disconnected host does not remove the vCenter Server agent from the managed host.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select a host in the inventory.
3. Optional: If the host is part of a cluster, put it in maintenance mode.
 - a) Right-click the host and select **Maintenance Mode > Enter Maintenance Mode** from the pop-up menu.
If not all virtual machines on the host are powered off, the host does not enter maintenance mode.
If the host is part of a DRS cluster, when the host enters maintenance mode, DRS attempts to evacuate powered on virtual machines from the host by using vMotion.
 - b) In the confirmation dialog box, click **Ok**.
If the host is part of a DRS cluster, you can evacuate powered off or suspended virtual machines to other hosts within the cluster. Select the check box **Move powered-off and suspended virtual machines to other hosts in the cluster**.
The host icon changes and the term Maintenance Mode is added to the name in parentheses.
4. Right-click the host you want to remove in the inventory pane, and select **Remove from Inventory** from the pop-up menu.
5. In the confirmation dialog box, click **Yes** to remove the host.
vCenter Server removes the host and the associated virtual machines from the vCenter Server instance. vCenter Server then returns the status of all associated processor and migration licenses to available.

How to Reboot or Shut Down Your ESXi Host

You can power off or restart any ESXi host using the vSphere Client. Powering off a managed host disconnects it from vCenter Server, but does not remove it from the inventory.

- Power off all virtual machines that are running on the host, or migrate the virtual machines to a new host by using vMotion.
- Put the host in maintenance mode.

To shut down or reboot an ESXi host, you can use the ESXCLI command set. See the *Stopping and Rebooting Hosts with ESXCLI* topic in the *ESXCLI Concepts and Examples* documentation at <https://developer.broadcom.com/tools/esxcli/latest>.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters** and select a host.
2. From the **Actions** menu, select **Power**.
3. Select the operation.
 - To power off and restart the ESXi host, click **Reboot**.
 - To power off the ESXi host, click **Shut Down**.
4. Provide a reason for the operation and click **OK**.
This information is added to the log file.

How Do You Verify SSL Certificates for Legacy ESXi Hosts

You can configure vCenter Server to check the SSL certificates of hosts to which it connects. If you configure this setting, vCenter Server and the vSphere Client check for valid SSL certificates before connecting to a host for operations such as adding a host or making a remote console connection to a virtual machine.

Starting with vCenter Server 6.0, SSL certificates are signed by VMware Certificate Authority by default. You can instead use certificates from a third-party CA. Thumbprint mode is supported only for legacy hosts.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. Select **SSL settings**.
6. Determine the host thumbprint for each legacy host that requires validation.
 - a) Log in to the direct console.
 - b) Select **View Support Information** on the **System Customization** menu.
The thumbprint is displayed in the column on the right.
7. Compare the thumbprint you obtained from the host with the thumbprint listed in the vCenter Server SSL settings dialog box.
8. If the thumbprints match, select the check box for the host.
Hosts that are not selected will be disconnected after you click **Save**.
9. Click **Save**.

Enabling SGX Remote Attestation on Multi-Socket ESXi Hosts

Intel® Software Guard Extensions (Intel® SGX) is a hardware-based security solution that allows you to isolate specific application code and data in private memory regions, called enclaves. Use the vSphere Client to register your SGX-capable hosts with multiple CPU sockets with the Intel Registration Server and use remote attestation for applications running inside your vSGX-enabled virtual machines.

Starting with vSphere 7.0, you can enable the Virtual Intel® Software Guard Extensions (vSGX) on virtual machines and provide additional security to your workloads. See [Securing Virtual Machines with Intel Software Guard Extensions](#) in the *vSphere Virtual Machine Administration* documentation. Furthermore, you can use remote attestation for the vSGX-enabled virtual machines. Intel SGX remote attestation is a security mechanism that allows you to establish an authenticated and secure communication channel with a trusted remote entity. To use remote attestation for virtual machines using SGX enclaves, hosts with a single CPU socket do not require Intel registration.

Starting with vSphere 8.0, to enable the remote attestation on a virtual machine running on a host with multiple CPU sockets, you must first register the host with the Intel Registration Server. If an SGX-capable host with multiple CPU sockets is not registered with the Intel Registration Server, you can only power on vSGX-enabled virtual machines that do not require remote attestation.

When you add a host with SGX-capable CPUs, vCenter Server accesses the Unified Extensible Firmware Interface (UEFI) variables provided by the BIOS and reads the current registration status of the host. To enable vCenter Server to retrieve information about the SGX status of a host, you must set the firmware boot mode of the host to UEFI mode. See [How to View the SGX Registration Status of Your ESXi Host](#).

You can change the current SGX registration status of the host by using the registration options in vSphere Client or by rebooting the ESXi host after microcode updates and adding or replacing a CPU package. After each host reboot, you can view the updated registration status of the host by using the vSphere Client.

SGX Registration Statuses of a Host

You can view the current status of an SGX-capable hosts using the vSphere Client and perform the necessary steps to register the hosts with the Intel Registration Server.

SGX Registration Status	Description
Not Applicable	SGX-capable hosts with a single CPU socket do not require registration with the Intel Registration Server to enable remote attestation.
Incomplete	The registration status is incomplete in one of the following use cases: - When you add a new host to a vCenter Server instance and the host is not registered yet. - After a host firmware update that performs an Intel SGX Trusted Computing Base (TCB) recovery. - For hosts which have multiple CPU packages, when a CPU package is added or replaced, you must manually perform an SGX factory reset on BIOS setup. Then the host must be registered as if a newly added host. - When you manually perform an SGX factory reset on the BIOS setup, you must register the host again.
Complete	The host is successfully registered with the Intel Registration Server.

How to View the SGX Registration Status of Your ESXi Host

You can view the current SGX registration status of an ESXi host by using the vSphere Client.

- Make sure that the host is installed on an Intel CPU with SGX capabilities and SGX is enabled.
- Set the firmware boot mode of the host to UEFI.

- In the vSphere Client, navigate to an SGX-capable host.
- On the **Summary** tab, navigate to the **Hardware** card.
- Expand the **SGX** node to view the value of the Registration Status property.

For more information about the different registration statuses, see [SGX Registration Statuses of a Host](#).

To use the remote attestation feature for vSGX-enabled virtual machines, you must register the host with the Intel Registration Server if the host registration is incomplete and the host has multiple CPU sockets. See [How to Register Your Multi-Socket ESXi Host with the Intel SGX Registration Server](#).

How to Register Your Multi-Socket ESXi Host with the Intel SGX Registration Server

To use the SGX remote attestation feature for a multi-socket host, register the ESXi host with the Intel Registration Server using the vSphere Client.

- Make sure that the host is installed on an Intel CPU with SGX capabilities and SGX is enabled.

- Set the firmware boot mode of the host to UEFI.

The Intel SGX attestation mechanism ensures the trust between the vSGX enclave and an external entity. To use this feature on a multi-socket host with enabled SGX capabilities, you must register the host with the Intel SGX Registration Server.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. Select an SGX-capable host from the inventory and click the **Configure** tab.
3. Under Hardware, select **SGX** and click **Register**.

Upon successful completion of the registration operation, the registration status of the host changes to Completed.

Enable remote attestation for a vSGX-enabled virtual machine. See [Securing Virtual Machines with Intel Software Guard Extensions](#) in the *vSphere Virtual Machine Administration* documentation.

vSphere License Management

vSphere provides a centralized license management system that you can use to manage licenses for ESXi hosts, vCenter Server systems, vSAN clusters, Supervisors, and solutions.

If your vSphere environment has several vCenter Server systems that are joined in Enhanced Linked Mode, the licensing inventory is replicated across all linked vCenter Server systems. This way, the licensing data for each asset and all available licenses are replicated across all linked vCenter Server systems. Each individual vCenter Server system contains a copy of that data and licenses for all linked vCenter Server systems.

Note: Licensing data is replicated across multiple linked vCenter Server systems on a 10-minute interval.

For example, suppose that your environment consists of eight vCenter Server systems that are joined in Enhanced Linked Mode, and every vCenter Server system has 10 hosts connected to it. You can manage the licensing for all eight vCenter Server systems and the 80 hosts that are connected to them through the vSphere Client.

vSphere Licensing and Subscription Terminology and Definitions

The licensing system in vSphere uses specific terminology and definitions to refer to different licensing-related objects.

License Key	A license key encodes details about the product it is associated with, the license expiration date, the license capacity, and other information. The license key is assigned to an object to activate the functionality of its associated product.
License	A container for a license key of a VMware product. To use a license key, you create a license object in the vSphere Client and insert the license key into the license. After the license is created, you can assign it to assets.
Solution License	Starting with vSphere 8.0 Update 2b, you can use a Solution License and assign it to all component products of a solution. Example of a Solution License is the VMware vSphere Foundation Solution License.
Product Edition	A set of specific features that are associated with a unique license key. When assigned, the license key unlocks the features in the product edition. Examples of product editions are vSphere Standard, VMware vSphere Foundation, and so on.
Feature	Activated or deactivated functionality by a license that is associated with a specific product edition. Examples of features are vSphere DRS, vSphere vMotion, and vSphere High Availability.

Solution	A product that is packed and distributed independently from vSphere such as vCenter Server extensions, and client plug-ins. You install a solution in vSphere to take advantage of certain functionality. Every solution has a licensing model specific for the solution, but uses the License Service for license management and reporting.
Asset	Any object in vSphere that requires licensing. If the license has sufficient capacity, the license administrator in vSphere can assign one license to one or multiple assets of the same type. Suite licenses can be assigned to all assets that are part of the suite. Assets are vCenter Server systems, ESXi hosts, and products that integrate with vSphere.
License Use	The number of units that an asset uses from the capacity of a license. For example, if you assign a VMware vSphere Foundation Solution License to ESXi hosts, the license use is the number of the physical CPU cores for each CPU on all ESXi hosts licensed with the Solution License, with a minimum of 16 cores licensed per CPU.
License Capacity	The number of units that you can assign to assets.

VMware vSphere Foundation (VVF) Licensing

Starting with vSphere 8.0 Update 2b, you can use a Solution License to license all components of VMware vSphere Foundation.

VMware vSphere Foundation is an integrated solution that includes the following components:

- vCenter Server
- ESXi
- vSphere IaaS control plane
- vSAN Enterprise (250 GiB per core per host)
- VMware Aria Operations
- VMware Aria Operations for Logs
- VMware Aria Suite Lifecycle
- VMware Aria Automation Orchestrator (Basic)

The Solution License is one license key for all components of VMware vSphere Foundation. In the Broadcom Support Portal and the vSphere Client, the Solution License is entitled vSphere 8 Enterprise Plus for vSphere Foundation.

Note:

Individual component license keys continue to be supported. In the Broadcom Support Portal, they are provided together with the Solution License. You can use the Solution License, the individual component licenses, or a mix of both.

You add and assign the Solution License to vCenter Server instances, ESXi hosts, and Tanzu Supervisor Clusters.

After you assign the license to all ESXi hosts in a vSAN cluster, and if the vSAN cluster has no vSAN license assigned at cluster level, the cluster is licensed automatically. With the vSphere Foundation Solution License, you receive capacity of 100 GiB per core per ESXi host. For more information, see the *vSAN Planning and Deployment* guide.

After you assign the license to vCenter Server, Aria Suite components that are registered with this vCenter Server are licensed automatically. For more information about managing and configuring licenses in the vSphere Client, see [Managing vSphere Licenses](#).

Note: The Solution License can be assigned to components, or applied automatically in the case of Aria components and vSAN, if they meet the minimum version requirements.

Component	Version
ESXi	8.0 Update 2b and later
vCenter Server	
vSphere IaaS control plane	
vSAN	
VMware Aria Operations	8.16 and later
VMware Aria Operations for Logs	
VMware Aria Suite Lifecycle	

If you cannot upgrade to the minimum required versions, you can downgrade the Solution License key in the Broadcom Support Portal.

You can assign one Solution License to multiple vCenter Server instances. vCenter Server does not consume Solution License capacity. The Solution License capacity is based on the per core licensing model. For more information about the per core licensing model for vSphere, see [Licensing and Subscription in vSphere](#).

In the vSphere Client, you can also view the information about the Solution Licenses that you add. For example, you can view the components licensed with the Solution License, the licensed features for each components, and the license capacity used by the ESXi hosts and vSAN. For more information about viewing licensing information in the vSphere Client, see [Viewing and Exporting vSphere Licensing Information](#).

Evaluation Mode

After you install a component, it operates in evaluation mode for up to 60 days. Evaluation mode provides all features of the components.

Licensing and Subscription in vSphere

To apply the licensing models of vSphere assets correctly, you must understand how the type of asset, for example ESXi hosts or vCenter Server instances, consume the license capacity.

If you have other VMware products with assigned licenses in your vCenter Server, you can view and manage their licenses in the vSphere Client.

Starting with vSphere 8.0 Update 2b, you can use a Solution License and assign it to all components of a solution. For example, you can license Supervisor Clusters and vSAN cluster with the VMware vSphere Foundation Solution License. For more information, see [VMware vSphere Foundation \(VVF\) Licensing](#).

If you use *VMware vSphere+*, you leverage several cloud capabilities while you manage your on-premises vSphere infrastructure. For more information, see the *VMware vSphere+* documentation.

Licensing for ESXi Hosts

ESXi hosts are licensed with vSphere licenses. Each vSphere license has a certain capacity that you can use to license multiple physical CPUs on ESXi hosts.

There are four main licensing models for vSphere:

- Per core licensing with a minimum of 16 cores licensed per CPU.
- Per virtual machine licensing.
- vSphere+ subscription capacity based licensing.
- Per CPU licensing that covers one CPU with up to 32 cores.

To license an ESXi host, you must assign to it a vSphere license that meets the following prerequisites:

- The license must have sufficient capacity depending on the licensing model.
- The license must support all the features that the host uses. For example, if the host is associated with a vSphere Distributed Switch, the license that you assign must support the vSphere Distributed Switch feature.

If you attempt to assign a license that has insufficient capacity or does not support the features that the host uses, the license assignment fails.

Per Core Licensing Model for vSphere

The per core licensing model is subscription based. To calculate the capacity you need for your environment, you need the total number of the physical CPU cores for each CPU on all ESXi hosts in your environment. Each core requires a single license, and the minimum license capacity you can purchase is 16 cores per CPU.

For example, if you have 1 ESXi host in your inventory, with 1 CPU, and 8 CPU cores per CPU, you must purchase the subscription capacity of 16 cores per CPU because it is the minimum license capacity.

Number of ESXi Hosts	Number of CPUs	Cores per CPU	Number of Core Licenses
1	1	8	16
2	2	8	64
2	2	16	64

For more information about calculating the number of licenses you need for your environment, see the VMware knowledge base article at <https://kb.vmware.com/s/article/95927>.

Per Virtual Machine Licensing Model for vSphere

The license use equals the total number of powered on desktop virtual machines running on the hosts that are assigned such a license.

Subscription Based Licensing Model for vSphere

If you currently use the vSphere+ workload platform, you have a subscription capacity based model for management of vSphere. For more information about your subscription model, see the *VMware vSphere +* documentation.

Per CPU Licensing Model for vSphere

If you use the per CPU licensing model, one PU license covers one CPU with up to 32 cores. If a CPU has more than 32 cores, you need additional CPU licenses.

Number of CPUs	Cores per CPU	Number of CPU Licenses
1	1-32	1
2	1-32	2
1	33-64	2
2	33-64	4

When you assign a vSphere license to a host, the amount of capacity consumed is determined by the number of physical CPUs on the host and the number of physical cores in each physical CPU.

If you use the licensing model with up to 32 cores, you can assign a vSphere license for 10 32-core CPUs to any of the following combinations of hosts:

- Five 2-CPU hosts with 32 cores per CPU

- Five 1-CPU hosts with 64 cores per CPU
- Two 2-CPU hosts with 48 cores per CPU and two 1-CPU hosts with 20 cores per CPU

Dual-core and quad-core CPUs, such as Intel CPUs that combine two or four independent CPUs on a single chip, count as one CPU.

Evaluation Mode for ESXi Hosts

After you install ESXi, it operates in evaluation mode for up to 60 consecutive days. Evaluation mode provides all features of vSphere.

After you assign a license to an ESXi host, at any time before the evaluation period expires, you can set the host back to evaluation mode to explore the entire set of features available for the remaining evaluation period.

For example, if you use an ESXi host in evaluation mode for 20 days, then assign a vSphere Standard license to the host, and 5 days later set the host back to evaluation mode, you can explore the entire set of features available for the host for the remaining 35 days of the evaluation period.

License and Evaluation Period Expiry for ESXi Hosts

For ESXi hosts, license or evaluation period expiry leads to disconnection from vCenter Server. All powered on virtual machines continue to work, but you cannot power on virtual machines after they are powered off. You cannot change the current configuration of the features that are in use. You cannot use the features that remained unused before the license expiration.

Note:

When there are expiring licenses, a notification appears 90 days before the license expiration.

Licensing ESXi Hosts After an Upgrade

If you upgrade an ESXi host to a version that starts with the same number, you do not need to replace the existing license with a new one. For example, if you upgrade a host from ESXi 8.0 to 8.1, you can use the same license for the host.

If you upgrade an ESXi host to a major version that starts with a different number, the evaluation period restarts and you must assign a new license. For example, if you upgrade an ESXi host from 7.x to 8.x, you must license the host with a vSphere 8 license.

Licensing for vCenter Server

You can license a vCenter Server systems with a vCenter Server license that has a per-instance capacity or with a Solution License.

Evaluation Mode for vCenter Server

When you install a vCenter Server system, it is in evaluation mode. The evaluation mode of a vCenter Server system expires 60 days after the product is installed no matter whether you assign a license to vCenter Server or not. You can set vCenter Server back to evaluation mode only within 60 days after its installation.

For example, suppose that you install a vCenter Server system and use it in evaluation mode for 20 days and assign the system an appropriate license. The vCenter Server evaluation mode expires after the remaining 40 days of the evaluation period.

License and Evaluation Period Expiry of vCenter Server

When the license or evaluation period of a vCenter Server system expires, all hosts disconnect from that vCenter Server system. The virtual machines running on the disconnected hosts remain intact. Until you assign a new license to the

vCenter Server system, you can manage the virtual machines on the disconnected hosts by logging into each host separately. After you assign a new license key to vCenter Server, all disconnected hosts reconnect to the vCenter Server system.

Note: When there are expiring licenses, a notification appears 90 days before the license expiration.

Licensing vCenter Server After an Upgrade

If you upgrade vCenter Server to a version that starts with the same number, you can keep the same license. For example, if you upgrade a vCenter Server system from vCenter Server 8.0 to 8.1, you can keep the same license on the system.

If you upgrade vCenter Server to a major version that starts with a different number, the evaluation period restarts and you must assign a new license. For example, if you upgrade a vCenter Server system from 7.x to 8.x, you must license the system with a vCenter Server 8 license.

If you upgrade the edition of the license, for example, from vCenter Server Foundation to vCenter Server Standard, replace the existing license on the system with the upgraded license.

Subscription Based Licensing for vCenter Server

If you currently use the vSphere+ workload platform, you have a subscription capacity based model for management of vSphere. For more information about your subscription model, see the *VMware vSphere+* documentation.

Managing vSphere Licenses

To license an asset in vSphere, you must assign to it a license that holds an appropriate product license key. You can use the license management functionality in the vSphere Client to license multiple assets at a time from a central place. Assets are vCenter Server systems, hosts, vSAN clusters, Supervisors, and solutions.

In vSphere, you can assign one license to multiple assets of the same type if the license has enough capacity.

How Do You Create a New License in the vSphere Client

When you purchase, split, or merge license keys in the Broadcom Support Portal, you must use the new keys to license assets in your vSphere environment. You must go to the vSphere Client and create a license object for every license key. A license is a container for a license key of a VMware product.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.

1. Click **Menu > Administration**.
2. Expand **Licensing** and click **Licenses**.
3. On the **Licenses** tab, click **ADD**.
4. On the **Enter licenses keys** page, enter one license key per line, and click **Next**.

The license key is a 25-symbol string of letters and digits in the format `XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX`. You can enter a list of keys in one operation. A new license will be created for every license key that you enter.

5. On the **Edit license names** page, rename the new licenses as appropriate and click **Next**.
6. On the **Ready to complete** page, review the new licenses and click **Finish**.

A new license is created for every license key that you entered.

After you create the new licenses, you can assign them to assets. Assign the new licenses to hosts, vCenter Server systems, or other products that you use with vSphere. You must not keep unassigned licenses in the inventory.

Configuring License Settings for Assets in the vSphere Client

To continue using product functionality, you must assign appropriate licenses to assets in evaluation mode, or assets with expiring licenses. When you upgrade a license edition, merge, or split licenses in the Broadcom Support Portal, you must assign the new licenses to assets. You can assign licenses that are already available or create licenses and assign them to the assets in a single workflow. Assets are vCenter Server systems, ESXi hosts, vSAN clusters, Supervisors, and other products that integrate with vSphere.

How to Assign a License to Multiple Assets in Your vSphere Environment

Learn how to assign licenses that are already available, or create licenses and assign them to vSphere assets in a single workflow.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.

To continue using product functionality, you must assign appropriate licenses to assets in evaluation mode, or assets with expiring licenses. When you upgrade a license edition, merge, or split licenses in the Broadcom Support Portal, you must assign the new licenses to assets. Assets are vCenter Server systems, ESXi hosts, vSAN clusters, Supervisors, and other products that integrate with vSphere.

- Click **Menu > Administration**.
- Expand **Licensing** and click **Licenses**.
- Select the **Assets** tab.
- On the **Assets** tab, click the **vCenter Server Systems, Hosts, vSAN Clusters, Supervisors, or Solutions** tab.
- Select the assets to license.
- Click **Assign License**.
- In the **Assign License** dialog box, select the task that you want to perform.
 - In the vSphere Client, select an existing license or select a newly created license.

Task	Steps
Select an existing license	Select an existing license from the list and click OK .
Select a newly created license	- Click the New License tab. - In the Assign License dialog box, type or copy and paste a license key. - Enter a name for the new license and click OK. Details about the product, product features, capacity, and expiration period appear on the page. - Click OK. - In the Assign License dialog box, select the newly created license, and click OK.

The license is assigned to the assets. Capacity from the license is allocated according to the license use of the assets. For example, if you assign the license to 3 hosts with 4 CPUs each, the consumed license capacity is 12 CPUs.

How to Configure License Settings for Your ESXi Host

You must assign a license to an ESXi host before its evaluation period expires or its currently assigned license expires. If you upgrade, merge, or split vSphere licenses in the Broadcom Support Portal, you must assign the new licenses to ESXi hosts and remove the old licenses.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
- Navigate to the host in the inventory.
 - Select the **Configure** tab.
 - Under **System**, select **Licensing**.
 - Click **Assign License**.
 - In the **Assign License** dialog box, select the task that you want to perform.
 - In the vSphere Client, select an existing license or select a newly created license.

Task	Steps
Select an existing license	Select an existing license from the list and click OK .
Select a newly created license	- Click the New License tab. - In the Assign License dialog box, type or copy and paste a license key. - Enter a name for the new license and click OK. Details about the product, product features, capacity, and expiration period appear on the page. - Click OK. - In the Assign License dialog box, select the newly created license, and click OK.

The license is assigned to the host. Capacity from the license is allocated according to the license use of the host.

How to Configure License Settings for Your vCenter Server

You must assign a license to a vCenter Server system before its evaluation period expires or its currently assigned license expires. If you upgrade, combine, or divide vCenter Server licenses on the Broadcom support portal, you must assign the new licenses to vCenter Server systems and remove the old licenses.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
- In the vSphere Client, navigate to the vCenter Server instance.
 - Select the **Configure** tab.
 - Under **Settings**, select **Licensing**.
 - Click **Assign License**.
 - In the **Assign License** dialog box, select the task that you want to perform.
 - In the vSphere Client, select an existing license or select a newly created license.

Task	Steps
Select an existing license	Select an existing license from the list and click OK .
Select a newly created license	- Click the New License tab. - In the Assign License dialog box, type or copy and paste a license key. - Enter a name for the new license and click OK. Details about the product, product features, capacity, and expiration period appear on the page. - Click OK.

Task	Steps
	5. In the Assign License dialog box, select the newly created license, and click OK .

The license is assigned to the vCenter Server system, and one instance from the license capacity is allocated for the vCenter Server system.

How to Set Your vSphere Assets to Evaluation Mode

To explore the complete set of features available for an asset, you can set it to evaluation mode.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.

Different products have different terms for using their evaluation mode. Before you set an asset to evaluation mode, you should consider the specifics for using the evaluation mode of its associated product. For details, see the licensing model documentation for the relevant product at [Licensing and Subscription in vSphere](#).

- Click **Menu > Administration**.
- Expand **Licensing** and click **Licenses**.
- Select the **Assets** tab.
- Select the **vCenter Server systems, Hosts, vSAN Clusters, Supervisors**, or **Solutions** tab.
- Select the asset that you want to set to evaluation mode.
- Click **Assign License**.
- Select the evaluation license, and click **OK** to save your changes.

The asset is in evaluation mode. You can explore the entire set of features that are available for the asset.

Note: You must assign an appropriate license to the asset before its evaluation period expires. Otherwise the asset gets into unlicensed state and some functionalities are blocked.

How Do You Rename a License in the vSphere Client

After you create a license, you can change its name by using the vSphere Client.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
- Click **Menu > Administration**.
 - Expand **Licensing** and click **Licenses**.
 - Select the **Licenses** tab.
 - Select the license to rename, and click **Rename License**.
 - Enter the new license name, and click **OK**.

How Do You Remove a License in the vSphere Client

To remain in compliance with the licensing models of products that you use with vSphere, you must remove all unassigned licenses from the inventory. If you have split, merged, or upgraded licenses in the Broadcom Support Portal, you must remove the old licenses.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
- Verify that the license is unassigned from the assets. To unassign the license from assets, navigate to **Administration > Pre-Version 9 Licenses > Assets**.

For example, suppose that you have upgraded a vSphere license from 7.0 to 8.0 in the Broadcom Support Portal. You assign the license to ESXi 8.0 hosts. After assigning the new vSphere 8.0 licenses, you must remove the old vSphere 7.0 license from the inventory.

1. Click **Menu > Administration**.
2. Select the **Licenses** tab.
3. Use the filters to display only the unassigned licenses.
If the license is assigned to an asset, you cannot remove it. You must first unassign it from all assets.
 - a) Click the filter icon (⌵) in the **State** column.
A text box appears.
 - b) Enter the state of the licenses that you want to display, unassigned or assigned.
If you want to remove the filter, delete your entry in the text box.
4. From the list, select a license to remove or press Ctrl+A to select all licenses.
5. Click **Remove Licenses**, review the confirmation message, and click **Yes**.

Viewing and Exporting vSphere Licensing Information

You can view the licensing state of the vSphere environment from a central place by using the license management functionality in the vSphere Client. You can view the licenses that are available in vSphere, current license assignments and usage, available license capacity, licensed features in use, and so on. You can export licensing information about vSphere licenses, products, or assets in a CSV file.

You can also export information about licenses and their expiration dates, capacity, and usage. You can export data about the available products and assets in the vSphere Client by downloading a .CSV file.

How to View Licensing Information About Your vSphere Environment

You can view the available licenses in vSphere and their expiration dates, available capacity, and usage. You can also view the available products and assets.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
1. Click **Menu > Administration**.
 2. Expand **Licensing** and click **Licenses**.
 3. Select a tab for the licensing information that you want to view.

Tab	Description
Licenses	Lists all licenses that are available in the vSphere environment. For every license, you can view the associated license key, license usage, license capacity, and expiration date.
Products	Lists the products that have licenses available in the vSphere environment. You can view the licenses that are available for every product, licensed features, license usage, and license capacity.

Tab	Description
Assets	Displays licensing information about the assets that are available in the vSphere environment. Assets are vCenter Server systems, hosts, vSAN clusters, Supervisors, and other products that you use with vSphere that are listed under Solutions.

If you have upgraded, split, or merged any licenses in the Broadcom Support Portal, you must not use these old license keys and should remove them from the inventory.

How to View Available Licenses and Features About Your vSphere Product

By using the vSphere Client, you can view information about a product, such as the available licenses, features, and license capacity.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
- Click **Menu > Administration**.
 - Expand **Licensing** and click **Licenses**.
 - Click the **Products** tab and select the product for which you want to view information.
 - Select the task that you want to perform.

Task	Description
View the licenses that are available for the selected product	Click the Licenses subtab below the list of products.
View the licensed features for the product	Click the Features subtab below the list of products.

How to View the vSphere Features That an Asset Can Use

Learn how to view the features that an asset in your vSphere environment can use based on the license that is assigned to the asset.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
- Click **Menu > Administration**.
 - Expand **Licensing** and click **Licenses**.
 - Select the **Assets** tab.
 - Select **vCenter Server systems**, **Hosts**, **vSAN Clusters**, **Supervisors**, or **Solutions**.
 - Select an asset from the list and click the **Open** button to the left of the asset name.
 - Click the **Features** subtab.

How to View the License Key of Your License

In vSphere, a license holds a license key for a product. Learn how to view the associated license key for every license.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
1. Click **Menu > Administration**.
 2. Expand **Licensing** and click **Licenses**.
 3. Select the **Licenses** tab.
 4. Select a license from the list, and click the **Open** button to the left of the asset name.
 5. Click the **Summary** subtab. Under General, you can see the license name, the license expiration date, the state of the license, and the license key.

How to View the Licensed Features for an Asset in Your vSphere Environment

Before you start to use a feature on an asset, you can check whether the asset is licensed to use this feature. For example, to use vSphere HA, you must check whether all hosts in a vSphere HA cluster are licensed for this feature.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
1. Navigate to the vCenter Server system, host, or cluster whose licensed features you want to view.
 2. Click the **Configure** tab.
 3. Under Settings, select **Licensing**.

The list of features that you can configure on the asset appears on the right.

How to Export Licensing Information for Your vSphere Environment

You can export licensing information about vSphere licenses, products, or assets. The information is saved on your local system as a **.CSV** file. You can later open the **.CSV** file with third-party applications.

- To view and manage licenses in the vSphere environment, you must have the **Global > Licenses** privilege on the system, where the vSphere Client runs.
1. Click **Menu > Administration**.
 2. Expand **Licensing** and click **Licenses**.
 3. Optional: Select an item to export.
If you do not select a particular license, product, or asset, all items from the respective list are exported.
 - Click the **Licenses** tab to select one or more licenses to export.
 - Click the **Products** tab to select a product to export.
 - Click the **Assets** tab to select the vCenter Server instance, hosts, clusters, or solutions whose licensing information you want to export.
 4. Click the **Export** button at the bottom of the list.
When prompted, save the CSV file in your file system.
 - If you selected one or more items to export, click **Selected Rows**.
 - If you selected one or more items but you want to export all items, click **All Rows**.
 - If you selected no items, all items from the respective list are exported.

vCenter Server Domain Repoint License Considerations

Domain repointing copies license keys to a new domain. Copying the license keys ensures that valid licensing of all assets is maintained after repointing.

vCenter Server tracks license usage on a per domain basis. If a key is used in more than one domain, you must ensure that the aggregate use of the key does not exceed its capacity. To simplify your license management, remove each license copied to a second domain and assign a new license to assets.

Consider the following two cases:

- License keys that are no longer in use (that is, assigned to assets) in the original domain post repointing.
- License keys that are in use (that is, assigned to assets) in multiple domains.

License Keys Not in Use in a Domain

License Keys in Use in Multiple Domains

If after completing repointing, a license key is in use (that is, assigned to assets) in more than one domain, to remove the license key from all but one domain, first a different license key must be assigned to each asset in domains from which the license key will be removed. Two common approaches:

vSphere Virtual Machines Migration

With vSphere vMotion, you can move virtual machines from one compute resource or storage location to another by using cold or hot migration.

For example, with vSphere vMotion you can move powered on virtual machines away from a host to perform maintenance, to balance loads, to colocate virtual machines that communicate with each other, to move virtual machines apart to minimize fault domain, to migrate to new server hardware, and so on.

Moving a virtual machine from one inventory folder to another folder or resource pool in the same data center is not a form of migration. Unlike migration, cloning a virtual machine or copying its virtual disks and configuration file in the same vCenter Server system are procedures that create a new virtual machine. Cloning and copying a virtual machine in the same vCenter Server system are also not forms of migration.

By using migration, you can change the compute resource that the virtual machine runs on. For example, you can move a virtual machine from one host to another host or cluster.

To migrate virtual machines with disks larger than 2 TB, the source and destination ESXi hosts must be version 6.0 and later.

Depending on the power state of the virtual machine that you migrate, migration can be cold or hot.

Cold Migration

Moving a powered off or suspended virtual machine to a new host. Optionally, you can relocate configuration and disk files for powered off or suspended virtual machines to new storage locations. You can also use cold migration to move virtual machines from one virtual switch to another, and from one data center to another. You can perform cold migration manually or you can schedule a task.

Hot Migration

Moving a powered on virtual machine to a new host. Optionally, you can also move the virtual machine disks or folder to a different datastore. Hot migration is also called live migration or vSphere vMotion. With vSphere vMotion, you migrate the virtual machine without any interruption in its availability.

Depending on the virtual machine resource type, you can perform three types of migration.

Change compute resource only

Moving a virtual machine, but not its storage, to another compute resource, such as a host, cluster, resource pool, or vApp. You can move the virtual machine to another compute resource by using

Change storage only

cold or hot migration. If you change the compute resource of a powered on virtual machine, you use vSphere vMotion.

Moving a virtual machine and its storage, including virtual disks, configuration files, or a combination of these, to a new datastore on the same host. You can change the datastore of a virtual machine by using cold or hot migration. If you move a powered on virtual machine and its storage to a new datastore, you use Storage vMotion.

Change both compute resource and storage

Moving a virtual machine to another host and at the same time moving its disk or virtual machine folder to another datastore. You can change the host and datastore simultaneously by using cold or hot migration.

In vSphere 6.0 and later, you can move virtual machines between vSphere sites by using migration between the following types of objects.

Migrate to another virtual switch

Moving the network of a virtual machine to a virtual switch of a different type. You can migrate virtual machines without reconfiguring the physical and virtual network. By using cold or hot migration, you can move the virtual machine from a standard to a standard or distributed switch, and from a distributed switch to another distributed switch. When you move a virtual machine network between distributed switches, the network configuration and policies that are associated with the network adapters of the virtual machine are transferred to the target switch.

Migrate to another data center

Moving a virtual machine to a different data center. You can change the data center of a virtual machine by using cold or hot migration. For networking in the target data center, you can select a dedicated port group on a distributed switch.

Migrate to another vCenter Server system

Moving a virtual machine to a different vCenter Server instance.

You can move a virtual machine to a vCenter Server instance that is connected to the source vCenter Server instance through vCenter Enhanced Linked Mode.

You can also move virtual machines between vCenter Server instances that are located across a long distance from each other.

Starting with vSphere 7.0 Update 1c, you can migrate workloads across vCenter Server systems by using Advanced Cross vCenter vMotion. You can initiate migration of workloads both from on-premise environments and from cloud environments. Advanced Cross vCenter vMotion does not depend on vCenter Enhanced Linked Mode or Hybrid Linked Mode and you can migrate virtual machines across vCenter Server systems in different vCenter Single Sign-On domains.

Starting with vSphere 7.0 Update 3, you can use the Advanced Cross vCenter vMotion feature to clone virtual machines across vCenter Server systems.

Cold Migration in Your vSphere Environment

Cold migration is the migration of powered off or suspended virtual machines between hosts across clusters, data centers, and vCenter Server instances. By using cold migration, you can also move associated disks from one datastore to another.

You can use cold migration to have the target host checked against fewer requirements than when you use vMotion. For example, if you use cold migration when a virtual machine contains a complex application setup, the compatibility checks during vMotion might prevent the virtual machine from moving to another host.

You must power off or suspend the virtual machines before you begin the cold migration process. Migrating a suspended virtual machine is considered a cold migration because although the virtual machine is powered on, it is not running.

Migrating a virtual machine across its network subnets may result in the loss of network connectivity.

CPU Compatibility Check During Cold Migration

If you attempt to migrate a powered off virtual machine that is configured with a 64-bit operating system to a host that does not support 64-bit operating systems, vCenter Server generates a warning. Otherwise, CPU compatibility checks do not apply when you migrate powered off virtual machines with cold migration.

When you migrate a suspended virtual machine, the new host for the virtual machine must meet CPU compatibility requirements. This requirement allows the virtual machine to resume execution on the new host.

Operations During Cold Migration

A cold migration consists of the following operations:

1. If you select the option to move to a different datastore, the configuration files, including the NVRAM file (BIOS settings), log files, and the suspend file, are moved from the source host to the destination host's associated storage area. You can choose to move the virtual machine's disks as well.
2. The virtual machine is registered with the new host.
3. After the migration is completed, the old version of the virtual machine is deleted from the source host and datastore if you selected the option to move to a different datastore.

Network Traffic for Cold Migration

By default, data for VM cold migration, cloning, and snapshots is transferred through the management network. This traffic is called provisioning traffic. It is not encrypted but uses run-length encoding of data.

On a host, you can dedicate a separate VMkernel network adapter to the provisioning traffic, for example, to isolate this traffic on another VLAN. On a host, you can assign no more than one VMkernel adapter for provisioning traffic. For information about enabling provisioning traffic on a separate VMkernel adapter, see the *vSphere Networking* documentation.

If you plan to transfer high volumes of virtual machine data that the management network cannot accommodate, redirect the cold migration traffic on a host to the TCP/IP stack that is dedicated to cold migration and cloning of powered off virtual machines. You can also redirect if you want to isolate cold migration traffic in a subnet different from the management network, for example, for migration over a long distance. See [How to Place Traffic for Cold Migration, Cloning, and Snapshots on the Provisioning TCP/IP Stack](#).

How to Migrate a Powered Off or Suspended Virtual Machine

You can use cold migration to move a virtual machine and its associated disks from one datastore to another in your vSphere environment. The virtual machines are not required to be on shared storage.

-

- Required privilege: **Resource > Migrate powered off virtual machine**

- Power off or suspend the virtual machine.
- Right-click the virtual machine and select **Migrate**.
 - To locate a virtual machine, select a data center, folder, cluster, resource pool, host, or vApp.
 - Click the **Virtual Machines** tab.
- Select the migration type and click **Next**.

Option	Description
Change compute resource only	Move the virtual machine to another host.
Change storage only	Move the virtual machine's configuration file and virtual disks.
Change both compute resource and storage	Move the virtual machine to another host and move its configuration file and virtual disks.

- If you change the compute resource of the virtual machine, select the destination compute resource for this virtual machine migration and click **Next**.

Any compatibility problem appears in the Compatibility panel. Fix the problem, or select another host or cluster.

Possible targets include hosts and DRS clusters with any level of automation. If a cluster has no DRS enabled, select a specific host in the cluster rather than selecting the cluster.

Important:

If the virtual machine that you migrate has an NVDIMM device and uses PMem storage, the destination host or cluster must have available PMem resources. Otherwise, the compatibility check fails and you cannot proceed further with the migration.

If the virtual machine that you migrate does not have an NVDIMM device but it uses PMem storage, you must select a host or cluster with available PMem resources, so that all PMem hard disks remain stored on a PMem datastore. Otherwise, all the hard disks use the storage policy and datastore selected for the configuration files of the virtual machine.

Important:

Migrating a virtual machine that has an NVDIMM device or a vPMem disk to a host that does not have the proper license fails and leaves the virtual machine in an unmanageable state for 90 seconds. You can afterwards retry the migration and select a destination host that is licensed to use PMem devices.

- If you change the storage of the virtual machine, enter the required details in the **Select Storage** page.
 - Select the storage type for the virtual machine configuration files and all the hard disks.
 - If you select the **Standard** mode, all virtual disks are stored on a standard datastore.
 - If you select the **PMem** mode, all virtual disks are stored on the host-local PMem datastore. Configuration files cannot be stored on a PMem datastore and you must additionally select a regular datastore for the configuration files of the virtual machine.
 - If you select the **Hybrid** mode, all PMem virtual disks remain stored on a PMem datastore. Non-PMem disks are affected by your choice of a VM storage policy and datastore or datastore cluster.

Selecting the type of storage is possible only if PMem or Hybrid storage types are available in the data center.

- b) Select the format for the virtual machine disks.
- c) Select a virtual machine storage policy from the **VM Storage Policy** drop-down menu.

Storage policies specify storage requirements for applications that run on the virtual machine. You can also select the default policy for vSAN or Virtual Volumes datastores.

Important:

If the virtual machine hard disks use different storage policies, the new policy that you select only applies to non-PMem hard disks. PMem hard disks are migrated to the host-local PMem datastore of the destination host.

- d) Select the datastore location where you want to store the virtual machine files.

Option	Action
Store all virtual machine files in the same location on a datastore.	Select a datastore from the list and click Next .
Store all virtual machine files in the same Storage DRS cluster.	1. Select a Storage DRS cluster. 2. To deactivate Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 3. Click Next.
Store virtual machine configuration files and disks in separate locations.	1. Click Configure per disk. Note: You can use the Configure per disk option to downgrade from or upgrade to PMem storage. 2. For the virtual machine configuration file and for each virtual disk, select Browse, and select a datastore or Storage DRS cluster. Note: Configuration files cannot be stored on a PMem datastore. 3. If you selected a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 4. Click Next.

- 6. If you change the compute resource of the virtual machine, select destination networks for the virtual machine migration.

You can migrate a virtual machine network to another distributed switch in the same or to another data center or vCenter Server.

Option	Action
Select a destination network for all VM network adapters connected to a valid source network.	1. Click the arrow in the Destination Network column and select Browse. 2. Select a destination network and click OK. 3. Click Next.
Select a new destination network for each VM network adapter connected to a valid source network.	1. Click Advanced. 2. Click the arrow in the Destination Network column and select Browse. 3. Select a destination network and click OK. 4. Click Next.

7. On the **Ready to complete** page, review the details and click **Finish**.

vCenter Server moves the virtual machine to the new host or storage location.

Event messages appear in the **Events** tab. The data displayed on the **Summary** tab shows the status and state throughout the migration. If errors occur during migration, the virtual machines revert to their original states and locations.

Migrating Virtual Machines with vSphere vMotion

If you must take a host offline for maintenance, you can move the virtual machine to another host. Migration with vSphere vMotion allows virtual machine processes to continue working throughout a migration.

When you migrate a virtual machine with vSphere vMotion, the new host for the virtual machine must meet compatibility requirements so that the migration can proceed.

vMotion Migration Types

With vSphere vMotion, you can change the compute resource on which a virtual machine is running. You also can change both the compute resource and the storage of the virtual machine.

When you migrate virtual machines with vSphere vMotion and choose to change only the host, the entire state of the virtual machine is moved to the new host. The associated virtual disk remains in the same location on storage that must be shared between the two hosts.

When you choose to change both the host and the datastore, the virtual machine state is moved to a new host and the virtual disk is moved to another datastore. vSphere vMotion to another host and datastore is possible in vSphere environments without a shared storage.

After the virtual machine state is migrated to the alternate host, the virtual machine runs on the new host. Migrations with vSphere vMotion are transparent to the running virtual machine.

When you choose to change both the compute resource and the storage, you can use vSphere vMotion to migrate virtual machines across vCenter Server instances, data centers, and subnets.

Transferred State Information

The state information includes the current memory content and all the information that defines and identifies the virtual machine. The memory content includes transaction data and the bits of the operating system and applications that are in the memory. The defining and identification information stored in the state includes all the data that maps to the virtual machine hardware elements. This information includes BIOS, devices, CPU, MAC addresses for the Ethernet cards, chipset states, registers, and so forth.

Stages in vSphere vMotion

Migration with vSphere vMotion occurs in three stages:

1. When you initiate a migration with vSphere vMotion, vCenter Server verifies that the existing virtual machine is in a stable state within its current host.
2. The virtual machine state information (memory, registers, and network connections) is copied to the target host.
3. The virtual machine resumes its activities on the new host.

If errors occur during migration, the virtual machine reverts to its original state and location.

Host Configuration for vSphere vMotion

Before using vSphere vMotion, you must configure your hosts correctly.

- Each host must be licensed for vSphere vMotion.

- Each host must meet shared storage requirements for vSphere vMotion. See [vSphere vMotion Shared Storage Requirements](#).
- Each host must meet the networking requirements for vSphere vMotion. See [What Are the vSphere vMotion Networking Requirements](#).

vSphere vMotion Across Long Distances

You can perform reliable migrations between hosts and sites that are separated by high network round-trip latency times. vSphere vMotion across long distances is enabled when the appropriate license is installed. No user configuration is necessary.

For long-distance migration, verify the network latency between the hosts and your license.

- The round-trip time between the hosts must be up to 150 milliseconds.
- Your license must cover vSphere vMotion across long distances.
- You must place the traffic related to transfer of virtual machine files to the destination host on the provisioning TCP/IP stack. See [How to Place Traffic for Cold Migration, Cloning, and Snapshots on the Provisioning TCP/IP Stack](#).

vSphere vMotion Shared Storage Requirements

Configure your hosts for vSphere vMotion with shared storage to ensure that virtual machines are accessible to both source and target hosts.

During a migration with vMotion, the migrating virtual machine must be on storage accessible to both the source and target hosts. Ensure that the hosts configured for vMotion use shared storage. Shared storage can be on a Fibre Channel storage area network (SAN), or can be implemented using iSCSI and NAS.

If you use vMotion to migrate virtual machines with raw device mapping (RDM) files, make sure to maintain consistent LUN IDs for RDMs across all participating hosts.

See the *vSphere Storage* documentation for information on SANs and RDMs.

What Are the vSphere vMotion Networking Requirements

Migration with vMotion requires correctly configured network interfaces on source and target hosts.

Configure each host with at least one network interface for vMotion traffic. To ensure secure data transfer, the vMotion network must be a secure network, accessible only to trusted parties. Additional bandwidth significantly improves vMotion performance. When you migrate a virtual machine with vMotion without using shared storage, the contents of the virtual disk is transferred over the network as well.

vSphere 6.5 and later allow the network traffic with vMotion to be encrypted. Encrypted vMotion depends on host configuration, or on compatibility between the source and destination hosts.

Requirements for Concurrent vMotion Migrations

You must ensure that the vMotion network has at least 250 Mbps of dedicated bandwidth per concurrent vMotion session. Greater bandwidth lets migrations complete more quickly. Gains in throughput resulting from WAN optimization techniques do not count towards the 250-Mbps limit.

To determine the maximum number of concurrent vMotion operations possible, see [vCenter Server Limits on Simultaneous Migrations](#). These limits vary with a host's link speed to the vMotion network.

Round-Trip Time for Long-Distance vMotion Migration

If you have the proper license applied to your environment, you can perform reliable migrations between hosts that are separated by high network round-trip latency times. The maximum supported network round-trip time for vMotion migrations is 150 milliseconds. This round-trip time lets you migrate virtual machines to another geographical location at a longer distance.

Multiple-NIC vMotion

You can configure multiple NICs for vMotion by adding two or more NICs to the required standard or distributed switch. For details, see Knowledge Base article [KB 2007467](#).

Network Configuration

Configure the virtual networks on vMotion enabled hosts as follows:

- On each host, configure a VMkernel port group for vMotion.
To have the vMotion traffic routed across IP subnets, enable the vMotion TCP/IP stack on the host. See [How to Place vSphere vMotion Traffic on the vMotion TCP/IP Stack of Your ESXi Host](#).
- If you are using standard switches for networking, ensure that the network labels used for the virtual machine port groups are consistent across hosts. During a migration with vMotion, vCenter Server assigns virtual machines to port groups based on matching network labels.

Note:

By default, you cannot use vMotion to migrate a virtual machine that is attached to a standard switch with no physical uplinks configured, even if the destination host also has a no-uplink standard switch with the same label.

To override the default behavior, set the `config.migrate.test.CompatibleNetworks.VMOnVirtualIntranet` advanced settings of vCenter Server to `false`. The change takes effect immediately. For information about configuring advanced settings of vCenter Server, see [vCenter Server Configuration](#).

- If you use Intrusion Detection Systems (IDS) and firewalls to protect your environment, make sure that you configure them to allow connections to the ports used for vMotion on the ESXi hosts. For the list of currently supported ports for vMotion, see the VMware Ports and Protocols Tool™ at <https://ports.esp.vmware.com/home/vSphere>.

For information about configuring the vMotion network resources, see [Networking Best Practices for vSphere vMotion](#).

For more information about vMotion networking requirements, see Knowledge Base article [KB 59232](#).

Networking Best Practices for vSphere vMotion

Consider certain best practices for configuring the network resources for vMotion on an ESXi host.

- Provide the required bandwidth in one of the following ways:

Physical Adapter Configuration	Best Practices
Dedicate at least one adapter for vMotion or reserve network bandwidth for vMotion.	Use at least one 10 GbE adapter for workloads that have a small number of memory operations, or if you migrate workloads that have many memory operations. If only two Ethernet adapters are available, for best availability, combine both adapters into a team, and use VLANs to divide traffic into networks: one or more for virtual machine traffic and one for vMotion. Alternatively, you can reserve at least 10 GbE of network bandwidth for vMotion on a high speed link.

Physical Adapter Configuration	Best Practices
Direct vMotion traffic to one or more physical NICs that have high-bandwidth capacity and are shared between other types of traffic as well	• To distribute a • nd allocate more bandwidth to vMotion traffic across several physical NICs, use multiple-NIC vMotion. • On a vSphere Distributed Switch 5.1 and later, use vSphere Network I/O Control shares to guarantee bandwidth to outgoing vMotion traffic. Defining shares also prevents from contention as a result from excessive vMotion or other traffic. • In vSphere 7.0 Update 1 or earlier, vMotion saturates 1 GbE and 10 GbE physical NICs with a single vMotion VMkernel NIC. Starting with vSphere 7.0 Update 2, vMotion saturates high speed links such as 25 GbE, 40 GbE and 100 GbE with a single vMotion VMkernel NIC. If you do not have dedicated uplinks for vMotion, you can use Network I/O Control to limit vMotion bandwidth use. • To avoid saturation of the physical NIC link as a result of intense incoming vMotion traffic, use traffic shaping in egress direction on the vMotion port group on the destination host. By using traffic shaping you can limit the average and peak bandwidth available to vMotion traffic, and reserve resources for other traffic types.

- Use jumbo frames for best vMotion performance.
Ensure that jumbo frames are enabled on all network devices that are on the vMotion path including physical NICs, physical switches, and virtual switches.
- Place vMotion traffic on the vMotion TCP/IP stack for migration across IP subnets that have a dedicated default gateway that is different from the gateway on the management network. See [How to Place vSphere vMotion Traffic on the vMotion TCP/IP Stack of Your ESXi Host](#).

For information about configuring networking on an ESXi host, see the *vSphere Networking* documentation.

Virtual Machine Conditions and Limitations for vSphere vMotion

To migrate virtual machines with vMotion, the virtual machine must meet certain network, disk, CPU, USB, and other device requirements.

The following virtual machine conditions and limitations apply when you use vSphere vMotion:

- The source and destination management network IP address families must match. You cannot migrate a virtual machine from a host that is registered to vCenter Server with an IPv4 address to a host that is registered with an IPv6 address.
- Using 1 GbE network adapters for the vSphere vMotion network might result in migration failure, if you migrate virtual machines with large vGPU profiles. Use 10 GbE network adapters for the vSphere vMotion network.
- If virtual CPU performance counters are enabled, you can migrate virtual machines only to hosts that have compatible CPU performance counters.
- You can migrate virtual machines that have 3D graphics enabled. If the 3D Renderer is set to Automatic, virtual machines use the graphics renderer that is present on the destination host. The renderer can be the host CPU or a GPU graphics card. To migrate virtual machines with the 3D Renderer set to Hardware, the destination host must have a GPU graphics card.
- Starting with vSphere 6.7 Update 1 and later, vSphere vMotion supports virtual machines with vGPU.
- vSphere DRS supports initial placement of vGPU virtual machines running vSphere 6.7 Update 1 or later without load balancing support.
- You can migrate virtual machines with USB devices that are connected to a physical USB device on the host. You must enable the devices for vSphere vMotion.

- You cannot use migration with vSphere vMotion to migrate a virtual machine that uses a virtual device backed by a device that is not accessible on the destination host. For example, you cannot migrate a virtual machine with a CD drive backed by the physical CD drive on the source host. Disconnect these devices before you migrate the virtual machine.
- You cannot use migration with vSphere vMotion to migrate a virtual machine that uses a virtual device backed by a device on the client computer. Disconnect these devices before you migrate the virtual machine.

Using vSphere vMotion to Migrate vGPU Virtual Machines

You can use vMotion to perform a live migration of NVIDIA vGPU-powered virtual machines without causing data loss.

To enable vMotion for vGPU virtual machines, you need to set the `vgpu.hotmigrate.enabled` advanced setting to `true`. For more information about how to configure the vCenter Server advanced settings, see *Configure Advanced Settings* in the *vCenter Server Configuration* documentation.

In vSphere 6.7 Update 1 and vSphere 6.7 Update 2, when you migrate vGPU virtual machines with vMotion and vMotion stun time exceeds 100 seconds, the migration process might fail for vGPU profiles with 24 GB frame buffer size or larger. To avoid the vMotion timeout, upgrade to vSphere 6.7 Update 3 or later.

During the stun time, you are unable to access the VM, desktop, or application. Once the migration is completed, access to the VM resumes and all applications continue from their previous state. For information on frame buffer size in vGPU profiles, refer to the [NVIDIA Virtual GPU documentation](#).

The expected VM stun times (the time when the VM is inaccessible to users during vMotion) and the estimated worst-case stun times are listed in the following tables. The expected stun times were tested over a 10Gb network with NVIDIA Tesla V100 PCIe 32 GB GPUs:

Table 146: Expected Stun Times for vMotion of vGPU VMs

Used vGPU Frame Buffer (GB)	VM Stun Time (sec)
1	2
2	4
4	6
8	12
16	22
32	39

Table 147: Estimated Worst-Case Stun Times (sec)

vGPU Memory	VM Memory 4 GB	VM Memory 8 GB	VM Memory 16 GB	VM Memory 32 GB
1 GB	5	6	8	12
2 GB	7	9	11	15
4 GB	13	14	16	21
8 GB	24	25	28	32
16 GB	47	48	50	54
32 GB	91	92	95	99

Note: When you consider the expected and estimated worst-case stun times, keep in mind the following points:

- The configured vGPU profile represents an upper bound to the used vGPU frame buffer. In many use cases, the amount of vGPU frame buffer memory used by the VM at any given time is below the assigned vGPU memory in the profile.
- Both expected and estimated worst-case stun times are only valid when migrating a single virtual machine. If you are concurrently migrating multiple virtual machines, that is, for a vSphere manual remediation process, the stun times will have adverse effects.
- The above estimates assume sufficient CPU, memory, PCIe, and network capacity to achieve 10 Gbps migration throughput.

DRS supports initial placement of vGPU VMs running vSphere 6.7 Update 1 and later without load balancing support.

VMware vSphere vMotion is supported only with and between compatible NVIDIA GPU device models and NVIDIA GRID host driver versions as defined and supported by NVIDIA. For compatibility information, refer to the [NVIDIA Virtual GPU User Guide](#).

To check compatibility between NVIDIA vGPU host drivers, vSphere, and Horizon, refer to the [VMware Compatibility Matrix](#).

Related tasks

- [How to Migrate Your Virtual Machine to a New Compute Resource](#)
- [How to Migrate Your Virtual Machine to a New Compute Resource and Storage](#)
- [How to Migrate Your Virtual Machine with Storage vMotion](#)

How to Set a Stun Time Limit for Your vGPU Virtual Machines

Learn how to set a stun time limit per virtual machine when you migrate NVIDIA vGPU-powered virtual machines with vSphere vMotion. Setting a stun time limit may prevent vCenter Server from powering on the virtual machine or migrating it to a host and network whose estimated maximum stun time exceeds that limit.

- Verify that the vCenter Server instance is version 8.0 Update 2.
- Verify that the source and destination ESXi hosts are version 8.0 Update 2.
- Verify that vMotion is enabled for vGPU virtual machines. The `vgpu.hotmigrate.enabled` advanced setting must be set to `true`. For more information about how to configure the vCenter Server advanced settings, see *Configure Advanced Settings* in the *vCenter Server Configuration* documentation.

Virtual machine stun time is the time when the virtual machine is inaccessible to users during vMotion. Starting with vSphere 8.0 Update 2, you can set a vMotion stun time limit for vGPU virtual machines.

The limit that you set must be higher than the estimated maximum stun time for the current device configuration. In case of multiple PCI devices, the stun time limit that you set must be higher than the sum of the contributions of all PCI devices. Setting a stun time limit that is lower than the estimated maximum stun time may prevent the virtual machine from powering on.

Note: Maximum stun time estimates are only available for C-Series and Q-Series vGPU types.

The maximum stun time is calculated on the basis of the bandwidth of the host on which the VM runs currently. The calculations may change at the time of migration if the destination host has lower bandwidth. For example, if a virtual machine is running on a host with a 25 Gbps vMotion NIC, but the destination host has a 10 Gbps vMotion NIC. At the time of migration the maximum stun time calculation is based on the 10 Gbps vMotion NIC.

1. In the vSphere Client, navigate to a vGPU virtual machine.
2. Right-click the virtual machine and select **Edit Settings**.

The **Edit Settings** dialog box appears.

3. Click the **VM Options** tab and expand **Advanced**.

Edit Settings | vgpu-multivgpu

1

Virtual HardwareVM OptionsAdvanced Parameters

> General Options

VM Name: vgpu-multivgpu

> VMware Remote Console Options

Expand for VMware Remote Console settings

> Encryption

Expand for encryption settings

> VMware Tools

Expand for VMware Tools settings

> Boot Options

Expand for boot options

> Power management

Expand for power management settings

< Advanced2

Settings

☒ Enable logging

Debugging and statistics

Run normally

Swap file location

☒ Default

CANCEL

4. In the **vMotion Stun Time Limit** input field, enter the stun time limit in seconds.

CANCEL

5. Click **OK**.

Virtual Machines Swap File Location and vSphere vMotion Compatibility

You can configure ESXi6.7 or later hosts to store virtual machine swap files with the virtual machine configuration file, or on a local swap file datastore specified for that host.

- For migrations between hosts running ESXi6.7 and later, vMotion and migrations of suspended and powered-off virtual machines are allowed.
- During a migration with vMotion, if the swap file location on the destination host differs from the swap file location on the source host, the swap file is copied to the new location. This activity can result in slower migrations with vMotion. If the destination host cannot access the specified swap file location, it stores the swap file with the virtual machine configuration file.

See the *vSphere Resource Management* documentation for information about configuring swap file policies.

vSphere vMotion Notifications for Latency Sensitive Applications

Starting with vSphere 8.0, you can notify an application running inside the guest OS of a virtual machine when a vSphere vMotion event starts and finishes. This notification mechanism allows latency sensitive applications to prepare or delay a vSphere vMotion operation.

In cases of latency sensitive applications, such as VoIP applications and high frequency trading applications, vSphere vMotion and vSphere DRS are usually deactivated. vSphere 8.0 introduces a notification mechanism which allows you to notify an application that a vSphere vMotion is about to happen so that the application takes the necessary steps to prepare. vSphere vMotion pauses after generating the start event and waits for the application to acknowledge the start notification before it proceeds.

Note:

Enabling the notification mechanism for vSphere vMotion operations might result in an increase of the overall vSphere vMotion time.

How to Configure a Virtual Machine for vSphere vMotion Notifications

By default, the notification mechanism is deactivated. To enable the notification mechanism on a virtual machine, the virtual machine must be of hardware version 20 or above and you must configure the following advanced virtual machine configuration options.

You can use vSphere Web Services API to enable sending notifications to applications running inside the virtual machine, and to specify the maximum period of time in seconds that an application has to prepare for the vMotion operation.

- To enable sending notifications to applications running inside the virtual machine, set the `vmx.vmOpNotificationToApp.enabled` virtual machine property to `true`
- To specify the maximum period of time in seconds that an application has to prepare for the vMotion operation, use the `vmx.vmOpNotificationToApp.timeout` virtual machine property. When a vSphere vMotion operation generates a start event, the vMotion operation pauses and waits for an acknowledgment from the application to proceed. vSphere vMotion waits for the notification timeout that you specify.

This property is optional and is not set by default. You can use it to configure a more restrictive timeout for a specific virtual machine. If left unset, the host notification configuration is used.

For more information, see the virtual machine data object properties in the [vSphere Web Services API](#) documentation. For information on how to use the vSphere Web Services SDK, see the [vSphere Management SDK guide](#).

How to Configure a Host for vSphere vMotion Notifications

You can use the `VmOpNotificationToApp.Timeout` advanced host configuration setting, to specify a notification timeout that applies to all virtual machines on a host. Use the `ConfigManager` APIs to set a value to this configuration property. The default notification timeout is 0. In this case, application notifications are generated but the vSphere vMotion operations are not delayed.

If you set a notification timeout on a host and a virtual machine running on the host at the same time, the smaller value is used.

For information on how to use the vSphere Web Services SDK, see the [vSphere Management SDK guide](#).

How to Register an Application for vSphere vMotion Notifications

To enable an application to receive notifications for vSphere vMotion events, you must have VMware Tools installed on the virtual machine hosting the application. For more information about installing and configuring VMware Tools, see the [VMware Tools](#) documentation.

With vSphere 8.0, you can register only one application per virtual machine. An application that wants to receive vSphere vMotion notifications, can use the following guest RPC calls and check for new vMotion events periodically, for example, every 1 or 2 seconds, using the `vm-operation-notification.check-for-event` guest RPC call.

For information on how to use guest RPC calls, see the [vSphere Guest SDK guide](#).

Command	Description
<code>vm-operation-notification.register</code>	Registers an application to start receiving notifications for vSphere vMotion events.
<code>vm-operation-notification.unregister</code>	Unregisters an application from receiving notifications for vSphere vMotion events.
<code>vm-operation-notification.list</code>	Retrieves information about the registered application running on a virtual machine on the host.
<code>vm-operation-notification.check-for-event</code>	Retrieves information about the vSphere vMotion event registered at the time of the call.
<code>vm-operation-notification.ack-event</code>	Acknowledges a vSphere vMotion start event.

How to Migrate Your Virtual Machine to a New Compute Resource

Learn how to use vSphere vMotion and the **Migration** wizard to migrate a powered-on virtual machine from one compute resource to another. To relocate only the disks of a powered-on virtual machine, migrate the virtual machine to a new datastore by using Storage vMotion.

Verify that your hosts and virtual machines meet the requirements for migration with vMotion with shared storage.

- Required privilege: **Resource > Migrate powered on virtual machine**

1. Right-click the virtual machine and select **Migrate**.
 - a) To locate a virtual machine, select a data center, folder, cluster, resource pool, host, or vApp.
 - b) Click the **Virtual Machines** tab.
2. Click **Change compute resource only** and click **Next**.

3. Select a host, cluster, resource pool, or vApp to run the virtual machine, and click **Next**.

Any compatibility problem appears in the Compatibility panel. Fix the problem, or select another host or cluster.

Possible targets include hosts and fully automated DRS clusters in the same or another vCenter Server system. If your target is a non-automated cluster, select a host within the non-automated cluster.

Important:

If the virtual machine that you migrate has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have available PMem resources. Otherwise, the compatibility check fails and you cannot proceed further with the migration.

If the virtual machine that you migrate does not have an NVDIMM device, but it has virtual PMem hard disks, the destination host or cluster must have available PMem resources, so that all PMem hard disks remain stored on a PMem datastore. Otherwise, all the hard disks use the storage policy and datastore selected for the configuration files of the virtual machine.

Important: Migrating a virtual machine that has an NVDIMM device or a vPMem disk to a host that does not have the proper license fails and leaves the virtual machine in an unmanageable state for 90 seconds. You can afterwards retry the migration and select a destination host that is licensed to use PMem devices.

4. Select a destination network for all VM network adapters connected to a valid source network and click **Next**.

You can click **Advanced** to select a new destination network for each VM network adapter connected to a valid source network.

You can migrate a virtual machine network to another distributed switch in the same or to another data center or vCenter Server.

5. Select the migration priority level and click **Next**.

Option	Description
Schedule vMotion with high priority	vCenter Server attempts to reserve resources on both the source and destination hosts to be shared among all concurrent migrations with vMotion. vCenter Server grants a larger share of host CPU resources. If sufficient CPU resources are not immediately available, vMotion is not initiated.
Schedule regular vMotion	vCenter Server reserves resources on both the source and destination hosts to be shared among all concurrent migration with vMotion. vCenter Server grants a smaller share of host CPU resources. If there is a lack of CPU resources, the duration of vMotion can be extended.

6. Review the page and click **Finish**.

vCenter Server moves the virtual machine to the new host or storage location.

Event messages appear in the **Events** tab. The data displayed on the **Summary** tab shows the status and state throughout the migration. If errors occur during migration, the virtual machines revert to their original states and locations.

What is Migration with Storage vMotion

With Storage vMotion, you can migrate a virtual machine and its disk files from one datastore to another while the virtual machine is running. You can also move virtual machines off of arrays for maintenance or to upgrade.

With Storage vMotion you also have the flexibility to optimize disks for performance, or to transform disk types, which you can use to reclaim space.

You can choose to place the virtual machine and all its disks in a single location, or you can select separate locations for the virtual machine configuration file and each virtual disk. The virtual machine does not change execution host during a migration with Storage vMotion.

During a migration with Storage vMotion, you can change the disk provisioning type.

Migration with Storage vMotion changes virtual machine files on the destination datastore to match the inventory name of the virtual machine. The migration renames all virtual disk, configuration, snapshot, and `.nvram` files. If the new names exceed the maximum filename length, the migration does not succeed.

Storage vMotion has several uses in administering virtual infrastructure, including the following examples of use.

- Storage maintenance and reconfiguration. You can use Storage vMotion to move virtual machines off a storage device to allow maintenance or reconfiguration of the storage device without virtual machine downtime.
- Redistributing storage load. You can use Storage vMotion to redistribute virtual machines or virtual disks to different storage volumes to balance capacity or improve performance.

Storage vMotion Requirements and Limitations

A virtual machine and its host must meet resource and configuration requirements for the virtual machine disks to be migrated with Storage vMotion.

Storage vMotion is subject to the following requirements and limitations:

- Virtual machine disks must be in persistent mode or be raw device mappings (RDMs). For virtual compatibility mode RDMs, you can migrate the mapping file or convert to thick-provisioned or thin-provisioned disks during migration if the destination is not an NFS datastore. If you convert the mapping file, a new virtual disk is created and the contents of the mapped LUN are copied to this disk. For physical compatibility mode RDMs, you can migrate the mapping file only.
- Migration of virtual machines during VMware Tools installation is not supported.
- Because VMFS3 datastores do not support large capacity virtual disks, you cannot move virtual disks greater than 2 TB from a VMFS5 datastore to a VMFS3 datastore.
- The host on which the virtual machine is running must have a license that includes Storage vMotion.
- ESXi 4.0 and later hosts do not require vMotion configuration to perform migration with Storage vMotion.
- The host on which the virtual machine is running must have access to both the source and target datastores.
- For limits on the number of simultaneous migrations with vMotion and Storage vMotion, see [vCenter Server Limits on Simultaneous Migrations](#).

How to Migrate Your Virtual Machine with Storage vMotion

Learn how to use Storage vMotion to migrate the configuration file of a virtual machine and its virtual disks to a new storage. You can migrate the virtual machine while it is powered on.

- For migration of a virtual machine with NVIDIA vGPU, verify that the ESXi host on which the virtual machine runs has a free vGPU slot when the host is with version 7.0 Update 2 and earlier. Starting with vSphere 7.0 Update 3, the source hosts are not required to have a free vGPU slot.
- Verify that the `vgpu.hotmigrate.enabled` advanced setting is set to `true`. For more information about configuring vCenter Server advanced settings, see *Configure Advanced Settings* in the *vCenter Server Configuration* documentation.
- Required privilege: **Resource > Migrate powered on virtual machine**

1. Right-click the virtual machine and select **Migrate**.
 - a) To locate a virtual machine, select a data center, folder, cluster, resource pool, host, or vApp.
 - b) Click the **Virtual Machines** tab.
2. Click **Change storage only** and click **Next**.
3. Select the format for the virtual machine's disks.

Option	Action
Same format as source	Use the same format as the source virtual machine.
Thick Provision Lazy Zeroed	Create a virtual disk in a default thick format. Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation. Instead, it is zeroed out on demand on first write from the virtual machine.
Thick Provision Eager Zeroed	Create a thick disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out during creation. It might take longer to create disks in this format than to create other types of disks.
Thin Provision	Use the thin provisioned format. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can expand to the maximum capacity allocated to it.

4. Select a virtual machine storage policy from the **VM Storage Policy** drop-down menu.

Storage policies specify storage requirements for applications that run on the virtual machine. You can also select the default policy for vSAN or Virtual Volumes datastores.

Important:

If the virtual machine hard disks use different storage policies, the new policy that you select only applies to non-PMem hard disks. PMem hard disks are migrated to the host-local PMem datastore of the destination host.

5. Select the datastore location where you want to store the virtual machine files.

Option	Action
Store all virtual machine files in the same location on a datastore.	Select a datastore and click Next .
Store all virtual machine files in the same Storage DRS cluster.	1. Select a Storage DRS cluster. 2. To deactivate Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 3. Click Next.
Store virtual machine configuration files and disks in separate locations.	1. Click Advanced. Note: You can use the Advanced option to downgrade from or upgrade to PMem storage. 2. For the virtual machine configuration file and for each virtual disk, select Browse, and select a datastore or Storage DRS cluster. Note: Configuration files cannot be stored on a PMem datastore. 3. If you selected a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 4. Click Next.

6. On the Ready to complete page, review the details and click **Finish**.

vCenter Server moves the virtual machine to the new storage location. Names of migrated virtual machine files on the destination datastore match the inventory name of the virtual machine.

Event messages appear in the **Events** tab. The data displayed on the **Summary** tab shows the status and state throughout the migration. If errors occur during migration, the virtual machines revert to their original states and locations.

How to Migrate Your Virtual Machine to a New Compute Resource and Storage

Learn how to move a virtual machine from your vSphere inventory to another compute resource and move its disks or virtual machine folder to another datastore. With vMotion, you can migrate a virtual machine and its disks and files while the virtual machine is powered on.

- Required privilege: **Resource > Migrate powered on virtual machine**

Simultaneous migration to a new compute resource and datastore provides greater mobility for virtual machines by eliminating the vCenter Server boundary. Virtual machine disks or contents of the virtual machine folder are transferred over the vMotion network to reach the destination host and datastores.

To make disk format changes and preserve them, you must select a different datastore for the virtual machine files and disks. You cannot select the same datastore on which the virtual machine disks currently reside.

1. Right-click the virtual machine and select **Migrate**.
 - a) To locate a virtual machine, select a data center, folder, cluster, resource pool, host, or vApp.
 - b) Click the **Virtual Machines** tab.
2. Select **Change both compute resource and storage** and click **Next**.
3. Select a destination resource for the virtual machine, and click **Next**.

Any compatibility problems appear in the Compatibility panel. Fix the problem, or select another host or cluster.

Possible targets include hosts and fully automated DRS clusters. If your target is a non-automated cluster, select a host within the non-automated cluster.

If your environment has more than one vCenter Server instances, you can move virtual machines from one vCenter Server inventory to another.

Important:

If the virtual machine that you migrate has an NVDIMM device and uses PMem storage, the destination host or cluster must have available PMem resources. Otherwise, the compatibility check fails and you cannot proceed further with the migration.

If the virtual machine that you migrate does not have an NVDIMM device but it uses PMem storage, you must select a host or cluster with available PMem resources, so that all PMem hard disks remain stored on a PMem datastore. Otherwise, all the hard disks use the storage policy and datastore selected for the configuration files of the virtual machine.

Important:

Migrating a virtual machine that has an NVDIMM device or a vPMem disk to a host that does not have the proper license fails and leaves the virtual machine in an unmanageable state for 90 seconds. You can afterwards retry the migration and select a destination host that is licensed to use PMem devices.

4. Select the format for the virtual machine's disks.

Option	Action
Same format as source	Use the same format as the source virtual machine.
Thick Provision Lazy Zeroed	Create a virtual disk in a default thick format. Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation. Instead, it is zeroed out on demand on first write from the virtual machine.
Thick Provision Eager Zeroed	Create a thick disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out during creation. It might take longer to create disks in this format than to create other types of disks.
Thin Provision	Use the thin provisioned format. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can expand to the maximum capacity allocated to it.

5. Select a virtual machine storage policy from the **VM Storage Policy** drop-down menu.

Storage policies specify storage requirements for applications that run on the virtual machine. You can also select the default policy for vSAN or Virtual Volumes datastores.

Important:

If the virtual machine hard disks use different storage policies, the new policy that you select only applies to non-PMem hard disks. PMem hard disks are migrated to the host-local PMem datastore of the destination host.

6. Select the datastore location where you want to store the virtual machine files.

Option	Action
Store all virtual machine files in the same location on a datastore.	Select a datastore and click Next .
Store all virtual machine files in the same Storage DRS cluster.	1. Select a Storage DRS cluster. 2. To deactivate Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 3. Click Next.
Store virtual machine configuration files and disks in separate locations.	1. Click Advanced. Note: You can use the Advanced option to downgrade from or upgrade to PMem storage. 2. For the virtual machine configuration file and for each virtual disk, select Browse, and select a datastore or Storage DRS cluster. Note: Configuration files cannot be stored on a PMem datastore. 3. If you selected a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 4. Click Next.

7. Select a destination network for all VM network adapters connected to a valid source network and click **Next**.

You can click **Advanced** to select a new destination network for each VM network adapter connected to a valid source network.

You can migrate a virtual machine network to another distributed switch in the same or to another data center or vCenter Server.

8. Select the migration priority level and click **Next**.

Option	Description
Schedule vMotion with high priority	vCenter Server attempts to reserve resources on both the source and destination hosts to be shared among all concurrent migrations with vMotion. vCenter Server grants a larger share of host CPU resources. If sufficient CPU resources are not immediately available, vMotion is not initiated.
Schedule regular vMotion	vCenter Server reserves resources on both the source and destination hosts to be shared among all concurrent migration with vMotion. vCenter Server grants a smaller share of host CPU resources. If there is a lack of CPU resources, the duration of vMotion can be extended.

9. On the Ready to complete page, review the details and click **Finish**.

vCenter Server moves the virtual machine to the new host or storage location.

Event messages appear in the **Events** tab. The data displayed on the **Summary** tab shows the status and state throughout the migration. If errors occur during migration, the virtual machines revert to their original states and locations.

vSphere vMotion in Environments Without Shared Storage

You can use vSphere vMotion to migrate virtual machines to a different compute resource and storage simultaneously. Unlike vSphere Storage vMotion, which requires a single host to have access to both the source and destination datastore, you can migrate virtual machines across storage accessibility boundaries.

vSphere vMotion does not require environments with shared storage. This is useful for performing cross-cluster migrations, when the target cluster machines might not have access to the storage of the source cluster. Processes that are working on the virtual machine continue to run during the migration with vSphere vMotion.

You can use vSphere vMotion to migrate virtual machines across vCenter Server instances.

You can place the virtual machine and all its disks in a single location or select separate locations for the virtual machine configuration file and each virtual disk. In addition, you can change virtual disks from thick-provisioned to thin-provisioned or from thin-provisioned to thick-provisioned. For virtual compatibility mode RDMs, you can migrate the mapping file or convert from RDM to VMDK.

vSphere vMotion without shared storage is useful for virtual infrastructure administration tasks similar to vSphere vMotion with shared storage or vSphere Storage vMotion tasks.

- Host maintenance. You can move virtual machines from a host to allow maintenance of the host.
- Storage maintenance and reconfiguration. You can move virtual machines from a storage device to allow maintenance or reconfiguration of the storage device without virtual machine downtime.
- Storage load redistribution. You can manually redistribute virtual machines or virtual disks to different storage volumes to balance capacity or improve performance.

Requirements and Limitations for vMotion Without Shared Storage

If you want to migrate a virtual machine with vMotion in the absence of shared storage, the virtual machine and its host must meet resource and configuration requirements for the virtual machine files and disks.

vMotion in an environment without shared storage is subject to the following requirements and limitations:

- The hosts must be licensed for vMotion.
- The hosts must be running ESXi 5.1 or later.
- The hosts must meet the networking requirement for vMotion. See [What Are the vSphere vMotion Networking Requirements](#).
- The virtual machines must be properly configured for vMotion. See [Virtual Machine Conditions and Limitations for vSphere vMotion](#).
- Virtual machine disks must be in persistent mode or be raw device mappings (RDMs). See [Storage vMotion Requirements and Limitations](#).
- The destination host must have access to the destination storage.
- When you move a virtual machine with RDMs and do not convert those RDMs to VMDKs, the destination host must have access to the RDM LUNs.
- Consider the limits for simultaneous migrations when you perform a vMotion migration without shared storage. This type of vMotion counts against the limits for both vMotion and Storage vMotion, so it consumes both a network resource and 16 datastore resources. See [vCenter Server Limits on Simultaneous Migrations](#).

When you migrate a powered-on virtual machine with vMotion in the absence of shared storage, the migration goes over the network using the vMotion vmknix, and the data is transferred directly between the two ESXi hosts.

Migrating Virtual Machines Between vCenter Server Systems

You can migrate powered-on and powered-off virtual machines between vCenter Server instances. With Advanced Cross vCenter vMotion, you can also move or clone workloads across vCenter Server systems both between on-premises environments and between cloud and on-premises environments.

Cold and Hot migration of virtual machines across vCenter Server systems is helpful in certain VM provisioning cases.

- Balance workloads across clusters and vCenter Server instances.
- Elastically expand or shrink capacity across resources in different vCenter Server instances in the same site or in another geographical area.
- Move virtual machines between environments that have different purposes, for example, from a development to production.
- Move virtual machines to meet different Service Level Agreements (SLAs) regarding storage space, performance, and so on.

Note: During the migration of a virtual machine to another vCenter Server system, the performance data that has been collected about the virtual machine is lost.

What Are the Requirements for vMotion Between vCenter Server Instances

You can use vSphere vMotion across vCenter Server instances if your system meets certain requirements.

The following list sums the requirements that your system must meet so that you can use migration across vCenter Server instances:

- When you migrate virtual machines across vCenter Server instances in Enhanced Linked Mode, without using Advanced Cross vCenter vMotion, verify that you meet the following requirements.
 - Both vCenter Server instances must be in the same vCenter Single Sign-On domain. Enhanced Link Mode lets the source vCenter Server authenticate to the destination vCenter Server.
For information about installing vCenter Server in Enhanced Linked Mode, see the *vCenter Server Installation and Setup* documentation.
 - The cross vCenter Server and long-distance vSphere vMotion features require a vSphere Enterprise Plus license.
 - The source and destination vCenter Server instances and ESXi hosts must be 6.7 or later.
 - Both vCenter Server instances must be time-synchronized with each other for correct vCenter Single Sign-On token verification.
 - For migration of compute resources only, both vCenter Server instances must be connected to the shared virtual machine storage.
- When you migrate virtual machines across vCenter Server instances in different vCenter Single Sign-On domains with Advanced Cross vCenter vMotion, verify that you meet the following requirements.
 - The vCenter Server instance from which you initiate the import or export of virtual machines must be version 7.0 Update 1c or later.
 - To vMotion powered-on virtual machines with the Advanced Cross vCenter vMotion feature, you must have a vSphere Enterprise Plus license on both the source and destination vCenter Server instances.
 - To migrate powered-off virtual machines with the Advanced Cross vCenter vMotion feature, you must have a vSphere Standard license.

Related Links

[Network Compatibility Checks During vMotion Between vCenter Server Instances on page 1878](#)

Migration of virtual machines between vCenter Server instances moves virtual machines to new networks. The migration process performs checks to verify that the source and destination networks are similar.

[MAC Address Management During Migration Between vCenter Server Systems on page 1878](#)

When you move a virtual machine between vCenter Server instances, the environment specifically handles MAC address migration to avoid address duplication and loss of data in the network.

[Export or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1879](#)

Learn how to export or clone workloads to a target vCenter Server system by using Advanced Cross vCenter vMotion. You can initiate migration of workloads both from on-premise environments and from cloud environments.

[Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1884](#)

Learn how to use Advanced Cross vCenter vMotion to import or clone workloads from a source vCenter Server system. You can initiate migration of workloads both from on-premise environments and from cloud environments.

Network Compatibility Checks During vMotion Between vCenter Server Instances

Migration of virtual machines between vCenter Server instances moves virtual machines to new networks. The migration process performs checks to verify that the source and destination networks are similar.

vCenter Server performs network compatibility checks to prevent the following configuration problems:

- vMotion between distributed switches of different versions
- vMotion from an external network to an internal network
- vMotion to a distributed switch that is not working properly

vCenter Server does not perform checks for and notify you about the following problems:

- If the source and destination distributed switches are not in the same broadcast domain, virtual machines lose network connectivity after migration.
- If the source and destination distributed switches do not have the same services configured, virtual machines might lose network connectivity after migration.

Related Links

[What Are the Requirements for vMotion Between vCenter Server Instances on page 1877](#)

You can use vSphere vMotion across vCenter Server instances if your system meets certain requirements.

[MAC Address Management During Migration Between vCenter Server Systems on page 1878](#)

When you move a virtual machine between vCenter Server instances, the environment specifically handles MAC address migration to avoid address duplication and loss of data in the network.

[Export or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1879](#)

Learn how to export or clone workloads to a target vCenter Server system by using Advanced Cross vCenter vMotion. You can initiate migration of workloads both from on-premise environments and from cloud environments.

[Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1884](#)

Learn how to use Advanced Cross vCenter vMotion to import or clone workloads from a source vCenter Server system. You can initiate migration of workloads both from on-premise environments and from cloud environments.

MAC Address Management During Migration Between vCenter Server Systems

When you move a virtual machine between vCenter Server instances, the environment specifically handles MAC address migration to avoid address duplication and loss of data in the network.

In an environment with multiple vCenter Server instances, when a virtual machine is migrated, its MAC addresses are transferred to the target vCenter Server. The source vCenter Server adds the MAC addresses to a denylist so that it does not assign them to newly created virtual machines.

To reclaim unused MAC addresses from the denylist, contact VMware Technical Support for assistance.

Related Links

[What Are the Requirements for vMotion Between vCenter Server Instances on page 1877](#)

You can use vSphere vMotion across vCenter Server instances if your system meets certain requirements.

[Network Compatibility Checks During vMotion Between vCenter Server Instances on page 1878](#)

Migration of virtual machines between vCenter Server instances moves virtual machines to new networks. The migration process performs checks to verify that the source and destination networks are similar.

[Export or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1879](#)

Learn how to export or clone workloads to a target vCenter Server system by using Advanced Cross vCenter vMotion. You can initiate migration of workloads both from on-premise environments and from cloud environments.

[Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1884](#)

Learn how to use Advanced Cross vCenter vMotion to import or clone workloads from a source vCenter Server system. You can initiate migration of workloads both from on-premise environments and from cloud environments.

Export or Clone a Virtual Machine with Advanced Cross vCenter vMotion

Learn how to export or clone workloads to a target vCenter Server system by using Advanced Cross vCenter vMotion. You can initiate migration of workloads both from on-premise environments and from cloud environments.

- Obtain the credentials for the administrator accounts of the vCenter Server instances to which you want to migrate virtual machines.
- Verify that the target vCenter Server instances to which you export or clone virtual machines are version 6.7 or later.
- Verify that the source vCenter Server instance from which you export virtual machines is version 7.0 Update 1c or later if you want to export virtual machines to other vCenter Server instances.
- Verify that the source vCenter Server instance from which you clone virtual machines is version 7.0 Update 3 or later.
- Verify that the destination ESXi hosts are version 6.7 or later.
- If the virtual machine that you migrate has an NVDIMM device and uses PMem storage, verify that the destination host or cluster have available PMem resources.
- If you migrate a virtual machine that has an NVDIMM device or a vPMem disk, verify that the destination host has the proper license.
- Required privilege: **Resource.Query vMotion**
- When you migrate powered on virtual machines, verify that you meet the following requirements.
 - Required privilege: **Resource.Migrate powered on virtual machine**
 - Verify that your hosts and virtual machines meet the requirements for migration with vMotion. See [Host Configuration for vMotion](#) and [Virtual Machine Conditions and Limitations for vMotion](#).
 - Verify that your hosts and virtual machines meet the requirements for live migration. See [Requirements and Limitations for vMotion Without Shared Storage](#).
- When you migrate powered off virtual machines, verify that you meet the following requirements.
 - Required privilege: **Resource.Migrate powered off virtual machine**
 - Make sure that you are familiar with the requirements for cold migration. See [Cold Migration](#).

- Required ports:
 - Port 8000 on the ESXi host for vMotion.
 - Port 902 on the ESXi host for NFC.
 - Port 443 between the two vCenter Server instances.

For the list of all supported ports and protocols in VMware products, see the VMware Ports and Protocols Tool™ at <https://ports.broadcom.com/>.

Advanced Cross vCenter vMotion does not depend on vCenter Enhanced Linked Mode or Hybrid Linked Mode and you can migrate virtual machines across vCenter Server systems in different vCenter Single Sign-On domains. For example, when you deploy a new vCenter Server instance, you can migrate virtual machines from the earlier version of your vCenter Server instance to the newly provisioned one.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. From the inventory tree, select the host or cluster that contains the virtual machines you want to migrate.
3. On the **VMs** tab, click the **Virtual Machines** button.
4. From the virtual machines list, select the virtual machines you want to migrate.
5. Right-click the selection and click **Migrate**.
6. If a confirmation dialog box appears, click **Yes**.

The migration wizard appears.

7. Select **Cross vCenter Server export** as the migration type.
- a) Optional: To create a clone of the virtual machine on the target vCenter Server, select the **Keep VMs on the source vCenter Server (performs a VM clone operation)**. check box.
- b) Click **Next** to continue.
8. Select the target vCenter Server instance to which you want to export or clone virtual machines.

Migrate | win10-vgpu

1 Select a migration type

2 Select a target vCenter Server

3 Select a compute resource

4 Select storage

5 Ready to complete

Select a target vCenter Server

Export Virtual Machines to the selected target vCenter Server.

SAVED VCENTER SERVERS

NEW VCENTER SERVER

vCenter Server address

vCenter Server FQDN or IP address

Username

example@domain.local

Password

Password

Save vCenter Server address

☒

LOGIN

Option	Action
Saved vCenter Servers	From the drop-down menu, select the IP address or FQDN of a saved vCenter Server.
New vCenter Server	1. Enter the IP address or FQDN, user name, and password of a vCenter Server instance. Note: The check box Save vCenter Server address is selected by default. If you do not want to save the vCenter Server address, deselect the check box. 2. Click Login.

The connection to a saved vCenter Server instance is valid only for the current user session. You must enter the credentials of the target vCenter Server instances for each login session.

9. If a **Security Alert** dialog box appears, click **Yes**.

10. Click **Next** to continue.

11. Select the destination compute resource for this virtual machine migration and click **Next**.

Possible targets include hosts and DRS clusters with any level of automation. If DRS is not enabled for a cluster, select a specific host in the cluster rather than selecting the cluster.

Important:

If the virtual machine that you migrate does not have an NVDIMM device but it uses PMem storage, you must select a host or cluster with available PMem resources, so that all PMem hard disks remain stored on a PMem datastore. Otherwise, all the hard disks use the storage policy and datastore selected for the configuration files of the virtual machine.

Important: When you try to migrate a virtual machine that has an NVDIMM device or a vPMem disk to a host that does not have the proper license, the operation fails and the virtual machine is in an unmanageable state for 90 seconds. After 90 seconds, you can retry the migration, and select a destination host that is licensed to use PMem devices.

12. If a compatibility problem appears in the Compatibility pane, fix the problem or select another host or cluster for the migration.

13. Select the destination storage for the virtual machine migration.

a) Select the storage type for the virtual machine configuration files and all the hard disks.

Mode	Description
Standard	All virtual disks are stored on a standard datastore.
PMem	All virtual disks are stored on the local PMem datastore of the host. Configuration files cannot be stored on a PMem datastore. You must additionally select a regular datastore for the configuration files of the virtual machine.
Hybrid	All PMem virtual disks remain stored on a PMem datastore. Your choice of a VM storage policy and your choice of a datastore affects non-PMem disks.

You can select the type of storage only if PMem or Hybrid storage types are available in the data center.

b) Select the format for the virtual machine disks.

c) Select a virtual machine storage policy from the **VM Storage Policy** drop-down menu.

Storage policies specify storage requirements for applications that run on the virtual machine. You can also select the default policy for vSAN or Virtual Volumes datastores.

Important: If the virtual machine hard disks use different storage policies, the new policy that you select only applies to non-PMem hard disks. PMem hard disks are migrated to the host-local PMem datastore of the destination host.

d) Select the datastore location for the virtual machine files.

Option	Action
Store all virtual machine files in the same location on a datastore.	Select a datastore from the list and click Next .
Store all virtual machine files in the same Storage DRS cluster.	1. Select a Storage DRS cluster.

Option	Action
	- If you want to migrate the virtual machine to a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. - Click Next.
Store virtual machine configuration files and disks in separate locations.	- Click Configure per disk. Note: You can use the Configure per disk option to downgrade from or upgrade to PMem storage. - For the virtual machine configuration file and for each virtual disk, select Browse, and select a datastore or a Storage DRS cluster. Note: Configuration files cannot be stored on a PMem datastore. - If you want to migrate the virtual machine to a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. - Click Next.

14. If a compatibility problem appears in the Compatibility pane, fix the problem or select another destination storage for the migration.

15. Select the destination folder for the virtual machine migration and click **Next**.

16. Select the destination networks for the virtual machine migration.

Option	Action
Select a destination network for all VM network adapters connected to a valid source network.	- Click the arrow in the Destination Network column and select Browse. - Select a destination network and click OK. - Click Next.
Select a different destination network for each VM network adapter connected to a valid source network.	- Click Advanced. - Click the arrow in the Destination Network column and select Browse. - Select a destination network and click OK. - Click Next.

17. On the **Ready to complete** page, review the details and click **Finish**.

The virtual machines are moved or cloned to the destination folder in the new vCenter Server instance. You can monitor the migration process in the **Recent tasks** pane. If errors occur during the migration, the virtual machines revert to their original states and locations.

Related Links

[What Are the Requirements for vMotion Between vCenter Server Instances on page 1877](#)

You can use vSphere vMotion across vCenter Server instances if your system meets certain requirements.

[Network Compatibility Checks During vMotion Between vCenter Server Instances on page 1878](#)

Migration of virtual machines between vCenter Server instances moves virtual machines to new networks. The migration process performs checks to verify that the source and destination networks are similar.

[MAC Address Management During Migration Between vCenter Server Systems on page 1878](#)

When you move a virtual machine between vCenter Server instances, the environment specifically handles MAC address migration to avoid address duplication and loss of data in the network.

[Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1884](#)

Learn how to use Advanced Cross vCenter vMotion to import or clone workloads from a source vCenter Server system. You can initiate migration of workloads both from on-premise environments and from cloud environments.

Import or Clone a Virtual Machine with Advanced Cross vCenter vMotion

Learn how to use Advanced Cross vCenter vMotion to import or clone workloads from a source vCenter Server system. You can initiate migration of workloads both from on-premise environments and from cloud environments.

- Obtain the credentials for the administrator account of the vCenter Server instance from which you want to import or clone virtual machines.
- Verify that the source vCenter Server instances from which you import or clone virtual machines are version 6.7 or later.
- Verify that the target vCenter Server instance to which you import virtual machines is version 7.0 Update 1c or later.
- Verify that the target vCenter Server instance to which you clone virtual machines is version 7.0 Update 1c or later.
- Verify that the destination ESXi hosts are version 6.7 or later.
- If the virtual machine that you migrate has an NVDIMM device and uses PMem storage, verify that the destination host or cluster have available PMem resources.
- If you migrate a virtual machine that has an NVDIMM device or a vPMem disk, verify that the destination host has the proper license.
- Required privilege: **Resource.Query vMotion**
- When you migrate powered on virtual machines, verify that you meet the following requirements.
 - Required privilege: **Resource.Migrate powered on virtual machine**
 - Verify that your hosts and virtual machines meet the requirements for migration with vMotion. See [Host Configuration for vMotion](#) and [Virtual Machine Conditions and Limitations for vMotion](#).
 - Verify that your hosts and virtual machines meet the requirements for live migration. See [Requirements and Limitations for vMotion Without Shared Storage](#).
- When you migrate powered off virtual machines, verify that you meet the following requirements.
 - Required privilege: **Resource.Migrate powered off virtual machine**
 - Make sure that you are familiar with the requirements for cold migration. See [Cold Migration](#).
- Required ports:
 - Port 8000 on the ESXi host for vMotion.
 - Port 902 on the ESXi host for NFC.
 - Port 443 between the two vCenter Server instances.

For the list of all supported ports and protocols in VMware products, see the VMware Ports and Protocols Tool™ at <https://ports.broadcom.com/>.

Advanced Cross vCenter vMotion does not depend on vCenter Enhanced Linked Mode or Hybrid Linked Mode and you can migrate virtual machines across vCenter Server systems in different vCenter Single Sign-On domains. For example, when you deploy a new vCenter Server instance, you can migrate virtual machines from the earlier version of your vCenter Server instance to the newly provisioned one.

1. In the vSphere Client home page, navigate to **Home > Hosts and Clusters**.
2. From the inventory tree, right-click the host or cluster to which you want to import virtual machines.
3. Click **Import VMs**.
The **Import VMs** wizard appears.

4. Select the source vCenter Server instance from which you want to import virtual machines.

Import VMs

1 Select a source vCenter Server

2 Import Virtual Machines

3 Select a compute resource

4 Select storage

5 Ready to complete

Select a source vCenter Server

Import Virtual Machines from the selected source vCenter Server

SAVED VCENTER SERVERS

NEW VCENTER SERVER

vCenter Server address

vCenter Server FQDN or IP address

Username

example@domain.local

Password

Password

Save vCenter Server address

i

☒

LOGIN

Option	Action
Saved vCenter Servers	From the drop-down menu, select the IP address or FQDN of a saved vCenter Server instance.
New vCenter Server	1. Enter the IP address or FQDN, user name, and password of a vCenter Server instance. Note: The check box Save vCenter Server address is selected by default. If you do not want to save the vCenter Server address, deselect the check box. 2. Click Login.

The connection to a saved vCenter Server instances is valid only for the current user session. You must enter the credentials of the source vCenter Server instances for each login session.

5. If a **Security Alert** dialog box appears, click **Yes**.
6. Click **Next**.
7. Select the virtual machines to import or clone.
 - a) If you want to migrate more than one virtual machine, the selected virtual machines must be in the same power state.
 - b) To keep a clone of the virtual machine on the source vCenter Server, select the **Keep VMs on the source vCenter Server (performs a VM clone operation)**. check box.
 - c) Click **Next** to continue.

If you want to migrate more than one virtual machine, the selected virtual machines must be in the same power state.

8. Select the destination compute resource for this virtual machine migration and click **Next**.

Possible targets include hosts and DRS clusters with any level of automation. If DRS is not enabled for a cluster, select a specific host in the cluster rather than selecting the cluster.

Important:

If the virtual machine that you migrate does not have an NVDIMM device but it uses PMem storage, you must select a host or cluster with available PMem resources, so that all PMem hard disks remain stored on a PMem datastore. Otherwise, all the hard disks use the storage policy and datastore selected for the configuration files of the virtual machine.

Important: When you try to migrate a virtual machine that has an NVDIMM device or a vPMem disk to a host that does not have the proper license, the operation fails and the virtual machine is in an unmanageable state for 90 seconds. After 90 seconds, you can retry the migration, and select a destination host that is licensed to use PMem devices.

9. If a compatibility problem appears in the Compatibility pane, fix the problem or select another host or cluster for the migration.
10. Select the destination storage for the virtual machine migration.
 - a) Select the storage type for the virtual machine configuration files and all the hard disks.

Mode	Description
Standard	All virtual disks are stored on a standard datastore.
PMem	All virtual disks are stored on the local PMem datastore of the host. Configuration files cannot be stored on a PMem datastore. You must additionally select a regular datastore for the configuration files of the virtual machine.
Hybrid	All PMem virtual disks remain stored on a PMem datastore. Your choice of a VM storage policy and your choice of a datastore affects non-PMem disks.

You can select the type of storage only if PMem or Hybrid storage types are available in the data center.

b) Select the format for the virtual machine disks.

c) Select a virtual machine storage policy from the **VM Storage Policy** drop-down menu.

Storage policies specify storage requirements for applications that run on the virtual machine. You can also select the default policy for vSAN or Virtual Volumes datastores.

Important: If the virtual machine hard disks use different storage policies, the new policy that you select only applies to non-PMem hard disks. PMem hard disks are migrated to the host-local PMem datastore of the destination host.

d) Select the datastore location where you want to store the virtual machine files.

Option	Action
Store all virtual machine files in the same location on a datastore.	Select a datastore from the list and click Next .
Store all virtual machine files in the same Storage DRS cluster.	1. Select a Storage DRS cluster. 2. If you want to migrate the virtual machine to a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 3. Click Next.
Store virtual machine configuration files and disks in separate locations.	1. Click Configure per disk. Note: You can use the Configure per disk option to downgrade from or upgrade to PMem storage. 2. For the virtual machine configuration file and for each virtual disk, select Browse, and select a datastore or Storage DRS cluster. Note: Configuration files cannot be stored on a PMem datastore. 3. If you want to migrate the virtual machine to a Storage DRS cluster and do not want to use Storage DRS with this virtual machine, select Disable Storage DRS for this virtual machine and select a datastore within the Storage DRS cluster. 4. Click Next.

11. If a compatibility problem appears in the Compatibility pane, fix the problem or select another destination storage for the migration.

12. Select the destination folder for the virtual machines and click **Next**.

13. Select the destination networks for the virtual machine migration.

Option	Action
Select a destination network for all VM network adapters connected to a valid source network.	1. Click the arrow in the Destination Network column and select Browse. 2. Select a destination network and click OK. 3. Click Next.
Select a different destination network for each VM network adapter connected to a valid source network.	1. Click Advanced. 2. Click the arrow in the Destination Network column and select Browse. 3. Select a destination network and click OK. 4. Click Next.

14. On the **Ready to complete** page, review the details and click **Finish**.

The virtual machines are moved or cloned to the destination folder in the current vCenter Server instance. You can monitor the migration process in the **Recent tasks** pane. If errors occur during the migration, the virtual machines revert to their original states and locations.

Related Links

[What Are the Requirements for vMotion Between vCenter Server Instances on page 1877](#)

You can use vSphere vMotion across vCenter Server instances if your system meets certain requirements.

[Network Compatibility Checks During vMotion Between vCenter Server Instances on page 1878](#)

Migration of virtual machines between vCenter Server instances moves virtual machines to new networks. The migration process performs checks to verify that the source and destination networks are similar.

[MAC Address Management During Migration Between vCenter Server Systems on page 1878](#)

When you move a virtual machine between vCenter Server instances, the environment specifically handles MAC address migration to avoid address duplication and loss of data in the network.

[Export or Clone a Virtual Machine with Advanced Cross vCenter vMotion on page 1879](#)

Learn how to export or clone workloads to a target vCenter Server system by using Advanced Cross vCenter vMotion. You can initiate migration of workloads both from on-premise environments and from cloud environments.

Migrating Virtual Machines Between Servers with DPU Devices

vSphere 8.0 introduces the VMware vSphere® Distributed Services Engine™ feature which allow offloading of some of your networking operations from an x86 host to a data processing unit (DPU). You can migrate virtual machines in a vSphere 8.0 environment with hosts that have data processing units (DPUs).

In vSphere 8.0, if you enable network offloads to a DPU device, you can use vSphere vMotion for migrating virtual machines when both hosts have DPU devices. To enable vMotion between hosts with DPU devices, you must perform multiple steps on the virtual machines, on the vCenter Server system, and on VMware NSX.

1. Create a vSphere Distributed Switch on a data center and enable the network offloads compatibility on the vSphere Distributed Switch. For more information about how to enable the network offloads compatibility on a host with a DPU device, see [What is Network Offloads Capability](#) in the *vSphere Networking* documentation.
2. Add to the switch only hosts with DPUs from the same vendor and of the same model.
3. Deploy NSX and configure an NSX transport node for the vSphere Distributed Switch. See the *NSX Installation Guide* documentation.
4. Enable the UPT support on the virtual machine that you want to migrate to and from a host with DPU devices. See *Add a Network Adapter to a Virtual Machine* in the *vSphere Virtual Machine Administration* documentation.

You can also use a DPU device on a host as a standard NIC. In this case, no additional configuration is required and you can use vSphere vMotion to migrate virtual machines to and from hosts with and without DPU devices.

What Is Encrypted vSphere vMotion

Encrypted vSphere vMotion secures confidentiality, integrity, and authenticity of data that is transferred with vSphere vMotion. vSphere supports encrypted vMotion of unencrypted and encrypted virtual machines across vCenter Server instances.

vSphere vMotion always uses encryption when migrating encrypted virtual machines. For virtual machines that are not encrypted, you can select one of the encrypted vSphere vMotion options.

What Is Encrypted in Encrypted vSphere vMotion

For encrypted disks, the data is transmitted encrypted in all cases. For unencrypted disks, the following applies:

- If disk data is transferred within a host, that is without changing the host, you change only the datastore, the transfer is unencrypted.
- If disk data is transferred between hosts and encrypted vMotion is used, the transfer is encrypted. If encrypted vMotion is not used the transfer is unencrypted.

For virtual machines that are encrypted, migration with vSphere vMotion always uses encrypted vSphere vMotion. You cannot turn off encrypted vSphere vMotion for encrypted virtual machines.

Encrypted vSphere vMotion States for Unencrypted Virtual Machines

For virtual machines that are not encrypted, you can set encrypted vSphere vMotion to one of the following states. The default is Opportunistic.

Disabled	Do not use encrypted vSphere vMotion.
Opportunistic	Use encrypted vSphere vMotion if the source and the destination hosts support it. Only ESXi hosts of version 6.5 and later support encrypted vSphere vMotion.
Required	Allow only encrypted vSphere vMotion. If the source or the destination host does not support encrypted vSphere vMotion, migration with vSphere vMotion is not allowed.

When you encrypt a virtual machine, the virtual machine keeps a record of the current encrypted vSphere vMotion setting. If you later deactivate encryption for the virtual machine, the encrypted vMotion setting remains at Required until you change the setting explicitly. You can change the settings using **Edit Settings**.

Note: Currently, you must use the vSphere APIs to migrate or clone encrypted virtual machines across vCenter Server instances. See *vSphere Web Services SDK Programming Guide* and *vSphere Web Services API Reference*.

Migrating or Cloning Encrypted Virtual Machines Across vCenter Server Instances

vSphere vMotion supports migrating and cloning encrypted virtual machines across vCenter Server instances.

When migrating or cloning encrypted virtual machines across vCenter Server instances, the source and destination vCenter Server instances must be configured to share the key provider that was used to encrypt the virtual machine. In addition, the key provider name must be the same on both the source and destination vCenter Server instances and have the following characteristics:

- Standard key provider: The same key server (or key servers) must be in the key provider.
- Trusted key provider: The same vSphere Trust Authority service must be configured on the destination host.
- vSphere Native Key Provider: Must have the same KDK.

Note: You cannot clone or migrate an encrypted virtual machine using vSphere Native Key Provider to a standalone host, no matter that the source host resides in a cluster.

The destination vCenter Server ensures the destination ESXi host has encryption mode set, ensuring the host is cryptographically "safe."

The following privileges are required when using vSphere vMotion to migrate or clone an encrypted virtual machine across vCenter Server instances.

- Migrating: **Cryptographic operations > Migrate** on the virtual machine
- Cloning: **Cryptographic operations > Clone** on the virtual machine

Also, the destination vCenter Server must have the **Cryptographic operations > EncryptNew** privilege. If the destination ESXi host is not in "safe" mode, the **Cryptographic operations > RegisterHost** privilege must also be on the destination vCenter Server.

Certain tasks are not allowed when migrating virtual machines (non-encrypted or encrypted), either on the same vCenter Server or across vCenter Server instances.

- You cannot change the VM Storage Policy.
- You cannot perform a key change.

Note: You can change the VM Storage Policy while cloning virtual machines.

Minimum Requirements for Migrating or Cloning Encrypted Virtual Machines Across vCenter Server Instances

The minimum version requirements for migrating or cloning standard key provider encrypted virtual machines across vCenter Server instances using vSphere vMotion are:

- Both the source and destination vCenter Server instances must be on version 7.0 or later.
- Both the source and destination ESXi hosts must be on version 6.7 or later.

The minimum version requirements for migrating or cloning trusted key provider encrypted virtual machines across vCenter Server instances using vSphere vMotion are:

- The vSphere Trust Authority service must be configured for the destination host and the destination host must be attested.
- Encryption cannot change on migration. For example, an unencrypted disk cannot be encrypted while the virtual machine is migrated to the new storage.
- You can migrate a standard encrypted virtual machine onto a Trusted Host. The key provider name must be the same on both the source and destination vCenter Server instances.
- You cannot migrate a vSphere Trust Authority encrypted virtual machine onto a non-Trusted Host.

Trusted Key Provider vMotion and Cross-vCenter Server vMotion

Trusted key provider fully supports vMotion across ESXi hosts.

Cross-vCenter Server vMotion is supported, but with the following restrictions.

1. The required trusted service must be configured on the destination host and the destination host must be attested.
2. Encryption cannot change on migration. For example, a disk cannot be encrypted while the virtual machine is migrated to the new storage.

When performing cross-vCenter Server vMotion, vCenter Server checks that the trusted key provider is available on the destination host, and if the host has access to it.

vSphere Native Key Provider vMotion and Cross-vCenter Server vMotion

vSphere Native Key Provider supports vMotion and Encrypted vMotion across ESXi hosts. Cross-vCenter Server vMotion is supported if vSphere Native Key Provider is configured on the destination host.

How to Activate Encrypted vSphere vMotion on Your Virtual Machine

You can activate encrypted vSphere vMotion during virtual machine creation. You can later change the encrypted vMotion state from the virtual machine settings.

Encrypted vMotion is supported only in vSphere 6.5 and later.

You can change the encrypted vMotion state only for virtual machines that are not encrypted.

1. Right-click the virtual machine and select **Edit Settings**.
2. Select **VM Options**.
3. Click **Encryption**, and select an option from the **Encrypted VMotion** drop-down menu.

Disabled

Do not use encrypted vMotion.

Opportunistic

Use encrypted vMotion if source and destination hosts support it. Only ESXi hosts of version 6.5 and later support encrypted vMotion.

Required

Allow only encrypted vMotion. If the source or destination host does not support encrypted vMotion, migration with vMotion fails.

How to Isolate Traffic for Migrating Your Virtual Machines

You can use the vSphere vMotion TCP/IP stack to isolate traffic for vMotion and the provisioning TCP/IP stack to isolate traffic for cold migration, VM clones, and snapshots.

How to Place vSphere vMotion Traffic on the vMotion TCP/IP Stack of Your ESXi Host

Learn how to use the vMotion TCP/IP stack to isolate traffic for vMotion and to assign a dedicated default gateway, routing table, and DNS configuration for this traffic. To enable the vMotion TCP/IP stack, assign a new VMkernel adapter to it.

Verify that the host is running 6.0 or later

By using a separate TCP/IP stack, you can handle vMotion and cold migration traffic according to the topology of the network and as required for your organization:

- Route the traffic for migration of powered on or powered off virtual machines by using a default gateway. The gateway must be different from the gateway assigned to the default stack on the host.
By using a separate default gateway, you can use DHCP for IP address assignment to the VMkernel adapters for migration in a flexible way.

- Assign a separate set of buffers and sockets.
- Avoid routing table conflicts that might otherwise appear when many features are using a common TCP/IP stack.
- Isolate traffic to improve security.

- In the vSphere Client, navigate to the host.
- Click the **Configure** tab.
- Select **Networking**, and click **VMkernel adapters**.
- Click **Add networking**.
- On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.
- On the Select target device page, select the switch for the VMkernel adapter, and click **Next**.

Option	Description
Select an existing network	Use the physical adapter configuration of an existing distributed port group to send data from the VMkernel adapter to the external network.
Select an existing standard switch	Use the physical adapter configuration for the VMkernel adapter of an existing standard switch.
New vSphere standard switch	Assign a new physical adapter configuration for the VMkernel adapter on a new standard switch.

- On the Port properties page, select **vMotion** from the **TCP/IP stack** drop-down menu.

The vMotion traffic becomes the only service that is enabled. You cannot use this VMkernel adapter for traffic types other than vMotion.

8. Set the label, VLAN ID, and IP mode of the VMkernel adapter, and click **Next**.
9. Optional: On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IPv4 settings automatically	Use DHCP to obtain IP settings. A DHCP server must be present on the network.
Use static IPv4 settings	Enter the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack. Select the Override default gateway for this adapter check box and enter a gateway address, if you want to specify a different gateway for the VMkernel adapter. Note: Selected netstack should have the default gateway explicitly on each host.

10. Optional: On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses. In ESXi 6.5 and later router advertisement is enabled by default and supports the M and O flags in accordance with RFC 4861.
Static IPv6 addresses	- Click Add IPv6 address to add a new IPv6 address. - Enter the IPv6 address and subnet prefix length, and click OK. - To change the VMkernel default gateway, click Override default gateway for this adapter. The VMkernel Default Gateway address for IPv6 is obtained from the selected TCP/IP stack.

11. Review your settings selections on the Ready to complete page and click **Finish**.

After you create a VMkernel adapter on the vMotion TCP/IP stack, you can use only this stack for vMotion on this host. The VMkernel adapters on the default TCP/IP stack are deactivated for the vMotion service. If a live migration uses the default TCP/IP stack while you are configuring VMkernel adapters with the vMotion TCP/IP stack, the migration completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are deactivated for future vMotion sessions.

Assign a default gateway, and configure the DNS settings, congestion control, and maximum number of connections for the vMotion TCP/IP stack.

For more information on how to change the configuration of a TCP/IP stack on a host, see the *vSphere Networking* documentation.

How to Place Traffic for Cold Migration, Cloning, and Snapshots on the Provisioning TCP/IP Stack

Learn how to use the provisioning TCP/IP stack to isolate traffic for cold migration, VM clones, snapshots, and cold data in hot migration, and to assign a dedicated default gateway, routing table, and DNS configuration for this traffic. To enable the Provisioning TCP/IP stack, assign a new VMkernel adapter to it.

Verify that the host is running 6.0 or later

By using a separate TCP/IP stack, you can handle vMotion and cold migration traffic according to the topology of the network and as required for your organization:

- Route the traffic for migration of powered on or powered off virtual machines by using a default gateway. The gateway must be different from the gateway assigned to the default stack on the host.

By using a separate default gateway, you can use DHCP for IP address assignment to the VMkernel adapters for migration in a flexible way.

- Assign a separate set of buffers and sockets.
- Avoid routing table conflicts that might otherwise appear when many features are using a common TCP/IP stack.
- Isolate traffic to improve security.

- In the vSphere Client, navigate to the host.
- Click the **Configure** tab.
- Select **Networking**, and click **VMkernel adapters**.
- Click **Add networking**.
- On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.
- On the Select target device page, select the switch for the VMkernel adapter, and click **Next**.

Option	Description
Select an existing network	Use the physical adapter configuration of an existing distributed port group to send data from the VMkernel adapter to the external network.
Select an existing standard switch	Use the physical adapter configuration for the VMkernel adapter of an existing standard switch.
New vSphere standard switch	Assign a new physical adapter configuration for the VMkernel adapter on a new standard switch.

- On the Port properties page, select **Provisioning** from the **TCP/IP stack** drop-down menu.

The provisioning traffic becomes the only service that is enabled. You cannot use this VMkernel adapter for traffic types other than provisioning.

- Set the label, VLAN ID, and IP mode of the VMkernel adapter, and click **Next**.
- Optional: On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IPv4 settings automatically	Use DHCP to obtain IP settings. A DHCP server must be present on the network.
Use static IPv4 settings	Enter the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack. Select the Override default gateway for this adapter check box and enter a gateway address, if you want to specify a different gateway for the VMkernel adapter. Note: Selected netstack should have the default gateway explicitly on each host.

10. Optional: On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses. In ESXi 6.5 and later router advertisement is enabled by default and supports the M and O flags in accordance with RFC 4861.
Static IPv6 addresses	1. Click Add IPv6 address to add a new IPv6 address. 2. Enter the IPv6 address and subnet prefix length, and click OK. 3. To change the VMkernel default gateway, click Override default gateway for this adapter. The VMkernel Default Gateway address for IPv6 is obtained from the selected TCP/IP stack.

11. Review your settings selections on the Ready to complete page and click **Finish**.

After you create a VMkernel adapter on the provisioning TCP/IP stack, you can use only this stack for cold migration, cloning, and snapshots on this host. The VMkernel adapters on the default TCP/IP stack are deactivated for the provisioning service. If a live migration uses the default TCP/IP stack while you configure VMkernel adapters with the provisioning TCP/IP stack, the data transfer completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are deactivated for future cold migration, cross-host cloning, and snapshot sessions.

vCenter Server Limits on Simultaneous Migrations

vCenter Server places limits on the number of simultaneous virtual machine migration and provisioning operations that can occur on each host, network, and datastore.

Each operation, such as a migration with vMotion or cloning a virtual machine, is assigned a resource cost. Each host, datastore, or network resource, has a maximum cost that it can support at any one time. Any new migration or provisioning operation that causes a resource to exceed its maximum cost does not proceed immediately, but is queued until other operations complete and release resources. Each of the network, datastore, and host limits must be satisfied for the operation to proceed.

vMotion without shared storage, migrating virtual machines to a different host and datastore simultaneously, is a combination of vMotion and Storage vMotion. This migration inherits the network, host, and datastore costs associated with those operations. vMotion without shared storage is equivalent to a Storage vMotion with a network cost of 1.

Network Limits

Network limits apply only to migrations with vMotion. Network limits depend on the version of ESXi and the network type. All migrations with vMotion have a network resource cost of 1.

Table 148: Network Limits for Migration with vMotion

Operation	ESXi Version	Network Type	Maximum Cost
vMotion	6.0, 6.5, 6.7, 7.0, 8.0	1GigE	4
vMotion	6.0, 6.5, 6.7, 7.0, 8.0	10GigE	8

Datastore Limits

Datastore limits apply to migrations with vMotion and with Storage vMotion. A migration with vMotion has a resource cost of 1 against the shared virtual machine's datastore. A migration with Storage vMotion has a resource cost of 16 against the source datastore and 16 against the destination datastore.

Table 149: Datastore Limits and Resource Costs for vMotion and Storage vMotion

Operation	ESXi Version	Maximum Cost Per Datastore	Datastore Resource Cost
vMotion	6.0, 6.5, 6.7, 7.0, 8.0	128	1
Storage vMotion	6.0, 6.5, 6.7, 7.0, 8.0	128	16

Host Limits

Host limits apply to migrations with vMotion, Storage vMotion, and other provisioning operations such as cloning, deployment, and cold migration. All hosts have a maximum cost per host of 8. For example, on an ESXi 7.0 host, you can perform 2 Storage vMotion operations, or 1 Storage vMotion and 4 vMotion operations.

Table 150: Host Migration Limits and Resource Costs for vMotion, Storage vMotion, and Provisioning Operations

Operation	ESXi Version	Derived Limit Per Host	Host Resource Cost
vMotion	6.0, 6.5, 6.7, 7.0, 8.0	8	1
Storage vMotion	6.0, 6.5, 6.7, 7.0, 8.0	2	4
vMotion Without Shared Storage	6.0, 6.5, 6.7, 7.0, 8.0	2	4
Other provisioning operations	6.0, 6.5, 6.7, 7.0, 8.0	8	1

vMotion Migration Compatibility Checks

During migration, use the **Migrate Virtual Machine** wizard to view information about the compatibility checks carried out using various criteria between the destination host or cluster and the migrating virtual machine.

When you select a host or a cluster, the Compatibility panel at the bottom of the **Migrate Virtual Machine** wizard displays information about the compatibility of the selected host or cluster with the virtual machine's configuration.

Compatibility Check Results

If the virtual machine is compatible, the panel displays the message `Compatibility checks succeeded`. If the virtual machine is not compatible with either the host's or cluster's configured networks or datastores, the compatibility window might display both warnings and errors:

- Warning messages do not stop migration. Often the migration is justified and you can continue with the migration despite the warnings.
- Errors might stop migration if no error-free destination hosts are available among the selected destination hosts. In this case, if you click **Next**, the wizard displays the compatibility errors again, and you cannot proceed to the next step.

Compatibility Checks During Migration Setup

When you attempt to move only the compute resource, the **Migrate Virtual Machine** wizard examines the source and destination hosts, the target resource pool, the datastore, and the network. When you attempt to move only the storage, the wizard checks the compatibility of the same objects except for the network.

When you move compute resources and storage together, the **Migrate Virtual Machine** wizard runs fewer compatibility checks. For example, if you move the compute resource, you select the target host or cluster under a vCenter Server instance. The wizard performs all necessary validation only against the selected host, and does not check the datastores available on the destination host. When you attempt to move the virtual machine to a cluster, the **Migrate Virtual Machine** wizard examines the compatibility against the host recommendation from vSphere DRS. The wizard directly validates the compatibility of the target datastore when you select it later.

Another compatibility check is whether vMotion is enabled on the source and target hosts.

Compatibility Checks for Virtual Hardware

Effects of a specific host CPU feature on compatibility depend on whether ESXi exposes or hides them from virtual machines.

- Features that are exposed to virtual machines are not compatible when they do not match on the source and target hosts.
- Features that are not exposed to virtual machines are considered as compatible regardless of whether they match on the hosts.

Specific items of virtual machine hardware can also cause compatibility problems. For example, a virtual machine using an Enhanced VMXNET virtual NIC cannot be migrated to a host running a version of ESXi that does not support Enhanced VMXNET.

CPU Compatibility and vSphere Enhanced vMotion Compatibility

vCenter Server performs compatibility checks before it allows migration of running or suspended virtual machines to ensure that the virtual machine is compatible with the target host.

vMotion transfers the running state of a virtual machine between underlying ESXi systems. Live migration requires that the processors of the target host provide the same instructions to the virtual machine after migration that the processors of the source host provided before migration. Clock speed, cache size, and number of cores can differ between source and target processors. However, the processors must come from the same vendor class (AMD or Intel) to be vMotion compatible.

Migrations of suspended virtual machines also require that the virtual machine be able to resume execution on the target host using equivalent instructions.

When you initiate a migration with vMotion or a migration of a suspended virtual machine, the **Migrate Virtual Machine** wizard checks the destination host for compatibility. If compatibility problems prevent migration, the wizard displays an error message.

The CPU instruction set available to the operating system and to applications running in a virtual machine is determined at the time that a virtual machine is powered on. This CPU feature set is based on the following items:

- Host CPU family and model
- Settings in the BIOS that might deactivate CPU features
- ESXi version running on the host
- The compatibility setting of the virtual machine
- The guest operating system of the virtual machine

To improve CPU compatibility between hosts of varying CPU feature sets, some host CPU features can be hidden from the virtual machine by placing the host in an Enhanced vMotion Compatibility (EVC) cluster.

Note:

You can hide host CPU features from a virtual machine by applying a custom CPU compatibility mask to the virtual machine, but this is not recommended. VMware, in partnership with CPU and hardware vendors, is working to maintain

vMotion compatibility across the widest range of processors. For additional information, search the VMware Knowledge Base for the *vMotion and CPU Compatibility FAQ*.

vSphere vMotion and CPU Compatibility Scenarios

Without the use of EVC, if the vCenter Server's CPU compatibility checks discover any mismatch between user-level features of the hosts, migration is blocked, whether or not the virtual machine itself has access to those features. A mismatch between kernel-level features of the hosts blocks migration only when the virtual machine has access to a feature that the target host does not provide.

vCenter Server's CPU compatibility checks compare the CPU features available on the source host, the subset of features that the virtual machine can access, and the features available on the target host.

User-level features are non-privileged instructions used by virtual machine applications. These include SSE3, SSSE3, SSE4.1, SSE4.2, and AES. Because they are user-level instructions that bypass the virtualization layer, these instructions can cause application instability if mismatched after a migration with vMotion.

Kernel-level features are privileged instructions used by the virtual machine operating system. These include the AMD No eXecute (NX) and the Intel eXecute Disable (XD) security features.

When you attempt to migrate a virtual machine with vMotion, one of the following scenarios applies:

- The destination host feature set matches the CPU feature set of the virtual machine. CPU compatibility requirements are met, and migration with vMotion proceeds.
- The CPU feature set of the virtual machine contains features not supported by the destination host. CPU compatibility requirements are not met, and migration with vMotion cannot proceed.
Note: EVC overcomes such incompatibility by providing a "baseline" feature set for all virtual machines running in a cluster. This baseline feature set hides the differences among the clustered hosts' CPUs from the virtual machines.
- The destination host supports the feature set of the virtual machine, plus additional user-level features (such as SSE4.1) not found in the feature set of the virtual machine. CPU compatibility requirements are not met, and migration with vMotion cannot proceed.
Note: This type of incompatibility is ignored for migrations among hosts in EVC clusters.
- The destination host supports the feature set of the virtual machine, plus additional kernel-level features (such as NX or XD) not found in the feature set of the virtual machine. CPU compatibility requirements are met, and migration with vMotion proceeds. The virtual machine retains its CPU feature set while it remains powered on, allowing it to migrate freely back to the original host. However, if the virtual machine is rebooted, it acquires a new feature set from the new host. This process might cause vMotion incompatibility if you attempt to migrate the virtual machine back to the original host.

CPU Families and Feature Sets Compatibility with vSphere vMotion

Processors are grouped into families. Processors within a given family generally have similar feature sets.

Processor vendors define processor families. You can distinguish different processor versions within the same family by comparing the processors' model, stepping level, and extended features. Sometimes, processor vendors have introduced significant architectural changes within the same processor family, such as the SSSE3 and SSE4.1 instructions, and NX/XD CPU security features.

By default, vCenter Server identifies mismatches on features accessible to applications as incompatible to guarantee the stability of virtual machines after migrations with vMotion.

Server hardware's CPU specifications usually indicate whether or not the CPUs contain the features that affect vMotion compatibility.

For more information on identifying Intel processors and their features, see *Application Note 485: Intel® Processor Identification and the CPUID Instruction*, available from Intel. For more information on identifying AMD processors and their features, see *CPUID Specification*, available from AMD.

CPU Compatibility Masks and vSphere vMotion

CPU compatibility masks allow customization of the CPU features visible to a virtual machine.

vCenter Server compares the CPU features available to a virtual machine with the CPU features of the destination host to determine whether or not to allow migrations with vMotion.

To guarantee the stability of virtual machines after a migration with vMotion, VMware sets the default values for CPU compatibility masks.

When a choice between CPU compatibility or guest operating system features (such as NX/XD) exists, VMware provides check-box options to configure individual virtual machines. You can access the configuration options through the Advanced Settings option for the CPU of the virtual machine. For more control over the visibility of CPU features, you can edit the CPU compatibility mask of the virtual machine at the bit level.

CAUTION: Changing the CPU compatibility masks can result in an unsupported configuration. Do not manually change the CPU compatibility masks unless instructed to do so by VMware Support or a VMware Knowledge base article.

CPU compatibility masks cannot prevent virtual machines from accessing masked CPU features in all circumstances. In some circumstances, applications can detect and use masked features even though they are hidden from the guest operating system. In addition, on any host, applications that use unsupported methods of detecting CPU features rather than using the CPUID instruction can access masked features. Virtual machines running applications that use unsupported CPU detection methods might experience stability problems after migration.

Enhanced vMotion Compatibility

You can use the Enhanced vMotion Compatibility (EVC) feature to help ensure vMotion compatibility for the hosts in a cluster. EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. Using EVC prevents migrations with vMotion from failing because of incompatible CPUs.

Configure EVC from the cluster settings dialog box. When you configure EVC, you configure all host processors in the cluster to present the feature set of a baseline processor. This baseline feature set is called the EVC mode. EVC uses AMD-V Extended Migration technology (for AMD hosts) and Intel FlexMigration technology (for Intel hosts) to mask processor features so that hosts can present the feature set of an earlier generation of processors. The EVC mode must be equivalent to, or a subset of, the feature set of the host with the smallest feature set in the cluster.

EVC masks only those processor features that affect vMotion compatibility. Enabling EVC does not prevent a virtual machine from taking advantage of faster processor speeds, increased numbers of CPU cores, or hardware virtualization support that might be available on newer hosts.

EVC cannot prevent virtual machines from accessing hidden CPU features in all circumstances. Applications that do not follow CPU vendor recommended methods of feature detection might behave unexpectedly in an EVC environment. VMware EVC cannot be supported with ill-behaved applications that do not follow the CPU vendor recommendations. For more information about creating well-behaved applications, search the VMware Knowledge Base for the article *Detecting and Using New Features in CPUs*.

Starting with vSphere 7.0 Update 1, you can take advantage of the EVC feature for Virtual Shared Graphics Acceleration (vSGA). vSGA allows multiple virtual machines to share GPUs installed on ESXi hosts and leverage the 3D graphics acceleration capabilities.

Enhanced vMotion Compatibility Requirements for Hosts

To improve CPU compatibility between hosts that have varying CPU feature sets, you can hide some host CPU features from the virtual machines by placing the host in an Enhanced vMotion Compatibility (EVC) cluster. Hosts in an EVC cluster and hosts that you add to an existing EVC cluster must meet EVC requirements.

- Power off all virtual machines in the cluster that are running on hosts with a feature set greater than the EVC mode that you intend to enable. You can also migrate these virtual machines out of the cluster.
- All hosts in the cluster must meet the following requirements:

Requirements	Description
Supported ESXi version	ESXi6.7 or later.
vCenter Server	The host must be connected to a vCenter Server system.
CPUs	A single vendor, either AMD or Intel.
Advanced CPU features enabled	Enable these CPU features in the BIOS if they are available: • Hardware virtualization support (AMD-V or Intel VT) • AMD No eXecute(NX) • Intel eXecute Disable (XD) Note: Hardware vendors sometimes deactivate particular CPU features in the BIOS by default. You might have problems enabling EVC because the EVC compatibility checks detect the absence of features that are expected to be present for a particular CPU. If you cannot enable EVC on a system with a compatible processor, ensure that all features are enabled in the BIOS.
Supported CPUs for the EVC mode that you want to enable	To check EVC support for a specific processor or server model, see the <i>VMware Compatibility Guide</i> at http://www.vmware.com/resources/compatibility/search.php .
Configured for vMotion	See Host Configuration for vSphere vMotion .

How Do You Activate vSphere EVC on an Existing Cluster

Learn how to activate vSphere Enhanced vMotion Compatibility (EVC) on an existing cluster to ensure vMotion CPU compatibility between the hosts in the cluster. Other cluster features such as vSphere DRS and vSphere HA are fully compatible with EVC.

Verify that the hosts in the cluster meet the requirements listed in [Enhanced vMotion Compatibility Requirements for Hosts](#).

1. Select a cluster in the vSphere inventory.
2. Power off all the virtual machines on the hosts with feature sets greater than the EVC mode.

To determine the EVC mode of all virtual machines on the hosts in the cluster, see [How to Determine the EVC Mode of Your Virtual Machine](#). You must power off all powered on virtual machine with no EVC mode or with an EVC mode that is greater than the EVC mode that you plan to enable on the cluster.

3. Click the **Configure** tab, select VMware EVC, and click **Edit**.
4. Enable EVC for the CPU vendor and feature set appropriate for the hosts in the cluster, and click **OK**.

How to Change the vSphere EVC Mode for Your Cluster

Learn how to configure vSphere EVC to ensure that virtual machine migrations between hosts in the cluster do not fail because of CPU feature incompatibilities.

- Verify that all hosts in the cluster have supported CPUs for the EVC mode you want to enable. See Knowledge Base article [KB 1003212](#) for a list of supported CPUs.
- Verify that all hosts in the cluster are connected and registered on vCenter Server. The cluster cannot contain a disconnected host.
- Virtual machines must be in the following power states, depending on whether you raise or lower the EVC mode.

EVC Mode	Virtual Machine Power Action
Raise the EVC mode to a CPU baseline with more features.	Running virtual machines can remain powered on. New EVC mode features are not available to the virtual machines until they are powered off and powered back on again. A full power cycling is required. Rebooting the guest operating system or suspending and resuming the virtual machine is not sufficient. Tip: You can configure a virtual machine to power cycle when its guest OS is rebooted, by setting the <code>vmx.reboot.powerCycle</code> VM parameter to <code>TRUE</code>. When the virtual machine is powered back on, this configuration parameter is removed. To edit the virtual machine configuration settings, you can use the VMware Host Client. See <i>Edit the Configuration File Parameters in the VMware Host Client</i> topic in the <i>vSphere Single Host Management - VMware Host Client</i> guide. See the VMware vSphere Blog post for information about how to set this parameter to a group of virtual machines using VMware PowerCLI.
Lower the EVC mode to a CPU baseline with fewer features.	Power off virtual machines if they are powered on and running at a higher EVC mode than the one you intend to enable.

Note:

When you take a snapshot of a virtual machine in an EVC cluster, keep in mind the following use cases:

- If you take a snapshot of a running virtual machine before an EVC downgrade, an attempt to revert to that snapshot after the EVC downgrade fails and the virtual machine is left in a suspended state.
- If you take a snapshot of a virtual machine before an EVC upgrade, you can revert to that snapshot before or after the EVC upgrade.

To verify the EVC mode for virtual machines, see [How to Determine the EVC Mode of Your Virtual Machine](#).

Several EVC approaches are available to ensure CPU compatibility:

- If all the hosts in a cluster are compatible with a newer EVC CPU or Graphics (vSGA) mode, you can change the EVC CPU or Graphics (vSGA) mode of an existing EVC cluster.
- You can enable EVC for a cluster that does not have EVC enabled.
- You can raise the EVC mode to expose more CPU features.

- You can lower the EVC mode to hide CPU features and increase compatibility.

- Select a cluster in the inventory.
- Click the **Configure** tab.
- Under **Configuration**, select **VMware EVC**, and click **Edit**.
- On the **Change EVC Mode** page, select whether to activate or deactivate EVC.

Option	Description
Disable EVC	The EVC feature is deactivated. CPU compatibility is not enforced for the hosts in this cluster.
Enable EVC for AMD Hosts	The EVC feature is enabled for AMD hosts.
Enable EVC for Intel Hosts	The EVC feature is enabled for Intel hosts.

- From the **CPU Mode** drop-down menu, select the baseline CPU feature set that you want to enable for the cluster.
If you cannot select the EVC CPU mode, the **Compatibility** pane displays the reason, and the relevant hosts for each reason.
- Optional: From the **Graphics Mode (vSGA)** drop-down menu, select a baseline graphics feature set.
If you cannot select the EVC vSGA mode, the **Compatibility** pane displays the reason, and the relevant hosts for each reason.

Option	Description
Baseline Graphics	Applies the Baseline Graphics feature set that includes features through Direct3D 10.1/OpenGL 3.3. Note: To configure the vSGA mode to apply the Baseline Graphics set that includes features through Direct3D 10.1/OpenGL 3.3, the virtual machine must be compatible with ESXi 6.5 or later.
D3D 11.0 class features	Applies the baseline graphics feature set that includes features through Direct3D 11.0/OpenGL 4.3. Note: When you create a cluster and you enable EVC graphics mode Direct3D 11.0/OpenGL 4.3 class features, you can add only ESXi 8.0 hosts that support Direct3D 11.0/OpenGL 4.3 to the cluster. When you try to add an ESXi 8.0 host that does not support Direct3D 11.0/OpenGL 4.3 to the cluster, the operation fails with an error message.

- Click **OK**.

How to Determine the EVC Mode of Your Virtual Machine

The EVC mode of a virtual machine determines the CPU and graphics features that a host must have in order for the virtual machine to migrate to that host and power on. The EVC mode of a virtual machine is independent from the EVC mode that you configure for the cluster in which the virtual machine runs.

The EVC mode of a virtual machine is determined when the virtual machine powers on. At power-on, the virtual machine also determines the EVC mode of the cluster in which it runs. If the EVC mode of a running virtual machine or the entire EVC cluster is raised, the virtual machine does not change its EVC mode until it is powered off and powered on again. This means that the virtual machine does not use any CPU features exposed by the new EVC mode until the virtual machine is powered off and powered on again.

For example, you create an EVC cluster that contains hosts with Intel processors and you set the EVC mode to Intel "Merom" Generation (Xeon Core 2). When you power on a virtual machine in this cluster, it runs in the Intel Merom Generation (Xeon Core 2) EVC mode. If you raise the EVC mode of the cluster to Intel "Penryn" Generation (Xeon 45 nm Core 2), the virtual machine retains the lower Intel "Merom" Generation (Xeon Core 2) EVC mode. To use the feature set of the higher EVC mode, such as SSE4.1, the virtual machine must be powered off and powered on again.

1. Navigate to a cluster or a host in the vCenter Server inventory.

2. Click the **VMs > Virtual Machines** tab.

A list of all virtual machines in the selected cluster or on the selected host appears.

3. To verify the status of the CPU mode, check the **EVC CPU Mode** column.

a) If the column is not visible, click **Manage Columns** at the bottom left corner of the **Virtual Machines** table.

The **Show Columns** pop-up appears.

b) To show the **EVC CPU Mode** column, locate and select the respective check box from the list.

The **EVC CPU Mode** column shows the CPU modes of all virtual machines in the cluster or on the host.

Important: For each virtual machine, the **EVC CPU Mode** column displays the EVC mode defined at the virtual machine level.

However, if you do not configure per-VM EVC for a virtual machine, the virtual machine inherits the EVC mode of its parent cluster or host. As a result, for all virtual machines that do not have per-VM EVC configured, the **EVC CPU Mode** column displays the inherited EVC mode of the parent host or cluster.

If the virtual machine is in an EVC cluster, the EVC mode that you see in the **EVC CPU Mode** column is defined in the following manner.

- When the virtual machine is powered on, the **EVC CPU Mode** column displays either the per-VM EVC mode, or the cluster-level EVC mode.

Per-VM EVC	Cluster-Level EVC	EVC Mode for the Virtual Machine
Activated	Activated	Enabled. The EVC CPU Mode column displays the EVC mode of the virtual machine.
Deactivated	Activated	Enabled. The EVC CPU Mode column displays the EVC mode of the EVC cluster.

- When the virtual machine is powered off, the **EVC CPU Mode** column displays the per-VM EVC mode. If per-VM EVC is deactivated, the **EVC CPU Mode** column for the virtual machine is empty.

When the virtual machine is not in an EVC cluster and per-VM EVC is not configured, the EVC mode that you see in the **EVC CPU Mode** column is defined in the following manner.

- When the virtual machine is powered on, the **EVC CPU Mode** column displays the EVC mode of the parent host.
- When the virtual machine is powered off, the **EVC CPU Mode** column is empty.

4. To verify the status of the graphics mode, check the **EVC Graphics Mode (vSGA)** column.

a) If the column is not visible, click **Manage Columns** at the bottom left corner of the **Virtual Machines** table.

The **Show Columns** pop-up appears.

b) To show the **EVC Graphics Mode (vSGA)** column, locate and select the respective check box from the list.

The **EVC Graphics Mode (vSGA)** column displays the baseline graphics features set. To view the baseline graphics, you must enable **3D graphics** in the virtual machine.

vMotion Migration and ESXi Hosts with AMD Processors Without 3DNow!

If hosts in a vSphere cluster have different generations of AMD processors, some with 3DNow! instruction sets and some without, you cannot successfully migrate virtual machines between the hosts. You must use an EVC mode or CPU compatibility mask to hide the instructions.

Ensure that the cluster contains only hosts with AMD Opteron Generation 3 or newer processors.

Newer generations of AMD processors do not include 3DNow! processor instructions. The vCenter Server **AMD Opteron Gen. 3 (no 3DNow!)** EVC mode masks the 3DNow! instructions from virtual machines. You can apply this EVC mode to EVC clusters containing only AMD Opteron Generation 3 hosts. Applying this mode allows the clusters to maintain vMotion compatibility with AMD Opteron hosts that do not have 3DNow! instructions. Clusters containing AMD Opteron Generation 1 or AMD Opteron Generation 2 hosts cannot be made vMotion-compatible with hosts that do not have 3DNow! instructions.

- Enable the **AMD Opteron Gen. 3 (no 3DNow!)** EVC mode for your EVC cluster.

The steps to enable the EVC mode differ depending on whether you are creating a cluster or enabling the mode on an existing cluster, and on whether the existing cluster contains powered-on virtual machines.

Option	Description
Creating a cluster	In the New Cluster wizard, enable EVC for AMD hosts and select the AMD Opteron Gen. 3 (no 3DNow!) EVC mode.
Editing a cluster without powered-on virtual machines	In the Cluster Settings dialog box, edit the VMware EVC settings and select the AMD Opteron Gen. 3 (no 3DNow!) EVC mode.
Editing a cluster with powered-on virtual machines	The AMD Opteron Gen. 3 (no 3DNow!) EVC mode cannot be enabled while there are powered-on virtual machines in the cluster. 1. Power-off any running virtual machines in the cluster, or migrate them out of the cluster using vMotion. Migrating the virtual machines out of the cluster with vMotion allows you to delay powering off the virtual machines until a more convenient time. 2. In the Cluster Settings dialog box, edit the VMware EVC settings and select the AMD Opteron Gen. 3 (no 3DNow!) EVC mode. 3. If you migrated virtual machines out of the cluster, power them off and cold migrate them back into the cluster. 4. Power on the virtual machines.

You can now add hosts with AMD processors without 3DNow! instructions to the cluster and preserve vMotion compatibility between the new hosts and the existing hosts in the cluster.

How to View Feature Capabilities Details for Your EVC Cluster

The feature set that is exposed by an EVC cluster corresponds to the feature set of a particular type of processor. Processor feature sets are described by a set of feature flags that you examine using the CPUID instruction.

You can view the CPUID feature flags currently exposed by the hosts in an EVC cluster.

1. Select a cluster in the inventory.
2. Click the **Configure** tab.
3. Select **VMware EVC** and expand **Current Feature Capabilities Details**.

This VMware EVC panel displays the CPUID feature flags that EVC enforces for the hosts in this cluster. For information about CPUID feature flags, see the Intel and AMD websites.

What Is vSphere Developer Center

The Developer Center is a single point of entry for developers that provides tools to manage API structure and to capture user actions to translate them into executable code.

What Is vSphere API Explorer

The API Explorer allows you to browse and invoke vSphere REST APIs supported by the system and provides information and context around the API calls.

With the API Explorer, you can choose an API endpoint from your environment and retrieve a list of vSphere REST APIs. You can review details like available parameters, expected responses, and response status codes, and you can invoke the APIs against the live environment. The available APIs depend on the role of the selected endpoint.

How Do You Retrieve APIs Using vSphere API Explorer

You can use the API Explorer to retrieve available vSphere REST APIs from a selected endpoint and receive information and context around API calls.

1. On the vSphere Client home page, click **Developer Center** and select the **API Explorer** tab.
2. From the **Select Endpoint** drop-down menu, select an endpoint from the environment.
3. From the **Select API** drop-down menu, select an API. The listed APIs are the ones publicly provided by the API explorer in vCenter Server.
4. Optional: You can use the filter text box to filter your results. For example, enter `health` to view a list of methods related to monitoring the health of the selected API.
5. Select an API category from the list.
6. Select a method from the list.
You can review deprecated APIs by using the toggle button next to each method from the list. Avoid using deprecated APIs. Deprecated APIs might become unresponsive in the future and cause unexpected failures in your automation scripts.
Detailed information about the method appears.
7. If a section about parameter details appears for the selected method, enter a method parameter value in the **Value** text box.
8. Optional: To invoke the method against the live environment, click **Execute**.
a) If a warning dialog box appears, click **Yes**.
The result for the invoked method appears in the response box.
9. Optional: To copy the result for the invoked method to your clipboard, click **Copy Response**.
10. Optional: To download the result for the invoked method, click **Download**.

What Is vSphere Code Capture

Code Capture records user actions and translates them into executable code.

How Do You Record Actions by Using vSphere Code Capture

To use Code Capture to record a session, you must first enable Code Capture.

Note: Calls made on operations regarding roles, privileges, tags, content libraries, and storage policies are not recorded. Sensitive data such as passwords is not recorded.

1. From the home sidebar menu, click **Developer Center** and go to the **Code Capture** tab.
2. Optional: If Code Capture is not enabled, click the toggle to enable Code Capture.
3. To start a recording, navigate to your desired pane and click the red record button in the top pane. To start recording immediately, click **Start Recording**.
While a recording is in progress, the red record button in the top pane blinks.
4. Optional: To clear the code captured in a previous session and start a new session, click **Clear and Start Another**.
5. To stop a recording, click the red record button in the top pane, or navigate to the **Code Capture** tab in the Developer Center and click **Stop Recording**.
The recorded code appears in the code pane.
6. Optional: Click **Copy** to copy the code or **Download** to download it as a PowerCLI script.
7. To clear the current code and start another recording, click **Clear and Start Another** or navigate to your desired pane and click the red record button in the top pane.

The recorded code appears in the code pane. You can copy the code, download it, or clear the code to start another recording.

Automating Management Tasks Using VMware Aria Automation Orchestrator

VMware Aria Automation Orchestrator is a platform that provides a library of extensible workflows. By using the workflow library, you can automate and configure processes to manage the vSphere infrastructure, other VMware technologies, and third-party technologies.

VMware Aria Automation Orchestrator exposes every operation in the vCenter Server API so that you can integrate all these operations into your own automated processes.

What Are VMware Aria Automation Orchestrator Workflows

A workflow is a series of actions and decisions that are automated to run sequentially after you run the workflow. VMware Aria Automation Orchestrator provides a library of workflows that perform common management tasks.

Basics of Workflows

Workflows consist of a schema, variables, and input and output parameters. The workflow schema is the main component of a workflow as it defines all the workflow elements and the logical task flow of the workflow. The workflow variables and parameters are used by workflows to transfer data. VMware Aria Automation Orchestrator saves a workflow token every time a workflow runs, recording the details of that specific run of the workflow. This token contains all parameters related to the workflow run. For example, if you run a workflow three times, three workflow tokens are saved.

With the vSphere Client, you can run and schedule workflows on selected objects from the vSphere inventory. You cannot create, delete, edit, and manage workflows in the vSphere Client. You develop and manage workflows in the VMware Aria Automation Orchestrator client. For more information about the VMware Aria Automation Orchestrator client, see .

Input Workflow Parameters

To run, most workflows require a certain set of input parameters. The workflow processes input parameters that the user, an application, another workflow, or an action passes to it.

For example, if a workflow resets a virtual machine, the workflow requires the name of the virtual machine as an input parameter.

Output Workflow Parameters

The workflow output parameters represent the result from the workflow run. Some workflows and workflow elements can change the output parameters of the workflow when they run. While they run, workflows can receive the output parameters of other workflows as input parameters.

For example, if a workflow creates a snapshot of a virtual machine, the output parameter for the workflow is the resulting snapshot.

Workflow Presentation

When you run a workflow in the vSphere Client, the client loads the workflow presentation. You provide the input parameters of the workflow in the workflow presentation.

Waiting for Input

Some workflows require user input during their run and the run is suspended either until the user provides the required information, or until the workflow run times out.

What Administration Tasks Can You Perform on vSphere Objects by Using VMware Aria Automation Orchestrator

By using VMware Aria Automation Orchestrator view in the vSphere Client, you can perform administration tasks such as running and scheduling workflows, and viewing the list of available workflows.

From the VMware Aria Automation Orchestrator view in the vSphere Client, you can work with workflows. Working with workflows includes the following tasks:

- Associating workflows with specific vSphere inventory objects such as virtual machines, ESXi hosts, clusters, resource pools, and folders.
- Exporting and importing existing associations of workflows with vSphere inventory objects for backup purposes or to import them to another vCenter Server instance.
- Editing associations of workflows with vSphere inventory objects such as virtual machines, ESXi hosts, clusters, resource pools, folders, and others.
- Viewing information about workflow runs and about workflows waiting for user intervention.
- Running and scheduling workflows on vSphere objects.

Managing Context Actions of Workflows on Your vSphere Inventory Objects

By using VMware Aria Automation Orchestrator, you can associate workflows with the different vSphere object types to see more workflows displayed in the context menu when you right-click a vSphere inventory object. You can also run these workflows on more object types.

You can add and edit context actions, and export and import configuration files containing the context actions of workflows on vSphere objects.

Workflows associated with inventory object types are listed in the context menu that appears when you right-click the inventory objects and in the **Actions** menu.

You can assign Automation Orchestrator groups. The users in the Automation Orchestrator groups and users from the VMware Aria Automation Orchestrator Administrator group in the vSphere Client can use the defined context actions on the vSphere inventory objects.

Assign an Automation Orchestrator Group in the vSphere Client

With VMware Aria Automation Orchestrator, you can assign an Automation Orchestrator group. The users in the group can use the defined context actions on the vSphere inventory objects.

- Verify that groups are created in the VMware Aria Automation Orchestrator client. For more information about Automation Orchestrator groups and how to create them, see the documentation.
 - Log in to the vSphere Client as a member of the VMware Aria Automation Orchestrator Administrators group. For more information about VMware Aria Automation Orchestrator roles and groups, see the documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click the **Group Assignments** tab.
A list of all groups which can use context actions is displayed.
 3. Click **Assign**.
 4. Select a group from the list and click **Assign**.
The group appears in the list of assigned groups.

Associate Your Workflows with vSphere Inventory Object Types

You can associate VMware Aria Automation Orchestrator workflows with a vSphere object type, such as host, to run the workflows directly on the inventory objects of that type.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
- Log in to the vSphere Client as a member of the VMware Aria Automation Orchestrator Administrators group. For more information about VMware Aria Automation Orchestrator roles and groups, see the documentation.

Workflows associated with inventory object types are listed in the context menu that appears when you right-click an inventory object, and in the **Actions** menu.

1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
2. Click the **Context Actions** tab.
3. Click **Add** to add a workflow.
4. Navigate through the workflow library to find the workflow to add, and select the workflow.
5. Under Type of Object, select the vSphere object types with which you want to associate the workflow.
6. Optional: Use the add or remove icons to add or remove object types.
To remove an object type, more than one object types must be selected.
7. Optional: Select the **Allow multiple selection** check box.
With multi-selection enabled, you can select multiple vSphere objects of the same type when you run the workflow.
8. Click **Assign**.
The workflow appears in the list of workflows.

Edit the Context Actions of Your Workflows on vSphere Objects

After you associate a VMware Aria Automation Orchestrator workflow with objects from the vSphere inventory, you can edit the context actions of workflows on the objects from the vSphere inventory.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
 - Log in to the vSphere Client as a member of the VMware Aria Automation Orchestrator Administrators group. For more information about VMware Aria Automation Orchestrator roles and groups, see the documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click the **Context Actions** tab.
 3. Click on the workflow you want to edit.
A dialog box appears.
 4. Change the context action properties.
 5. Click **Save**.

Export the Context Actions of Your Workflows on vSphere Objects

You can transfer the context actions of VMware Aria Automation Orchestrator workflows on objects in the vSphere inventory by using a JSON file.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
 - Log in to the vSphere Client as a member of the VMware Aria Automation Orchestrator Administrators group. For more information about VMware Aria Automation Orchestrator roles and groups, see the documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click the **Context Actions** tab.
 3. Click **MORE > Export**.
 4. Select a location where you want to save the JSON file, and click **Save**.

Import Context Actions of Your Workflows on vSphere Objects

You can import a JSON file that contains the context actions of VMware Aria Automation Orchestrator workflows on objects in the vSphere inventory.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.

- Log in to the vSphere Client as a member of the VMware Aria Automation Orchestrator Administrators group. For more information about VMware Aria Automation Orchestrator roles and groups, see the documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click the **Context Actions** tab.
 3. Click **MORE > Import**.
 4. Browse to select the JSON file to import, and click **Open**.

VMware Aria Automation Orchestrator compares the two sets of workflow context actions and imports the missing workflow context actions.

Working with VMware Aria Automation Orchestrator Workflows in the vSphere Client

In the vSphere Client, you can view information about VMware Aria Automation Orchestrator workflows, run, and schedule workflows.

You can perform some scheduling and running tasks on the VMware Aria Automation Orchestrator workflows from the VMware Aria Automation Orchestrator view in the vSphere Client.

You can schedule a workflow to run at a specified time or start a workflow directly, by right-clicking a vSphere inventory object and selecting **VMware Aria Automation Orchestrator**.

Workflow tasks include:

- Running workflows on vSphere inventory objects, such as virtual machines, ESXi hosts, clusters, resource pools, and folders.
- Viewing information about workflow runs.
- Viewing information about workflows waiting for user interaction.
- Searching for a specific workflow from the list of available workflows.
- Scheduling workflows.

Run Your Workflows on vSphere Inventory Objects

You can automate management tasks in vSphere by running VMware Aria Automation Orchestrator workflows directly on objects from the vSphere inventory.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
 - Verify that you have workflows associated with the vSphere inventory objects. See [Associate Your Workflows with vSphere Inventory Object Types](#).
1. In the vSphere Client home page, select **Global Inventory Lists**.
 2. From the left navigation pane, select an inventory category.
 3. Right-click the object that you want to run the workflow on, and click **VMware Aria Automation Orchestrator > Run Workflow**.
A dialog box appears.

4. From the drop-down menu, select the workflow that you want to run.

Note: If you cannot find the expected workflows, you might have to associate them to the specified vSphere inventory object.

5. Provide the required workflow parameters.
6. Click **Run**.

Schedule Your Workflows on vSphere Inventory Objects

You can automate management tasks in vSphere by scheduling VMware Aria Automation Orchestrator workflows directly on objects from the vSphere inventory.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
- Verify that you have workflows associated with the vSphere inventory objects. See [Associate Your Workflows with vSphere Inventory Object Types](#).

1. In the vSphere Client home page, select **Global Inventory Lists**.
2. From the left navigation pane, select an inventory category.
3. Right-click the object that you want to run the workflow on, and click **VMware Aria Automation Orchestrator > Schedule Workflow**.

A dialog box appears.

4. In the **Task name** text box, enter the name of the scheduled task.
5. Optional: In the **Description** text box, enter a description of the scheduled task.
6. Schedule the date and time of the workflow run.
7. Optional: Specify the end date.
8. Optional: From the **Schedule** drop-down menu, select a recurrence option.
9. Click **Next**.
10. From the drop-down menu, select the workflow that you want to run.

Note: If you cannot find the expected workflows, you might have to associate them to the specified vSphere inventory object.

11. Provide the required workflow parameters.
12. Click **Finish**.

Scheduling Workflows from the VMware Aria Automation Orchestrator view in the vSphere Client

From the VMware Aria Automation Orchestrator view in the vSphere Client, you can create tasks to schedule workflows, edit scheduled tasks, suspend scheduled tasks, and resume suspended scheduled tasks.

Schedule a VMware Aria Automation Orchestrator Workflow

In the vSphere Client, you can schedule a VMware Aria Automation Orchestrator workflow to run at a specified time. You can also set the recurrence for a scheduled workflow.

1. Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
2. Click **Workflows**.
3. Select the workflow that you want to schedule and click **Schedule**.
4. In the **Task name** text box, enter the name of the scheduled task.
5. Optional: In the **Description** text box, enter a description of the scheduled task.
6. Schedule the date and time of the workflow run.
7. Optional: Specify the end date.
8. Optional: From the **Schedule** drop-down menu, select a recurrence option.
9. Click **Next**.
10. Provide the required workflow parameters.
11. Click **Finish**.

Edit the Schedule of Your VMware Aria Automation Orchestrator Workflow

In the vSphere Client, you can modify the schedule of a VMware Aria Automation Orchestrator workflow and set it to run at an earlier or later time.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click **Scheduled workflows**.
A list of the scheduled workflows appears.

3. Select the workflow whose schedule you want to edit and click **Edit**.
4. Modify the parameters of the scheduled workflow and click **Finish**.

Suspend a Scheduled VMware Aria Automation Orchestrator Workflow Run

In the vSphere Client, you can suspend a VMware Aria Automation Orchestrator scheduled task for a workflow run. You can also resume suspended scheduled tasks.

1. Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
2. Click **Scheduled workflows**.
A list of the scheduled workflows appears.
3. Select a workflow and click **Pause**.
The workflow schedule is suspended.

The state of the scheduled task changes to Suspended.

Resume a Suspended Scheduled VMware Aria Automation Orchestrator Workflow Run

In the vSphere Client, you can resume a scheduled VMware Aria Automation Orchestrator workflow run that has been suspended.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click **Scheduled workflows**.
A list of the scheduled workflows appears.
 3. Select a workflow and click **Resume**.
The workflow schedule is resumed from suspension.

The state of the scheduled task changes to Pending.

View Information About Your VMware Aria Automation Orchestrator Workflow Runs

In the vSphere Client, you can view information about the workflow runs for the connected VMware Aria Automation Orchestrator server. The available information includes the workflow name, start and end date, state of the workflow, and the user who started the workflow.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.

- Run a VMware Aria Automation Orchestrator workflow.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click **Workflow Runs**.
The list of workflow runs appears.

You can review the list of workflow runs, cancel a running workflow, delete a workflow run, or respond to a workflow that requires interaction.

View Information About the Runs of a Specific Workflow

In the vSphere Client, you can view information about the runs of a single VMware Aria Automation Orchestrator workflow such as start and end date, state of the workflow, and user who has started the workflow.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
 - Run the specific VMware Aria Automation Orchestrator workflow at least once.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click **Workflows**.
A list of the available workflows appears.
 3. Click the name of a workflow, and click the **All Runs** tab.
A list of workflow runs appears.

You can review the list of workflow runs, cancel a running workflow, or respond to a workflow that requires interaction.

View Workflows That Are Waiting for User Input

In the vSphere Client, you can view the VMware Aria Automation Orchestrator workflows that are waiting for user input.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click **Waiting for input**.

A list of workflows that are waiting for user input appears.

You can provide values for the required parameters of workflows that are waiting for user input. For example, the user might have to provide an approval or further details about the workflow.

Find a VMware Aria Automation Orchestrator Workflow

You can browse for workflows in the inventory of the VMware Aria Automation Orchestrator server or filter the available workflows by a search keyword to find a particular workflow. In vSphere environments with many available workflows, filtering provides a quick way to navigate through the list of workflows.

- Verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.
1. In the object navigator, click **vRealize Orchestrator**.
Depending on your environment and the integration between vCenter Server and VMware Aria Automation Orchestrator, the name of the interface control might be **VMware Aria Automation Orchestrator**.
 2. Click **Workflows**.
 3. In the filter text box, enter a search term or the name of the workflow that you want to find.
You can see a list with the workflows that contain the search term in the workflow name or description.

Workflows for Managing Your vSphere Inventory Objects

The vCenter Server plug-in workflow library contains VMware Aria Automation Orchestrator workflows that you can use to run automated processes related to the vCenter Server and host management.

To access workflows in the vSphere Client, verify that you have configured the VMware Aria Automation Orchestrator plug-in for the vSphere Client. For information about configuring the VMware Aria Automation Orchestrator plug-in for the vSphere Client, see the *Using vCenter Orchestrator Plug-Ins* documentation.

The default workflows for managing vSphere inventory objects are the workflows included in the vCenter Server plug-in workflow library.

Note: Only a predefined set of vCenter Server workflows are available by default in the context menu. You can associate additional workflows with each vSphere object. See [Associate Your Workflows with vSphere Inventory Object Types](#).

For more information about the workflows available in the vCenter Server plug-in workflow library, see *vCenter Plug-In Workflow Library* in the *Using VMware Aria Automation Orchestrator Plug-Ins* documentation.

Working with Headless Systems

ESXi supports the detection and configuration of headless systems.

A headless system is a system that can be operated without a monitor, keyboard, or mouse. Network Appliance boxes do not have Video Graphics Array (VGA), the primary interface is a single serial port. You can set up your existing headless systems to use ESXi. You can add ESXi appliances to a data center where virtual machines are managed with vCenter Server. All existing ESXi features can be used with a headless system that is configured with either embedded flash or minimal local storage. ESXi allows for dynamic switching between different serial modes, which is useful for diagnosing and debugging problems. You can switch between modes to view or modify system parameters.

How to Detect a Headless System

ESXi automatically detects headless systems.

ESXi automatically redirects the DCUI over a serial port connection to improve headless detection. When ESXi automatically detects a headless system, ESXi will set up the serial port as COM1, 115200 baud, and redirects the DCUI over this serial port. The specific settings of com port and baud rate are read from the Serial Port Console Redirection (SPCR) table, if it exists. This behavior can be deactivated using new boot parameters if the default settings are not acceptable. You can set the `headless` flag in the ACPI FADT table to mark a system as headless.

Serial Mode Dynamic Switching

ESXi supports dynamic switching between four different serial port modes.

ESXi supports serial mode dynamic switching to provide maximum platform flexibility, and to allow debugging and supportability in the text box. ESXi examines the input characters for any serial port mode and switches the modes based

on the input key sequence. DCUI, Shell, GDB, and Logging modes are supported. If you have two serial ports, only one of the four modes is allowed on each port. Two serial ports cannot be in the same mode. If you attempt a dynamic switch to a mode in use by the other serial port, the request is ignored. Dynamic switching eliminates the need to interrupt the boot process manually or to create a custom image to redirect to a serial port. It also addresses supportability issues regarding headless systems that only have one serial port, by making it possible to switch the serial port between different modes of operation.

What Serial Port Modes Does ESXi Support

ESXi supports four serial port modes — logging mode, GDB mode, Shell mode, and DCUI mode.

There are four serial port modes in ESXi:

Logging mode – Logging mode is the default mode in a debug build. Logging mode sends the `vmkernel.log` over the serial port.

GDB mode – Use GDB mode for dedicated debugging.

Shell mode – Shell mode is the shell port access, which is similar to SSH.

DCUI mode – DCUI mode is a Direct Console User Interface. This is the user interface that is displayed when you boot ESXi using a monitor.

Note:

Only COM1 and COM2 ports are supported. USB serial or PCI serial cards are not supported.

ESXi Keystrokes for Dynamic Serial Mode Switching

ESXi includes a unique keystroke sequence that allows dynamic serial mode switching.

Dynamic Switching Keystrokes

Once the correct keystroke sequence is entered, the system switches the serial port to the desired mode.

Logging mode: Ctrl+G, Ctrl+B, 1

Shell mode: Ctrl+G, Ctrl+B, 2

DCUI mode: Ctrl+G, Ctrl+B, 3

GDB mode: Ctrl+G, Ctrl+B, ?

Note: Once in GDB mode, you cannot switch modes again using a key sequence. You must use the CLI to switch modes.

How Do You Switch Serial Modes Through the ESXi CLI

You can switch serial modes using the CLI.

Dynamic Switching Using the CLI

Use `esxcfg-advcfg` to set the current mode to **none**. Then set the new desired mode using the CLI.

Logging mode: `esxcfg-advcfg -s com1 /Misc/LogPort`

Shell mode: `esxcfg-advcfg -s com1 /Misc/ShellPort`

DCUI mode: `esxcfg-advcfg -s com1 /Misc/ConsolePort`

GDB mode: `esxcfg-advcfg -s com1 /Misc/GDBPort`

Example

If the serial mode is set to logging mode, enter these two commands to switch it to DCUI mode.

```
$ . > esxcfg-advcfg -s none /Misc/LogPort
$ . > esxcfg-advcfg -s com1 /Misc/ConsolePort
```

What Are the Keystrokes for Controlling the Serial DCUI

You can use alternate keystrokes to control the DCUI over a serial port. These alternate are useful when F2 or other function keys cannot be used.

Controlling the Serial DCUI

Alternate keystroke mappings for DCUI mode:

- Esc + 1 -> F1
- Esc + 2 -> F2
- Esc + 3 -> F3
- Esc + 4 -> F4
- Esc + 5 -> F5
- Esc + 6 -> F6
- Esc + 7 -> F7
- Esc + 8 -> F8
- Esc + 9 -> F9
- Esc + 0 -> F10
- Esc + ! -> F11
- Esc + @ -> F12

Troubleshooting Overview

vSphere Troubleshooting contains common troubleshooting scenarios and provides solutions for each of these problems. You can also find guidance here for resolving problems that have similar origins. For unique problems, consider developing and adopting a troubleshooting methodology.

The following approach for effective troubleshooting elaborates on how to gather troubleshooting information, such as identifying symptoms and defining the problem space. Troubleshooting with log files is also discussed.

Guidelines for Troubleshooting Your Implementation of vSphere

To troubleshoot your implementation of vSphere, identify the symptoms of the problem, determine which of the components are affected, and test possible solutions.

Identifying Symptoms

A number of potential causes might lead to the under-performance or nonperformance of your implementation. The first step in efficient troubleshooting is to identify exactly what is going wrong. After you have isolated the symptoms of the problem, you must define the problem space. Identify the software or hardware components that are affected and might be causing the problem and those components that are not involved.

Defining the Problem Space

Testing Possible Solutions

When you know what the symptoms of the problem are and which components are involved, test the solutions systematically until the problem is resolved.

How Do You Identify the Symptoms

Before you attempt to resolve a problem in your vSphere implementation, you must identify precisely how it is failing.

The first step in the troubleshooting process is to gather information that defines the specific symptoms of what is happening. You might ask these questions when gathering this information:

- What is the task or expected behavior that is not occurring?
- Can the affected task be divided into subtasks that you can evaluate separately?
- Is the task ending in an error? Is an error message associated with it?
- Is the task completing but in an unacceptably long time?
- Is the failure consistent or sporadic?
- What has changed recently in the software or hardware that might be related to the failure?

How Do You Define the Problem Space

After you identify the symptoms of the problem in your vSphere implementation, determine which components in your setup are affected, which components might be causing the problem, and which components are not involved.

To define the problem space in an implementation of vSphere, be aware of the components present. In addition to VMware software, consider third-party software in use and which hardware is being used with the VMware virtual hardware.

Recognizing the characteristics of the software and hardware elements and how they can impact the problem, you can explore general problems that might be causing the symptoms.

- Misconfiguration of software settings
- Failure of physical hardware
- Incompatibility of components

Break down the process and consider each piece and the likelihood of its involvement separately. For example, a case that is related to a virtual disk on local storage is probably unrelated to third-party router configuration. However, a local disk controller setting might be contributing to the problem. If a component is unrelated to the specific symptoms, you can probably eliminate it as a candidate for solution testing.

Think about what changed in the configuration recently before the problems started. Look for what is common in the problem. If several problems started at the same time, you can probably trace all the problems to the same cause.

How Do You Test Possible Solutions

After you know the symptoms of the problem in your vSphere implementation and which software or hardware components are most likely involved, you can systematically test solutions until you resolve the problem.

With the information that you have gained about the symptoms and affected components, you can design tests for pinpointing and resolving the problem. These tips might make this process more effective.

- Generate ideas for as many potential solutions as you can.
- Verify that each solution determines unequivocally whether the problem is fixed. Test each potential solution but move on promptly if the fix does not resolve the problem.
- Develop and pursue a hierarchy of potential solutions based on likelihood. Systematically eliminate each potential problem from the most likely to the least likely until the symptoms disappear.
- When testing potential solutions, change only one thing at a time. If your setup works after many things are changed at once, you might not be able to discern which of those things made a difference.
- If the changes that you made for a solution do not help resolve the problem, return the implementation to its previous status. If you do not return the implementation to its previous status, new errors might be introduced.

- Find a similar implementation that is working and test it in parallel with the implementation that is not working properly. Make changes on both systems at the same time until few differences or only one difference remains between them.

Troubleshooting with vCenter Server Logs

You can often obtain valuable troubleshooting information by looking at the logs provided by the various services and agents that your vSphere implementation is using.

Most logs are located in `/var/log/vmware/<service_name>` for vCenter Server deployments.

Common Logs

The following logs are common to all vCenter Server deployments.

Table 151: Common Log Directories

Log Directory	Description
<code>../firstboot</code>	Stores first boot logs
<code>applmgmt</code> and <code>applmgmt-audit</code>	Stores logs related to the VMware Appliance Management Service
<code>cloudvm</code>	Stores logs for allotment and distribution of resources between services
<code>rhttpproxy</code>	Stores logs for the VMware HTTP Reverse Proxy service
<code>sca</code>	Stores logs for the VMware Service Control Agent service
<code>vapi</code>	Stores logs for the VMware vAPI Endpoint service
<code>vmafdd</code>	Stores logs for the VMware Authentication Framework - LDAP service
<code>vmdird</code>	Stores logs for the VMware Directory Service - LDAP service
<code>vmon</code>	Stores logs for the VMware Service Lifecycle Manager service

Management Node Logs

The following logs are available if a management node deployment is selected.

Table 152: Management Node Log Directories

Log Directory	Service
<code>rbd</code>	VMware vSphere Auto Deploy
<code>content-library</code>	VMware Content Library Service
<code>eam</code>	VMware ESX Agent Manager
<code>netdumper</code>	VMware vSphere ESXi Dump Collector
<code>perfcharts</code>	VMware Performance Charts Service
<code>vmcam</code>	VMware vSphere Authentication Proxy
<code>vmdird</code>	VMware Directory Service - LDAP
<code>vmware-sps</code>	VMware vSphere Profile-Driven Storage Service
<code>vpzd</code>	VMware vCenter Server
<code>vpostgres</code>	VMware Postgres service

Log Directory	Service
vcha	VMware vCenter High Availability service

Troubleshooting vCenter Server and ESXi Host Certificates

Certificates are automatically generated when you install vCenter Server. These default certificates are not signed by a commercial certificate authority (CA) and might not provide strong security. You can replace default vCenter Server certificates with certificates signed by a commercial CA. When you replace vCenter Server and ESXi certificates, you might encounter errors.

New vCenter Server Certificate Does Not Appear to Load

After you replace default vCenter Server certificates, the new certificates might not appear to load.

When you install new vCenter Server certificates, you might not see the new certificate.

Existing open connections to vCenter Server are not forcibly closed and might still use the old certificate.

To force all connections to use the new certificate, use one of the following methods.

- Restart the network stack or network interfaces on the server.
- Restart the vCenter Serverservice.

vCenter Server Cannot Connect to Managed Hosts

After you replace default vCenter Server certificates and restart the system, vCenter Server might not be able to connect to managed hosts.

vCenter Server cannot connect to managed hosts after server certificates are replaced and the system is restarted.

Log into the host as the root user and reconnect the host to vCenter Server.

Troubleshooting ESXi Hosts

The host troubleshooting topics provide solutions to potential problems that you might encounter when using your vCenter Server and ESXi hosts.

Troubleshooting vSphere HA Host States

vCenter Server reports vSphere HA host states that indicate an error condition on the host. Such errors can prevent vSphere HA from fully protecting the virtual machines on the host and can impede vSphere HA's ability to restart virtual machines after a failure. Errors can occur when vSphere HA is being configured or unconfigured on a host or, more rarely, during normal operation. When this happens, you should determine how to resolve the error, so that vSphere HA is fully operational.

vSphere HA Agent Is in the Agent Unreachable State

The vSphere HA agent on a host is in the Agent Unreachable state for a minute or more. User intervention might be required to resolve this situation.

vSphere HA reports that an agent is in the Agent Unreachable state when the agent for the host cannot be contacted by the primary host or by vCenter Server. Consequently, vSphere HA is not able to monitor the virtual machines on the host and might not restart them after a failure.

A vSphere HA agent can be in the Agent Unreachable state for several reasons. This condition most often indicates that a networking problem is preventing vCenter Server or the primary host from contacting the agent on the host, or that all hosts in the cluster have failed. This condition can also indicate the unlikely situation that vSphere HA was deactivated

and then reactivated on the cluster while vCenter Server could not communicate with the vSphere HA agent on the host, or that the ESXi host agent on the host has failed, and the watchdog process was unable to restart it. In any of these cases, a failover event is not triggered when a host goes into the Unreachable state.

Determine if vCenter Server is reporting the host as not responding. If so, there is a networking problem, an ESXi host agent failure, or a total cluster failure. After the condition is resolved, vSphere HA should work correctly. If not, reconfigure vSphere HA on the host. Similarly, if vCenter Server reports the hosts are responding but a host's state is Agent Unreachable, reconfigure vSphere HA on that host.

vSphere HA Agent is in the Uninitialized State

The vSphere HA agent on a host is in the Uninitialized state for a minute or more. User intervention might be required to resolve this situation.

vSphere HA reports that an agent is in the Uninitialized state when the agent for the host is unable to enter the run state and become the primary host or to connect to the primary host. Consequently, vSphere HA is not able to monitor the virtual machines on the host and might not restart them after a failure.

A vSphere HA agent can be in the Uninitialized state for one or more reasons. This condition most often indicates that the host does not have access to any datastores. Less frequently, this condition indicates that the host does not have access to its local datastore on which vSphere HA caches state information, the agent on the host is inaccessible, or the vSphere HA agent is unable to open required firewall ports. It is also possible that the ESXi host agent has stopped.

Search the list of the host's events for recent occurrences of the event `vSphere HA Agent for the host has an error`. This event indicates the reason for the host being in the uninitialized state. If the condition exists because of a datastore problem, resolve whatever is preventing the host from accessing the affected datastores. If the ESXi host agent has stopped, you must restart it. After the problem has been resolved, if the agent does not return to an operational state, reconfigure vSphere HA on the host.

Note: If the condition exists because of a firewall problem, check if there is another service on the host that is using port 8182. If so, shut down that service, and reconfigure vSphere HA.

vSphere HA Agent is in the Initialization Error State

The vSphere HA agent on a host is in the Initialization Error state for a minute or more. User intervention is required to resolve this situation.

vSphere HA reports that an agent is in the Initialization Error state when the last attempt to configure vSphere HA for the host failed. vSphere HA does not monitor the virtual machines on such a host and might not restart them after a failure.

This condition most often indicates that vCenter Server was unable to connect to the host while the vSphere HA agent was being installed or configured on the host. This condition might also indicate that the installation and configuration completed, but the agent did not become a primary host or a secondary host within a timeout period. Less frequently, the condition is an indication that there is insufficient disk space on the host's local datastore to install the agent, or that there are insufficient unreserved memory resources on the host for the agent resource pool. Finally, for ESXi 5.x hosts, the configuration fails if a previous installation of another component required a host reboot, but the reboot has not yet occurred.

When a Configure HA task fails, a reason for the failure is reported.

Reason for Failure	Action
Host communication errors	Resolve any communication problems with the host and retry the configuration operation.

Reason for Failure	Action
Timeout errors	Possible causes include that the host crashed during the configuration task, the agent failed to start after being installed, or the agent was unable to initialize itself after starting up. Verify that vCenter Server is able to communicate with the host. If so, see vSphere HA Agent Is in the Agent Unreachable State or vSphere HA Agent is in the Uninitialized State for possible solutions.
Lack of resources	Free up approximately 75MB of disk space. If the failure is due to insufficient unreserved memory, free up memory on the host by either relocating virtual machines to another host or reducing their reservations. In either case, retry the vSphere HA configuration task after resolving the problem.
Reboot pending	If an installation for a 5.0 or later host fails because a reboot is pending, reboot the host and retry the vSphere HA configuration task.

vSphere HA Agent is in the Uninitialization Error State

The vSphere HA agent on a host is in the Uninitialization Error state. User intervention is required to resolve this situation.

vSphere HA reports that an agent is in the Uninitialization Error state when vCenter Server is unable to unconfigure the agent on the host during the Unconfigure HA task. An agent left in this state can interfere with the operation of the cluster. For example, the agent on the host might elect itself as primary host and lock a datastore. Locking a datastore prevents the valid cluster primary host from managing the virtual machines with configuration files on that datastore.

This condition usually indicates that vCenter Server lost the connection to the host while the agent was being unconfigured.

Add the host back to vCenter Server (version 5.0 or later). The host can be added as a stand-alone host or added to any cluster.

vSphere HA Agent is in the Host Failed State

The vSphere HA agent on a host is in the Host Failed state. User intervention is required to resolve the situation.

Usually, such reports indicate that a host has actually failed, but failure reports can sometimes be incorrect. A failed host reduces the available capacity in the cluster and, in the case of an incorrect report, prevents vSphere HA from protecting the virtual machines running on the host.

This host state is reported when the vSphere HA primary host to which vCenter Server is connected is unable to communicate with the host and with the heartbeat datastores that are in use for the host. Any storage failure that makes the datastores inaccessible to hosts can cause this condition if accompanied by a network failure.

Check for the noted failure conditions and resolve any that are found.

vSphere HA Agent is in the Network Partitioned State

The vSphere HA agent on a host is in the Network Partitioned state. User intervention might be required to resolve this situation.

While the virtual machines running on the host continue to be monitored by the primary hosts that are responsible for them, vSphere HA's ability to restart the virtual machines after a failure is affected. First, each primary host has access to a subset of the hosts, so less failover capacity is available to each host. Second, vSphere HA might be unable to restart a FT Secondary VM after a failure. See also *vSphere Availability* troubleshooting.

A host is reported as partitioned if both of the following conditions are met:

- The vSphere HA primary host to which vCenter Server is connected is unable to communicate with the host by using the management (or VMware vSAN™) network, but is able to communicate with that host by using the heartbeat datastores that have been selected for it.
- The host is not isolated.

A network partition can occur for a number of reasons including incorrect VLAN tagging, the failure of a physical NIC or switch, configuring a cluster with some hosts that use only IPv4 and others that use only IPv6, or the management networks for some hosts were moved to a different virtual switch without first putting the host into maintenance mode.

Resolve the networking problem that prevents the hosts from communicating by using the management networks.

vSphere HA Agent is in the Network Isolated State

The vSphere HA agent on a host is in the Network Isolated state. User intervention is required to resolve this situation.

When a host is in the Network Isolated state, there are two things to consider -- the isolated host and the vSphere HA agent that holds the primary role.

- On the isolated host, the vSphere HA agent applies the configured isolation response to the running VMs, determining if they should be shut down or powered off. It does this after checking whether a primary agent is able to take responsibility for each VM (by locking the VM's home datastore.) If not, the agent defers applying the isolation response for the VM and rechecks the datastore state after a short delay.
- If the vSphere HA primary agent can access one or more of the datastores, it monitors the VMs that were running on the host when it became isolated and attempts to restart any that were powered off or shut down.

A host is network isolated if both of the following conditions are met:

- Isolation addresses have been configured and the host is unable to ping them.
- The vSphere HA agent on the host is unable to access any of the agents running on the other cluster hosts.

Note: If your vSphere HA cluster has vSAN enabled, a host is determined to be isolated if it cannot communicate with the other vSphere HA agents in the cluster and cannot reach the configured isolation addresses. Although the vSphere HA agents use the vSAN network for inter-agent communication, the default isolation address is still the gateway of the host. Hence, in the default configuration, both networks must fail for a host be declared isolated.

Resolve the networking problem that is preventing the host from pinging its isolation addresses and communicating with other hosts.

Configuration of vSphere HA on Hosts Times Out

The configuration of a vSphere HA cluster might time out on some of the hosts added to it.

When you enable vSphere HA on an existing cluster with a large number of hosts and virtual machines, the setup of vSphere HA on some of the hosts might fail.

This failure is the result of a time out occurring before the installation of vSphere HA on the host(s) completes.

Set the vCenter Server advanced option `config.vpxd.das.electionWaitTimeSec` to `value=240`. Once this change is made, the time outs do not occur.

Authentication Token Manipulation Error

When you try to create a password that does not meet the authentication requirements of the ESXi host, an error occurs.

When you create a password on the host, the following fault message appears: `A general system error occurred: passwd: Authentication token manipulation error.`

The following message is included: `Failed to set the password. It is possible that your password does not meet the complexity criteria set by the system.`

The host checks for password compliance using the default authentication plug-in, `pam_passwdqc.so`. If the password is not compliant, the error appears.

When you create a password, include a mix of characters from four character classes: lowercase letters, uppercase letters, numbers, and special characters such as an underscore or dash.

Note: An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used.

Your user password must meet the following length requirements.

- Passwords containing characters from three character classes must be at least eight characters long.
- Passwords containing characters from all four character classes must be at least seven characters long.

When the PAM module determines if the password consists of dictionary words, the module may treat some numbers as letters and may reject a password which seems to meet the complexity requirements. For example, a password such as `P4$$w0rd` may be rejected because it is based on a dictionary word.

If the module rejects your password, you can transfer the password to the `/bin/pwqcheck` application on the ESXi Shell, change some numbers or letters, and check if it is accepted.

For more information, see the *vSphere Security* documentation.

Unable to Download VIBs When Using vCenter Server Reverse Proxy

You are unable to download VIBs if vCenter Server is using a custom port for the reverse proxy.

If you configure vCenter Server reverse proxy to use a custom port, the VIB downloads fail.

If vCenter Server is using a custom port for the reverse proxy, the custom port is not automatically enabled in the ESXi firewall and the VIB downloads fail.

1. Open an SSH connection to the host and log in as root.
2. Optional: List the existing firewall rules.

```
esxcli network firewall ruleset list
```

3. Optional: Back up the `/etc/vmware/firewall/service.xml` file.

```
cp /etc/vmware/firewall/service.xml /etc/vmware/firewall/service.xml.bak
```

4. Edit the access permissions of the `service.xml` file to allow writes by running the `chmod` command.

- To allow writes, run `chmod644 /etc/vmware/firewall/service.xml`.
- To toggle the sticky bit flag, run `chmod+t /etc/vmware/firewall/service.xml`.

5. Open the `service.xml` file in a text editor.

6. Add a new rule to the `service.xml` file that enables the custom port for the vCenter Server reverse proxy .

```
<service id='id_value'>
  <id>vcenterhttpreverseproxy</id>
  <rule id='0000'>
    <direction>outbound</direction>
    <protocol>tcp</protocol>
    <port type='dst'>custom_reverse_proxy_port</port>
  </rule>
  <enabled>true</enabled>
  <required>false</required>
```

```
</service>
```

Where `id_value` must be a unique value, for example, if the last listed service in the `service.xml` file has ID 0040, you must enter id number 0041.

7. Revert the access permissions of the `service.xml` file to the default read-only setting.

```
chmod 444 /etc/vmware/firewall/service.xml
```

8. Refresh the firewall rules for the changes to take effect.

```
esxcli network firewall refresh
```

9. Optional: List the updated rule set to confirm the change.

```
esxcli network firewall ruleset list
```

10. Optional: If you want the firewall configuration to persist after a reboot of the ESXi host, copy the `service.xml` onto persistent storage and modify the `local.sh` file.

- a) Copy the modified `service.xml` file onto persistent storage, for example `/store/`, or onto a VMFS volume, for example `/vmfs/volumes/volume/`.

```
cp /etc/vmware/firewall/service.xml location_of_xml_file
```

You can store a VMFS volume in a single location and copy it to multiple hosts.

- b) Add the `service.xml` file information to the `local.sh` file on the host.

```
cp location_of_xml_file /etc/vmware/firewall
esxcli network firewall refresh
```

Where `location_of_xml_file` is the location to which the file was copied.

Troubleshooting Licensing of ESXi Hosts and vCenter Server

The troubleshooting licensing topics provide solutions to problems that you might encounter as a result of an incorrect or incompatible license setup in vSphere.

Troubleshooting ESXi Host Licensing

You might encounter different problems that result from an incompatible or incorrect license configuration of ESXi hosts.

Unable to Assign a License to an ESXi Host

Under certain conditions, you might be unable to assign a license to an ESXi host.

You try to assign a license to an ESXi host, but you cannot perform the operation and you receive an error message.

You might be unable to assign a license to an ESXi host because of the following reasons:

- The calculated license usage for the host exceeds the license capacity. For example, you have a vSphere license key with capacity for two CPUs. You try to assign the key to a host that has four CPUs. You cannot assign the license, because the required license usage for the host is greater than the license capacity.
- The features on the host do not match the license edition. For example, you might configure hosts with vSphere Distributed Switch and vSphere DRS while in evaluation mode. Later, you try to assign vSphere Standard license to the hosts. This operation fails because the vSphere Standard edition does not include vSphere Distributed Switch and vSphere DRS.
- The host is connected to a vCenter Server system that is assigned a license that restricts the edition of the license that you want to assign.
- Assign a license with larger capacity.

- Upgrade the license edition to match the resources and features on the host, or deactivate the features that do not match the license edition.
- Assign a vSphere license whose edition is compatible with the license edition of vCenter Server.

ESXi Host Disconnects from vCenter Server

An ESXi host might disconnect from vCenter Server or all ESXi hosts might disconnect from vCenter Server at the same time.

An ESXi host disconnects from vCenter Server when the host evaluation period or license expires. All ESXi hosts disconnect from vCenter Server when the evaluation period or the license of vCenter Server expire. You receive a licensing-related error message both when a single host disconnects and when all hosts disconnect. You cannot add hosts to the vCenter Server inventory. The hosts and the virtual machines on the hosts continue to run.

- The 60-day evaluation period of the host has expired or the host license has expired.
- The 60-day evaluation period of vCenter Server is expired or the vCenter Server license is expired.
- Assign a vSphere license to the ESXi host and try to reconnect it to vCenter Server.
- Assign a vCenter Server license to the vCenter Server system.

Unable to Power On a Virtual Machine

You try to power on a virtual machine in your vSphere implementation, but the operation is unsuccessful and you receive an error message.

You cannot power on a virtual machine on an ESXi host.

You might be unable to power on a virtual machine because of the following reasons.

- The 60-day evaluation period of the host is expired.
- The license of the host is expired.

Table 153: Power on a Virtual Machine

Cause	Solution
The evaluation period of the host is expired	Assign a vSphere license to the ESXi host
The license of the host is expired	Assign a vSphere license to the ESXi host

Unable to Configure or Use a Feature in Your vSphere Implementation

You cannot use a feature or change its configuration and a licensing-related error message appears.

You cannot use or configure a feature and a licensing-related error message appears.

The ESXi host or the vCenter Server system is assigned a license that does not support the features that you want to configure.

Check the licensed features on the ESXi host and on the vCenter Server system. Upgrade the edition of the license assigned to the host or vCenter Server if they do not include the features that you try to configure or use.

vSphere Virtual Machine Administration

vSphere Virtual Machine Administration describes how to create, configure, and manage virtual machines in the VMware vSphere® environment.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we have updated this guide to remove instances of non-inclusive language.

This guide provides introductions to the tasks that you can perform within the system and also cross-references to the documentation that describes the tasks in detail.

This information focuses on managing virtual machines and includes the following information.

- Creating and deploying virtual machines, templates, and clones
- Deploying OVF templates
- Using content libraries to manage templates and other library items
- Configuring virtual machine hardware and virtual machine options
- Managing multi-tiered applications with VMware vSphere vApp
- Monitoring solutions with the vCenter Solutions Manager
- Managing virtual machines, including using snapshots
- Upgrading virtual machines
- Troubleshooting virtual machines

vSphere Virtual Machine Administration covers VMware ESXi™ and VMware vCenter Server®.

Intended Audience

This information is written for experienced Windows or Linux system administrators who are familiar with virtualization.

Updated Information

This *vSphere Virtual Machine Administration* is updated with each release of the product or when necessary.

This table provides the update history of the .

Revision	Description
10 SEP 2024	Added information about the Virtual Shared Graphics Acceleration (vSGA): • Virtual Shared Graphics Acceleration • Setup Your Environment for vSGA
19 AUG 2024	Updated the table in Required vCenter Server Privileges for Common Tasks .
18 JUL 2024	Updated the Install the VMware Remote Console Application procedure.
25 JUN 2024	Initial release.

Introduction to VMware vSphere Virtual Machines

A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. The virtual machine consists of a set of specification and configuration files and is backed by the physical resources of a host.

Every virtual machine has virtual devices that provide the same functionality as physical hardware but are more portable, more secure, and easier to manage.

Before you start creating and managing virtual machines, you benefit from some background information, for example, the virtual machine files, life cycle, components, and so on.

Virtual Machine Files

A virtual machine consists of several files that are stored on a storage device. The key files are the configuration file, virtual disk file, NVRAM setting file, and log file. You configure virtual machine settings through the vSphere Client, ESXCLI, or the vSphere Web Services SDK.

CAUTION: Do not change, move, or delete virtual machine files without instructions from a VMware Technical Support representative.

Table 154: Virtual Machine Files

File	Usage	Description
.vmx	<i>vmname.vmx</i>	Virtual machine configuration file
.vmxf	<i>vmname.vmx</i> f	Additional virtual machine configuration files
.vmdk	<i>vmname.vmdk</i>	Virtual disk characteristics
-flat.vmdk	<i>vmname-flat.vmdk</i>	Virtual machine data disk
.nvram	<i>vmname.nvram</i> or <i>nvram</i>	Virtual machine BIOS or EFI configuration
.vmem	<i>vmname.vmem</i>	Virtual machine paging backup file
.vmsd	<i>vmname.vmsd</i>	Virtual machine snapshots information (metadata) file
.vmsn	<i>vmname.vmsn</i>	Virtual machine memory snapshot file
.vswp	<i>vmname.vswp</i>	Virtual machine swap file
.vmss	<i>vmname.vmss</i>	Virtual machine suspend file
.log	<i>vmware.log</i>	Current virtual machine log file
-#.log	<i>vmware-#.log</i> (where # is a number starting with 1)	Old virtual machine log files

Additional files are created when you perform certain tasks with the virtual machine.

- A *.hlog* file is a log file that is used by vCenter Server to keep track of virtual machine files that must be removed after a certain operation completes.
- A *.vmtx* file is created when you convert a virtual machine to a template. The *.vmtx* file replaces the virtual machine configuration file (*.vmx* file).

Virtual Machines and the Virtual Infrastructure

The infrastructure that supports virtual machines consists of at least two software layers, virtualization and management. In vSphere, ESXi provides the virtualization capabilities that aggregate and present the host hardware to virtual machines as a normalized set of resources. Virtual machines run on ESXi hosts that vCenter Server manages.

vCenter Server can pool the resources of multiple hosts and lets you effectively monitor and manage your data center infrastructure. You can manage resources for virtual machines, provision virtual machines, schedule tasks, collect statistics logs, create templates, and more. vCenter Server also provides vSphere vMotion™, vSphere Storage vMotion, vSphere Distributed Resource Scheduler (DRS), vSphere High Availability (HA), and vSphere Fault Tolerance. These services enable efficient and automated resource management and high availability for virtual machines.

The vSphere Client is the primary interface for managing vCenter Server, ESXi hosts, and virtual machines. The vSphere Client also provides console access to virtual machines.

In the vCenter Server hierarchy that you see in the vSphere Client, a data center is the top-level container of ESXi hosts, folders, clusters, resource pools, vSphere vApps, virtual machines, and so on.

Datastores are virtual representations of underlying physical storage resources. Datastores hide the idiosyncrasies of the underlying physical storage and present a uniform model for the storage resources required by virtual machines. A datastore is the storage location (for example, a physical disk or LUN on a RAID, or a SAN) for virtual machine files.

Virtual Machine Lifecycle

You have many options for creating and deploying virtual machines. You can create a single virtual machine and install a guest operating system and VMware Tools on it. You can clone an existing virtual machine or convert it to a template. You can also deploy OVF or OVA templates.

The vSphere Client **New Virtual Machine** wizard and the **Edit Settings** dialog box let you add, configure, or remove most of the virtual machine's hardware, options, and resources. You monitor CPU, memory, disk, network, and storage metrics through the performance charts in the vSphere Client. Snapshots let you capture the state of the virtual machine, including the virtual machine memory, settings, and virtual disks. You can roll back to the previous virtual machine state when needed.

With vSphere vApps, you can manage multi-tiered applications. You use vSphere Lifecycle Manager to perform orchestrated upgrades to upgrade the virtual hardware and VMware Tools of virtual machines in the inventory at the same time.

When a virtual machine is no longer needed, you can remove it from the inventory without deleting it from the datastore, or you can delete the virtual machine and all its files.

Virtual Machine Components

Virtual machines typically have an operating system, VMware Tools, and virtual resources and hardware. You manage these components like you manage the components of a physical computer.

Operating System

You install a guest operating system on a virtual machine the same way you install an operating system on a physical computer. You must have a CD/DVD-ROM or ISO image containing the installation files from an operating system vendor.

After installation, you are responsible for securing and patching the operating system.

VMware Tools

VMware Tools is a suite of utilities that enhances the performance of the virtual machine guest operating system and improves management of the virtual machine. It includes device drivers and other software that is essential for your VM. With VMware Tools, you have more control over the virtual machine interface.

Compatibility Setting

In the vSphere Client, you assign each virtual machine to a compatible ESXi host version, cluster, or data center by applying a compatibility setting. The compatibility setting determines which ESXi host versions the virtual machine can run on and the hardware features available to the virtual machine.

Hardware Devices

Each virtual hardware device performs the same function for the virtual machine as hardware on a physical computer does. Every virtual machine has CPU, memory, and disk resources. CPU virtualization emphasizes performance and runs directly on the processor whenever possible. The underlying physical resources are used whenever possible. The virtualization layer runs instructions only as needed to make virtual machines operate as if they were running directly on a physical machine.

All recent operating systems provide support for virtual memory, allowing software to use more memory than the machine physically has. Similarly, the ESXi hypervisor provides support for overcommitting virtual machine memory, where the amount of guest memory configured for all virtual machines might be larger than the amount of the host's physical memory.

You access the hardware devices in the **Edit Settings** dialog box. Not all devices are configurable. Some hardware devices are part of the virtual motherboard and appear in the expanded device list of the **Edit Settings** dialog box, but you cannot modify or remove them. For a list of hardware devices and their functions, see [Virtual Machine Hardware Available to vSphere Virtual Machines](#).

In the **Edit Settings** dialog box you can also add virtual hardware devices to the virtual machine. You can use the memory or CPU hotplug options to add memory or CPU resources to a virtual machine while the virtual machine is running. You can deactivate Memory or CPU hotplug to avoid adding memory or CPUs while the virtual machine is running. Memory hotplug is supported on all 64 bit operating systems, but to use the added memory, the guest operating system must also support this feature. See the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>

Virtual Machine Hardware Available to vSphere Virtual Machines

VMware provides devices, resources, profiles, and vServices that you can configure or add to your virtual machine.

Table 155: Virtual Machine Hardware and Descriptions

Hardware Device	Description
CPU	You can configure a virtual machine that runs on an ESXi host to have one or more virtual processors. A virtual machine cannot have more virtual CPUs than the actual number of logical CPUs on the host.
Chipset	The motherboard uses VMware proprietary devices based on the following chips: - Intel 440BX AGPset 82443BX Host Bridge/Controller - Intel 82371AB (PIIX4) PCI ISA IDE Xcelerator - National Semiconductor PC87338 ACPI 1.0 and PC98/99 Compliant SuperI/O - Intel 82093AA I/O Advanced Programmable Interrupt Controller
DVD/CD-ROM Drive	Installed by default when you create a virtual machine. You can configure DVD/CD-ROM devices to connect to client devices, host devices, or datastore ISO files. You can add, remove, or configure DVD/CD-ROM devices.
Hard Disk	Stores the operating system of a virtual machine, program files, and other data associated with its activities. A virtual disk is a large physical file, or a set of files, that can be copied, moved, archived, and backed up as easily as any other file.
IDE 0, IDE 1	By default, two Integrated Drive Electronics (IDE) interfaces are presented to the virtual machine. The IDE interface (controller) is a standard way for storage devices (floppy drives, hard drives, and CD-ROM drives) to connect to the virtual machine.
Keyboard	Provides keyboard input from any virtual machine consoles.
Memory	The virtual hardware memory size determines how much memory applications that are running inside the virtual machine have available to them. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size.

Hardware Device	Description
Network Adapter	ESXi networking features provide communication between virtual machines on the same host, between virtual machines on different hosts, and between other virtual and physical machines. When you configure a virtual machine, you can add network adapters (NICs) and specify the adapter type.
Parallel port	Interface for connecting peripherals to the virtual machine. The virtual parallel port can connect to a file. You can add, remove, or configure virtual parallel ports.
PCI controller	Bus on the virtual machine motherboard that communicates with components such as hard disks and other devices. One PCI controller is presented to the virtual machine. You cannot configure or remove this device.
PCI Device	You can add up to 16 PCI vSphere DirectPath devices to a virtual machine. The devices must be reserved for PCI passthrough on the host on which the virtual machine runs. Snapshots are not supported with DirectPath I/O passthrough devices.
Pointing device	Mirrors the pointing device that is connected to the virtual machine console when you first connect to the console.
Serial Port	Interface for connecting peripherals to the virtual machine. The virtual serial port can connect to a physical serial port, to a file on the host computer, or over the network. You can also use it to establish a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. You can configure a virtual machine with up to 32 serial ports. You can add, remove, or configure virtual serial ports.
SATA controller	Provides access to virtual disks and DVD/CD-ROM devices. The SATA virtual controller appears to a virtual machine as an AHCI SATA Controller.
SCSI controller	Provides access to virtual disks. The SCSI virtual controller appears to a virtual machine as different types of controllers, including LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual. You can change the SCSI controller type, allocate bus sharing for a virtual machine, or add a paravirtualized SCSI controller.
SIO controller	Provides serial and parallel ports, floppy devices, and performs system management activities. One SIO controller is available to the virtual machine. You cannot configure or remove this device.
USB controller	The USB hardware chip that provides USB 1.x and USB 2.0 function to the USB ports that it manages. The virtual USB Controller is the software virtualization of the USB 1.x and USB 2.0 host controller function in the virtual machine.
USB xHCI controller	The USB hardware chip that provides USB 3 function to the USB ports that it manages. The virtual USB xHCI controller is the software virtualization of the USB 3 host controller function in the virtual machine.
USB device	
VMCI	Virtual Machine Communication Interface device. Provides a high-speed communication channel between a virtual machine and the hypervisor. You cannot add or remove VMCI devices.
NVMe controller	NVM Express controller. NVMe is a logical device interface specification for accessing nonvolatile storage media attached through a PCI Express (PCIe) bus in real and virtual hardware.
NVDIMM controller	Provides access to the non-volatile memory resources of the host.
NVDIMM device	Non-Volatile Dual In-Line Memory Module. NVDIMM modules are memory devices that sit on an ordinary memory channel, but contain non-volatile memory. You can add up to 64 virtual NVDIMM devices to a virtual machine.
TPM device	Trusted Platform Module. When you add a virtual TPM 2.0 device to a virtual machine, the guest OS uses the device to store sensitive information, perform cryptographic tasks, or attest the integrity of the guest platform.

Hardware Device	Description
Virtual Precision Clock device	A virtual clock device that provides a virtual machine with access to the system time of the primary ESXi host.
Virtual Watchdog Timer device	To ensure self-reliance related to the system performance within a virtual machine. If the guest operating system stops responding and cannot recover on its own due to software glitches or errors, the watchdog timer waits for a predefined period of time and then restarts the system.
vSGX device	Virtual Intel® Software Guard Extensions (vSGX) provides a virtual machine additional security to your workloads. Intel SGX is a processor-specific technology that defines private regions of memory, called enclaves. Intel SGX protects the enclave contents from disclosure and modification in such a way that code running outside the enclave cannot access them.
RDM disk	You can use a raw device mapping (RDM) to store virtual machine data directly on a SAN LUN, instead of storing it in a virtual disk file.
PS2 controller	PS2 controller provides access to the virtual keyboard and point to the PS2 interface.
Video card	A virtual graphics card that provides graphics acceleration and display capabilities for virtual machine consoles.

Virtual Machine Options

Use the available virtual machine options to fine-tune the settings and behavior of your virtual machine and to ensure maximum performance.

A virtual machine might be running in any of several locations, such as ESXi hosts, data centers, clusters, or resource pools. Many of the options and resources that you configure have dependencies on and relationships with these objects.

VMware virtual machines have the following options.

General Options

View or modify the virtual machine name, and check the location of the configuration file and the working location of the virtual machine.

Encryption Options

Enable or deactivate encryption for the virtual machine if the vCenter Server instance is in a trusted relationship with a KMS server. For more information, see the *vSphere Security* documentation.

You can also enable or deactivate encrypted vMotion for virtual machines that are not encrypted. You can set encrypted vMotion to the deactivated, opportunistic, or required state. You can enable encrypted vMotion during virtual machine creation. Alternatively, you can change the encrypted vMotion state later. For more information, see the *vCenter Server and Host Management* documentation.

Power Management

Manage guest power options. Suspend the virtual machine or leave the virtual machine powered on when you put the guest operating system into standby.

VMware Tools

Manage the power controls for the virtual machine, run VMware Tools scripts, and upgrade VMware Tools during power cycling. Automatically synchronize the time between guest and host operating systems, and periodically synchronize the guest time with the host if the virtual machine guest operating system does not have a native time synchronization software.

Virtualization Based Security (VBS)

Enable VBS to provide an additional level of protection to the virtual machine. VBS is available on the latest Windows OS versions. For more information, see the *vSphere Security* documentation.

Boot Options	Set the boot delay when powering on virtual machines or to force BIOS setup and configure failed boot recovery.
Advanced Options	Deactivate acceleration and enable logging, configure debugging and statistics, and change the swap file location. You can also change the latency sensitivity and add configuration parameters.
Fibre Channel NPIV	Control virtual machine access to LUNs on a per-virtual machine basis. N-port ID virtualization (NPIV) provides the ability to share a single physical Fibre Channel HBA port among multiple virtual ports, each with unique identifiers.
vApp Options	Enable or deactivate the vApp functionality in a virtual machine. When you enable vApp options, you can view and edit vApp properties, vApp Deployment options, and vApp Authoring options. For example, you can configure an IP allocation policy or a network protocol profile for the vApp. A vApp option that is specified at the level of a virtual machine overrides the settings specified at the level of the vApp.

ThevSphere Client

All administrative functions are available through thevSphere Client .

ThevSphere Client is a cross-platform application that can connect only tovCenter Server. It has a full range of administrative functionality and an extensible plug-in-based architecture. Typical users are virtual infrastructure administrators, help desk, network operations center operators, and virtual machine owners.

Users can use thevSphere Client to accessvCenter Server through a Web browser.vSphere Client uses the VMware API to mediate the communication between the browser and thevCenter Server .

ThevSphere Client introduces some productivity enhancements and usability improvements. Watch the following video to learn about the usability improvements to working with virtual machine hard disks.

Where to Go From Here

You must create, provision, and deploy your virtual machines before you can manage them.

To begin provisioning virtual machines, determine whether to create a single virtual machine and install an operating system and VMware tools, work with templates and clones, or deploy virtual machines, virtual appliances, or vApps stored in Open Virtual Machine Format (OVF).

After you provision and deploy virtual machines into the vSphere infrastructure, you can configure and manage them. You can configure existing virtual machines by modifying or adding hardware or install or upgrade VMware Tools. You might need to manage multitiered applications with VMware vApps or change virtual machine startup and shutdown settings, use virtual machine snapshots, work with virtual disks, or add, remove, or delete virtual machines from the inventory.

Deploying Virtual Machines

VMware supports several methods to provision vSphere virtual machines. What works best in your environment depends on factors such as the size and type of your infrastructure and the goals that you want to achieve.

Create a single virtual machine if no other virtual machines in your environment have the requirements you are looking for, such as a particular operating system or hardware configuration. You can also create a single virtual machine and install an operating system on it, and then use that virtual machine as a template from which to clone other virtual machines. See [Create a Virtual Machine with the New Virtual Machine Wizard](#).

To use a preconfigured virtual machine, deploy and export virtual machines, virtual appliances, and vApps stored in Open Virtual Machine Format (OVF). A virtual appliance is a virtual machine that typically has an operating system and other

software installed. You can deploy virtual machines from local file systems and from shared network drives. See [Deploy and Export OVF and OVA Templates](#).

Create a Virtual Machine with the New Virtual Machine Wizard

If no virtual machines in your environment meet your needs, you can create a single virtual machine, for example of a particular operating system or hardware configuration.

Verify that you have the following privileges:

- **Virtual machine > Edit inventory > Create new** on the destination folder or data center.
- **Virtual machine > Change Configuration > Add new disk** on the destination folder or data center, if you are adding a new disk.
- **Virtual machine > Change Configuration > Add existing disk** on the destination folder or data center, if you are adding an existing disk.
- **Virtual machine > Configuration > Configure Raw device** on the destination folder or data center, if you are using an RDM or SCSI pass-through device.
- **Virtual machine > Change Configuration > Configure Host USB device** on the destination folder or data center, if you are attaching a virtual USB device backed by a host USB device.
- **Virtual machine > Change Configuration > Advanced configuration** on the destination folder or data center, if you are configuring advanced virtual machine settings.
- **Virtual machine > Configuration > Change Swapfile placement** on the destination folder or data center, if you are configuring swap file placement.
- **Virtual machine > Change Configuration > Toggle disk change tracking** on the destination folder or data center, if you are enabling change tracking on the virtual machine's disks.
- **Resource > Assign virtual machine to resource pool** on the destination host, cluster, or resource pool.
- **Datastore > Allocate space** on the destination datastore or datastore folder.
- **Network > Assign network** on the network that the virtual machine will be assigned to.

To verify the privileges assigned to your role, click **Menu > Administration > Roles** and select the role.

If you want to create a virtual machine that uses persistent memory, choose a host or a cluster with an available PMem resource.

When you create a virtual machine without a template or clone, you can configure the virtual hardware, including processors, hard disks, and memory. You open the New Virtual Machine wizard from any object in the inventory that is a valid parent object of a virtual machine.

During the creation process, a default disk is configured for the virtual machine. You can remove this disk and add a new hard disk, select an existing disk, or add an RDM disk on the Virtual Hardware page of the wizard.

1. On the **Select a creation type** page, select **Create a new virtual machine** and click **Next**.
2. On the **Select a name and folder** page, enter a unique name for the virtual machine and select a deployment location.
3. Click **Next**.
4. On the **Select a compute resource** page, select the host, cluster, resource pool, or vApp where the virtual machine will run and click **Next**.

If creating the virtual machine at the selected location causes compatibility problems, an alarm appears in the **Compatibility** pane.

5. On the **Select storage** page, choose the storage type, the storage policy, and a datastore or datastore cluster where to store the virtual machine files.

Option	Description
Create a virtual machine on a host that has PMem resource	1. Choose the type of storage by selecting the Standard or the PMem radio button. With the PMem storage option, every virtual machine disk file is stored on the host-local PMem datastore by default. You can change the datastore at a later time. The virtual machine home location must be on a non-PMem datastore. For more information about persistent memory and PMem storage, see the <i>vSphere Resource Management</i> guide. 2. From the VM Storage Policy drop-down menu, select a virtual machine storage policy or leave the default one. 3. To encrypt the virtual machine, select the Encrypt this virtual machine check box. 4. Select a datastore or a datastore cluster. 5. If you do not want to use storage DRS with the virtual machine, select the Disable Storage DRS for this virtual machine check box.
Create a virtual machine on a host that does not have PMem resource	1. Select a VM storage policy or leave the default one. 2. To encrypt the virtual machine, select the Encrypt this virtual machine check box. 3. Select a datastore or a datastore cluster.

For information about creating an encrypted virtual machine, see *vSphere Security*.

6. On the **Select compatibility** page, select the virtual machine compatibility with ESXi host versions and click **Next**.

To have access to the latest hardware features, select the latest ESXi host version.

7. On the **Select a guest OS** page, select the guest OS family and version and click **Next**.

When you select a guest operating system, BIOS or Extensible Firmware Interface (EFI) is selected by default, depending on the firmware supported by the operating system. Mac OS X Server guest operating systems support only EFI. If the operating system supports BIOS and EFI, you can change the default by editing the virtual machine after you create it and before you install the guest operating system. If you select EFI, you cannot boot an operating system that supports only BIOS, and the reverse.

Important: Do not change the firmware after the guest operating system is installed. The guest operating system installer partitions the disk in a particular format, depending on which firmware the installer was booted from. If you change the firmware, you will not be able to boot the guest.

8. Optional: Enable **Windows Virtualization Based Security**.

When you enable this option, hardware virtualization, IOMMU, EFI, and secure boot become available to the guest operating system. You must also enable **Virtualization Based Security** within the guest operating system of this virtual machine.

The **Enable Windows Virtualization Based Security** option is available only for the latest Windows OS versions, for example Windows 10 and Windows Server 2016. For more information about VBS, see the *vSphere Security* documentation.

9. Click **Next**.
10. On the **Customize hardware** page, configure the virtual machine hardware and options and click **Next**.
You can leave the defaults and configure the virtual machine hardware and options later. For more information, see [Configuring Virtual Machine Hardware](#) and [Configuring Virtual Machine Options](#).
Important: If you chose to use PMem storage for the virtual machine, its default hard disk, the new hard disks that you configure, and the NVDIMM devices that you add to the virtual machine all share the same PMem resources. So, you must adjust the size of the newly added devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.
11. On the **Ready to complete** page, review the details and click **Finish**.

The virtual machine appears in the vSphere Client inventory.

Clone a Virtual Machine to a Template

After you create a virtual machine, you can clone it to a template. Templates are primary copies of virtual machines that let you create ready-for-use virtual machines.

If a load generator is running in the virtual machine, stop it before you perform the clone operation.

Verify that you have the following privileges:

- **Virtual machine > Provisioning > Create template from virtual machine** on the source virtual machine.
- **Virtual machine > Edit inventory > Create from existing** on virtual machine folder where the template is created.
- **Resource > Assign virtual machine to resource pool** on the destination host, cluster, or resource pool.
- **Datastore > Allocate space** on all datastores where the template is created.

You can make changes to the template, such as installing additional software in the guest operating system, while preserving the original virtual machine.

You cannot modify templates after you create them. To alter an existing template, you must convert it to a virtual machine, make the required changes, and convert the virtual machine back to a template. To preserve the original state of a template, clone the template to a template.

1. Start the **Clone Virtual Machine To Template** wizard.

Option	Description
From a valid parent object of a virtual machine	1. Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host, and select New Virtual Machine. 2. On the Select a creation type page, select Clone virtual machine to template and click Next. 3. On the Select a virtual machine page, select the virtual machine that you want to clone.
From a virtual machine	Right-click the virtual machine and select Clone > Clone to Template .

2. On the Select a name and folder page, enter a name for the template and select a data center or a folder in which to deploy it.

The template name determines the name of the files and folder on the disk. For example, if you name the template win8tmp, the template files are named win8tmp.vmdk, win8tmp.nvram, and so on. If you change the template name, the names of the files on the datastore do not change.

Folders provide a way to store virtual machines and templates for different groups in an organization and you can set permissions on them. If you prefer a flatter hierarchy, you can put all virtual machines and templates in a datacenter and organize them a different way.

- On the Select a compute resource, select a host or a cluster resource for the template.

The **Compatibility** pane shows the result from the compatibility checks.

Important:

If the virtual machine that you clone has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, you cannot proceed with the task.

If the virtual machine that you clone does not have an NVDIMM device, but it has virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, all the hard disks of the template will use the storage policy and datastore selected for the configuration files of the source virtual machine.

- On the Select storage page, select the datastore or datastore cluster in which to store the template configuration files and all of the virtual disks. Click **Next**.

Option	Description
Clone a virtual machine that has vPMem hard disks	- Choose the type of storage for the template by selecting the Standard, the PMem, or the Hybrid radio button. If you select the Standard mode, all virtual disks are stored on a standard datastore. If you select the PMem mode, all virtual disks are stored on the host-local PMem datastore. Configuration files cannot be stored on a PMem datastore and you must additionally select a regular datastore for the configuration files of the virtual machine. If you select the Hybrid mode, all PMem virtual disks remain stored on a PMem datastore. Non-PMem disks are affected by your choice of a VM storage policy and datastore or datastore cluster. For more information about persistent memory and PMem storage, see the <i>vSphere Resource Management</i> guide. - From the Select virtual disk format drop-down menu, select a new virtual disk format for the template or keep the same format as the source virtual machine. - From the VM Storage Policy drop-down menu, select a virtual machine storage policy or leave the default one. - Select a datastore or a datastore cluster. - Select the Disable Storage DRS for this virtual machine check box if you do not want to use storage DRS with the virtual machine. - Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.
Clone a virtual machine that does not have vPMem hard disks	- Select the disk format for the virtual machine virtual disks. Same format as source uses the same disk format as the source virtual machine. The Thick Provision Lazy Zeroed format creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out later, on demand, on first write from the virtual machine. Thick Provision Eager Zeroed is a type of thick virtual disk that supports clustering features such as Fault tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat file format (<code>-flat.vmdk</code>), the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types of disks. The Thin Provision format saves storage space. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can expand to the maximum capacity allocated to it. - Select a VM storage policy or leave the default one. - Select a datastore or a datastore cluster.

Option	Description
	4. Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.

Important: You cannot change the storage policy if you clone an encrypted virtual machine. For information about cloning an encrypted virtual machine, see *vSphere Security*.

- On the Ready to complete page, review the template settings and click **Finish**.

The progress of the clone task appears in the **Recent Tasks** pane. When the task completes, the template appears in the inventory.

Deploy a Virtual Machine from a Template

Deploying a virtual machine from a template creates a virtual machine that is a copy of the template. The new virtual machine has the virtual hardware, installed software, and other properties that are configured for the template.

You must have the following privileges to deploy a virtual machine from a template:

- Virtual machine > Edit inventory > Create from existing** on the data center or virtual machine folder.
- Virtual machine > Change Configuration > Add new disk** on the data center or virtual machine folder. Required only if you customize the original hardware by adding a new virtual disk.
- Virtual machine > Provisioning > Deploy template** on the source template.
- Resource > Assign virtual machine to resource pool** on the destination host, cluster, or resource pool.
- Datastore > Allocate space** on the destination datastore.
- Network > Assign network** on the network to which the virtual machine is assigned. Required only if you customize the original hardware by adding a new network card.
- Virtual machine > Provisioning > Customize guest** on the template or template folder if you are customizing the guest operating system.
- Virtual machine > Provisioning > Read customization specifications** on the root vCenter Server if you are customizing the guest operating system.
- Virtual machine > Edit inventory > Create new** on the destination folder or data center.
- vApp > Import**

- Start the **Deploy From Template** wizard.

Option	Description
From a valid parent object of a virtual machine	- Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host, and select New Virtual Machine. - On the Select a creation type page, select Deploy from template and click Next. - On the Select a template page, select the template that you want to use.
From a template	Right-click a template and select New VM from This Template .

2. On the **Select a name and folder** page, enter a unique name for the virtual machine and select a deployment location.
3. On the **Select a compute resource** page, select the host, cluster, resource pool, or vApp where the virtual machine will run and click **Next**.

The virtual machine will have access to the resources of the selected object.

Important:

If the template that you deploy has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, you cannot proceed with the task.

If the template that you deploy does not have an NVDIMM device, but it has virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, all the hard disks of the virtual machine will use the storage policy and datastore selected for the configuration files of the source template.

If creating the virtual machine at the selected location causes compatibility problems, an alarm appears in the **Compatibility** pane.

4. On the **Select storage** page, select the datastore or datastore cluster in which to store the virtual machine configuration files and all of the virtual disks. Click **Next**.

Option	Description
Deploy a virtual machine from a template that has vPMem hard disks	1. Choose the type of storage for the template by selecting the Standard, the PMem, or the Hybrid radio button. If you select the Standard mode, all virtual disks will be stored on a standard datastore. If you select the PMem mode, all virtual disks will be stored on the host-local PMem datastore. Configuration files cannot be stored on a PMem datastore and you must additionally select a regular datastore for the configuration files of the virtual machine. If you select the Hybrid mode, all PMem virtual disks will remain stored on a PMem datastore. Non-PMem disks are affected by your choice of a VM storage policy and datastore or datastore cluster. For more information about persistent memory and PMem storage, see the <i>vSphere Resource Management</i> guide. 2. From the Select virtual disk format drop-down menu, select a new virtual disk format for the template or keep the same format as the source virtual machine. 3. From the VM Storage Policy drop-down menu, select a virtual machine storage policy or leave the default one. 4. Select a datastore or a datastore cluster. 5. Select the Disable Storage DRS for this virtual machine check box if you do not want to use storage DRS with the virtual machine. 6. Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.
Deploy a virtual machine from a template that does not have vPMem hard disks	1. Select the disk format for the virtual machine virtual disks. Same format as source uses the same disk format as the source virtual machine. The Thick Provision Lazy Zeroed format creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out later, on demand, on first write from the virtual machine. Thick Provision Eager Zeroed is a type of thick virtual disk that supports clustering features such as Fault tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types o disks.

Option	Description
	The Thin Provision format saves storage space. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can grow to the maximum capacity allocated to it. 2. Select a VM storage policy or leave the default one. 3. Select a datastore or a datastore cluster. 4. Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.

5. On the **Select clone options** page, select additional customization options for the new virtual machine.
You can choose to customize the guest operating system or the virtual machine hardware. You can also choose to power on the virtual machine after its creation.
6. Optional: On the **Customize guest OS** page, apply a customization specification to the virtual machine.
Customizing the guest OS prevents from conflicts that might occur if you deploy virtual machines with identical settings, such as duplicate computer names.

Note: To access customization options for Windows guest operating systems, Microsoft Sysprep tools must be installed on the vCenter Server system. The Sysprep Tool is built into the Windows Vista and Windows 2008 and later operating systems. For details about this and other customization requirements, see [Guest Operating System Customization Requirements](#).

Option	Description
Select an existing specification	Select a customization specification from the list.
Create a specification	Click the Create a new specification icon, and complete the steps in the wizard.
Create a specification from an existing specification	1. Select a customization specification from the list. 2. Click the Create a spec from an existing spec icon, and complete the steps in the wizard.

7. Optional: On the **Customize hardware** page, configure the virtual machine hardware and options and click **Next**.
You can leave the defaults and configure the virtual machine hardware and options later. For more information, see [Configuring Virtual Machine Hardware](#) and [Configuring Virtual Machine Options](#).
Important: If you chose to use PMem storage for the virtual machine, its default hard disk, the new hard disks that you configure, and the NVDIMM devices that you add to the virtual machine all share the same PMem resources. So, you must adjust the size of the newly added devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.
8. On the Ready to complete page, review the information and click **Finish**.

Clone an Existing Virtual Machine

Cloning a virtual machine creates a virtual machine that is a copy of the original. The new virtual machine is configured with the same virtual hardware, installed software, and other properties that were configured for the original virtual machine.

If a load generator is running in the virtual machine, before you perform the clone operation, you must stop the load generator.

You must have the following privileges to clone a virtual machine:

- **Virtual machine > Provisioning > Clone virtual machine** on the virtual machine you are cloning.

- **Virtual machine > Edit inventory > Create from existing** on the datacenter or virtual machine folder.
- **Virtual machine > Change Configuration > Add new disk** on the datacenter or virtual machine folder.
- **Resource > Assign virtual machine to resource pool** on the destination host, cluster, or resource pool.
- **Datastore > Allocate space** on the destination datastore or datastore folder.
- **Network > Assign network** on the network to which you assign the virtual machine.
- **Virtual machine > Provisioning > Customize guest** on the virtual machine or virtual machine folder if you are customizing the guest operating system.
- **Virtual machine > Provisioning > Read customization specifications** on the root vCenter Server if you are customizing the guest operating system.
- If the virtual machine that you clone has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have an available PMem resource. Otherwise, you cannot proceed with the task
- If the virtual machine that you clone does not have an NVDIMM device, but has virtual PMem hard disks, the destination host or cluster must have an available PMem resource. Otherwise, all hard disks of the destination virtual machine will use the storage policy and datastore selected for the configuration files of the source virtual machine.
- To access customization options for Windows guest operating systems, Microsoft Sysprep tools must be installed on the vCenter Server system. Sysprep Tool is built into the Windows Vista and Windows 2008 and later operating systems. For details about this and other customization requirements, see [Guest Operating System Customization Requirements](#).

For information about persistent memory and PMem storage, see the *vSphere Resource Management* guide.

For information how to configure the virtual machine hardware options, see [Configuring Virtual Machine Hardware](#) and [Configuring Virtual Machine Options](#)

Note: When heavily loaded applications, such as load generators, are running in the guest operating system during a clone operation, the virtual machine quiesce operation might fail. VMware Tools might be denied CPU resources and time out. You can quiesce the virtual machines running lower I/O disk operation.

Important: If you clone an encrypted virtual machine, you cannot change the storage policy. For information about cloning an encrypted virtual machine, see *vSphere Security*.

1. Start the **Clone Existing Virtual Machine** wizard.

Option	Action
From a valid parent object of a virtual machine	1. Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, cluster, vApp, resource pool, or host, and select New Virtual Machine. 2. On the Select a creation type page, select Clone an existing virtual machine, and click Next. 3. On the Select a virtual machine page, select the virtual machine that you want to clone.
From a virtual machine	Right-click a virtual machine and select Clone > Clone to Virtual Machine .

2. On the **Select a name and folder** page, enter a unique name for the new virtual machine, select a deployment location, and click **Next**.

The template name determines the name of the files and folder on the disk. For example, if you name the template `win8tmp`, the template files are named `win8tmp.vmdk`, `win8tmp.nvram`, and so on. If you change the template name later, the names of the files on the datastore do not change.

Folders provide a way to store virtual machines and templates for different groups in an organization and you can set permissions on them. If you prefer a flatter hierarchy, you can put all virtual machines and templates in a data center and organize them in a different way.

- On the **Select a compute resource** page, select the host, cluster, resource pool, or vApp where the virtual machine will run and click **Next**.

The **Compatibility** pane shows the result from the compatibility checks.

- On the **Select storage** page, select the datastore or datastore cluster in which to store the template configuration files and all virtual disks.

Option	Action
Clone a virtual machine that has vPMem hard disks	- Select the type of storage for the template by clicking the Standard, the PMem, or the Hybrid radio button. - If you select the Standard mode, all virtual disks are stored on a standard datastore. - If you select the PMem mode, all virtual disks are stored on the host-local PMem datastore. Configuration files cannot be stored on a PMem datastore and you must additionally select a regular datastore for the configuration files of the virtual machine. - If you select the Hybrid mode, all PMem virtual disks remain stored on a PMem datastore. Your choice of a VM storage policy and datastore or datastore cluster affects the non-PMem disks. - From the Select virtual disk format drop-down menu, select a new virtual disk format for the template or keep the same format as the source virtual machine. - From the VM Storage Policy drop-down menu, select a virtual machine storage policy or leave the default one. - Select a datastore or a datastore cluster. - Select the Disable Storage DRS for this virtual machine check box if you do not want to use storage DRS with the virtual machine. - Enable the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.
Clone a virtual machine that does not have vPMem hard disks	- Select the disk format for the virtual machine virtual disks. - The Same format as source option uses the same disk format as the source virtual machine. - The Thick Provision Lazy Zeroed format creates a virtual disk in a default thick format. The space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out later, on demand, on first write from the virtual machine. Thick Provision Eager Zeroed is a type of thick virtual disk that supports clustering features such as Fault tolerance. The space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types of disks. - The Thin Provision format saves storage space. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can grow to the maximum capacity allocated to it. - Select a VM storage policy or leave the default one. - Select a datastore or a datastore cluster. - Enable the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.

5. Click **Next**.
6. On the **Select clone options** page, select additional customization options for the new virtual machine and click **Next**. You can choose to customize the guest operating system or the virtual machine hardware. You can also choose to power on the virtual machine after its creation.
7. Optional: On the **Customize guest OS** page, apply a customization specification to the virtual machine and click **Next**. Customizing the guest operating system prevents from conflicts that might occur if you or other users clone virtual machines with identical settings, such as duplicate computer names.

Option	Action
Select an existing specification	Select a customization specification from the list.
Override	To change the guest customization specification for this deployment only, click Override , complete the steps in the Override VM Customization Specification wizard, and click OK .

8. Optional: On the **User settings** page, specify the required settings for the virtual machine. This page of the wizard appears only if the selected specification requires additional customization.
9. Optional: On the **Customize hardware** page, configure the virtual machine hardware and options and click **Next**. You can leave the defaults and configure the virtual machine hardware and options later.

Important: If you chose to use PMem storage for the virtual machine, its default hard disk, the new hard disks that you configure, and the NVDIMM devices that you add to the virtual machine all share the same PMem resources. You must adjust the size of the newly added devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.
10. On the **Ready to complete** page, review the virtual machine settings and click **Finish**.

The new virtual machine appears in the inventory.

Cloning a Virtual Machine with Instant Clone

With the Instant Clone technology, you can create powered on virtual machines from the running state of another powered on virtual machine. The result is a new virtual machine that is identical to the source virtual machine.

With Instant Clone you can create new virtual machines from a controlled point in time. Instant cloning is very convenient for large scale application deployments because it ensures memory efficiency and allows for creating numerous virtual machines on a single host.

The result of an Instant Clone operation is a virtual machine that is called a destination virtual machine. The processor state, virtual device state, memory state, and disk state of the destination virtual machine are identical to those of the source virtual machine. To avoid network conflicts, you can customize the virtual hardware of the destination virtual machine during an Instant Clone operation. For example, you can customize the MAC addresses of the virtual NICs or the serial port configurations of the destination virtual machine. vSphere 7.0 and later does not support customization of the guest OS of the destination virtual machine. For information about manual guest OS customization, see the *vSphere Web Services SDK Programming Guide*.

During an Instant Clone operation, the source virtual machine is stunned for a short period of time, less than 1 second. While the source virtual machine is stunned, a new writable delta disk is generated for each virtual disk and a checkpoint is taken and transferred to the destination virtual machine. The destination virtual machine then powers on by using the source's checkpoint. After the destination virtual machine is fully powered on, the source virtual machine also resumes running.

Instant Cloned virtual machines are fully independent vCenter Server inventory objects. You can manage Instant Cloned virtual machines like regular virtual machines without any restrictions.

Starting with vSphere 7.0, you can Instant Clone a virtual machine only through the API calls.

For information about Instant Clone, see the *vSphere Web Services SDK Programming Guide*.

Clone a Template to a Template

Templates are primary copies of virtual machines that let you create ready-for-use virtual machines. You can change the template, for example, by installing additional software in the guest operating system, while preserving the state of the primary template.

Verify that you have the following privileges:

- **Virtual machine > Provisioning > Clone template** on the source template.
- **Virtual machine > Edit inventory > Create from existing** on the folder where the template is created.
- **Datastore > Allocate space** on all datastores where the template is created.

1. Start the **Clone Template to Template** wizard.

Option	Description
From a valid parent object of a virtual machine	1. Right-click any inventory object that is a valid parent object of a virtual machine and select New Virtual Machine. 2. Select Clone template to template and click Next. 3. On the Select a template to clone page, browse to the template that you want to clone or accept the default one.
From a template	Right-click a template and select Clone to Template .

2. On the **Select a name and folder** page, enter a unique name for the template and select the data center or folder in which to deploy it.

The template name determines the name of the files and folder on the disk. For example, if you name the template win10tmp, the template files are named win10tmp.vmdk, win10tmp.nvram, and so on. If you change the template name, the names of the files on the datastore do not change.

Folders provide a way to store virtual machines and templates for different groups in an organization and you can set permissions on them. If you prefer a flatter hierarchy, you can put all virtual machines and templates in a data center and organize them in a different way.

3. Click **Next**.
4. On the **Select a compute resource** page, select a host or cluster resource for the template.
The **Compatibility** pane shows the result from the compatibility checks.

Important:

If the template that you clone has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, you cannot proceed with the task.

If the template that you clone does not have an NVDIMM device, but it has virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, all the hard disks of the template use the storage policy and datastore selected for the configuration files of the source template.

5. On the **Select storage** page, select the datastore or datastore cluster in which to store the virtual machine configuration files and all of the virtual disks.

Option	Description
Clone a virtual machine that has vPMem hard disks	1. Choose the type of storage for the template by selecting the Standard, the PMem, or the Hybrid radio button. If you select the Standard mode, all virtual disks are stored on a standard datastore. If you select the PMem mode, all virtual disks are stored on the host-local PMem datastore. Configuration files cannot be stored on a PMem datastore and you must additionally select a regular datastore for the configuration files of the virtual machine.

Option	Description
	If you select the Hybrid mode, all PMem virtual disks remain stored on a PMem datastore. Non-PMem disks are affected by your choice of a VM storage policy and datastore or datastore cluster. For more information about persistent memory and PMem storage, see the <i>vSphere Resource Management</i> guide. - From the Select virtual disk format drop-down menu, select a new virtual disk format for the template or keep the same format as the source virtual machine. - From the VM Storage Policy drop-down menu, select a virtual machine storage policy or leave the default one. - Select a datastore or a datastore cluster. - Select the Disable Storage DRS for this virtual machine check box if you do not want to use storage DRS with the virtual machine. - Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance issues. You can also convert a standard hard disk to a PMem hard disk.
Clone a virtual machine that does not have vPMem hard disks	- Select the disk format for the virtual machine virtual disks. Same format as source uses the same disk format as the source virtual machine. The Thick Provision Lazy Zeroed format creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out later, on demand, on first write from the virtual machine. Thick Provision Eager Zeroed is a type of thick virtual disk that supports clustering features such as Fault tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types o disks. The Thin Provision format saves storage space. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can expand to its maximum capacity allocated to it. - Select a VM storage policy or leave the default one. - Select a datastore or a datastore cluster. - Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance issues. You can also convert a standard hard disk to a PMem hard disk.

Important: If you clone an encrypted virtual machine, you cannot change the storage policy. For information about cloning an encrypted virtual machine, see *vSphere Security*.

- Click **Next**.
- On the **Ready to complete** page, review the template settings and click **Finish**.

The progress of the clone task appears in the **Recent Tasks** pane. When the task completes, the template appears in the inventory.

Convert a Template to a Virtual Machine

Converting a template to a virtual machine (VM) does not make a copy. You convert a template to a VM to edit the template or if you do not need to preserve the VM as a golden image for deploying VMs.

Verify that you have the following privileges:

- **Virtual machine > Provisioning > Mark as virtual machine** on the source template.
- **Resource > Assign virtual machine to resource pool** on the resource pool where the virtual machine runs.

1. Start the **Convert Template to Virtual Machine** wizard.

Option	Description
From a valid parent object of a virtual machine	1. Right-click any inventory object that is a valid parent object of a virtual machine and select New Virtual Machine. 2. On the Select a creation type page, select Convert template to virtual machine, and click Next. 3. On the Select a template page of the wizard, select a template to deploy from the list.
From a template	Right-click the template and select Convert to Virtual Machine .

2. On the **Select a compute resource** page, select the host, cluster, vApp, or resource pool for the virtual machine to run in.

Important:

If the template that you convert has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, you cannot proceed with the task.

If the template that you convert does not have an NVDIMM device, but it has virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, all the hard disks of the virtual machine use the storage policy and datastore selected for the configuration files of the source template.

The **Compatibility** pane shows the result from the compatibility checks.

3. Click **Next**.
4. On the **Ready to complete** page, review the settings, and click **Finish**.

The virtual machine appears in the inventory.

Deploy and Export OVF and OVA Templates

You can export virtual machines, virtual appliances, and vApps in Open Virtual Format (OVF) and Open Virtual Appliance (OVA). You can then deploy the OVF or OVA template in the same environment or in a different environment.

Note: In vSphere 6.5 and later, you cannot export OVA templates, OVF templates is the only option.

Alternatively, you can use the VMware OVF Tool to deploy and export OVF templates.

OVF and OVA File Formats and Templates

OVF is an open standard which contains multiple files as a package. For example, `.ovf`, `.vmdk`, `.nvram`, and so on. OVF supports exchange of virtual appliances across products and platforms. OVA is a single-file distribution of the OVF file package. When you export a virtual machine as an OVF file, you download an OVF file and the virtual disk files.

The OVF and OVA formats offer the following advantages:

- OVF and OVA files are compressed, allowing for faster downloads.
- The vSphere Client validates an OVF or OVA file before importing it, and ensures that it is compatible with the intended destination server. If the appliance is incompatible with the selected host, it cannot be imported and an error message appears.
- OVF and OVA can encapsulate multi-tiered applications and more than one virtual machine.

Exporting OVF templates allows you to create virtual appliances that other users can import later. You can use the export function to distribute pre-installed software as a virtual appliance, or to distributing template virtual machines to users. You can make the OVF or OVA file available to users who cannot access your vCenter Server inventory.

Deploying an OVF or OVA template allows you to add pre-configured virtual machines or vApps to your vCenter Server or ESXi inventory. Deploying an OVF or OVA template is similar to deploying a virtual machine from a template. However, you can deploy an OVF or OVA template from any local file system accessible from the vSphere Web Client, or from a remote Web server. The local file systems can include local disks (such as C:), removable media (such as CDs or USB keychain drives), and shared network drives.

Deploy an OVF or OVA Template

You can deploy an OVF or OVA template from a local file system or from a URL.

Some of the pages in the **Deploy OVF Template** wizard only appear if the OVF template that you deploy requires additional customization, contains deployment options or has one or multiple vService dependencies.

1. Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host, and select **Deploy OVF Template**.
The **Deploy OVF Template** wizard opens.
2. On the **Select an OVF template** page, specify the location of the source OVF or OVA template and click **Next**.

Option	Action
URL	Type a URL to an OVF or OVA template located on the Internet. Supported URL sources are HTTP and HTTPS. Note: Verify that all files from the OVF package are at the same location of the OVF descriptor file. Example: <code>http://vmware.com/VMTN/appliance.ovf</code> .
Local file	Click Browse and select all the files associated with an OVF template or OVA file. This includes files such as <code>.ovf</code> , <code>.vmdk</code> , etc. If you do not select all the required files, a warning message displays.

3. On the **Select a name and folder** page, enter a unique name for the virtual machine or vApp, select a deployment location, and click **Next**.
The default name for the virtual machine is the same as the name of the selected OVF or OVA template. If you change the default name, choose a name that is unique within each vCenter Server virtual machine folder.
The default deployment location for the virtual machine is the inventory object where you started the wizard.
4. On the **Select a compute resource** page, select a resource where to run the deployed VM template,.
 - a. To automatically power on the virtual machine that you deploy on the host, select the **Automatically power on deployed VM** check box.
 - b. Click **Next**.
5. On the **Review details** page, verify the OVF or OVA template details and click **Next**.

Option	Description
Publisher	Publisher of the OVF or OVA template, if a certificate included in the OVF or OVA template file specifies a publisher.
Download size	Size of the OVF or OVA file.
Size on disk	Size on disk after you deploy the OVF or OVA template.

6. On the **Configuration** page, select a deployment configuration and click **Next**.

7. On the **Select storage** page, define where and how to store the files for the deployed OVF or OVA template.
 - a. Select the disk format for the virtual machine virtual disks.

Format	Description
Thick Provision Lazy Zeroed	Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out later, on demand, on first write from the virtual machine.
Thick Provision Eager Zeroed	A type of thick virtual disk that supports clustering features such as Fault tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat file format (-flat.vmdk), the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types of disks.
Thin Provision	Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk requires based on the value that you enter for the disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations.

- b. Select a VM Storage Policy.
This option is available only if storage policies are enabled on the destination resource.
 - c. Enable the **Show datastores from Storage DRS clusters** check box to choose individual datastores from Storage DRS clusters for the initial placement of the virtual machine.
 - d. Select a datastore to store the deployed OVF or OVA template.
The configuration file and virtual disk files are stored on the datastore. Select a datastore large enough to accommodate the virtual machine or vApp and all associated virtual disk files.

Note: If you want to use the API calls to deploy an OVF template that contains vPMem hard disks and that has been exported from a content library, consult <https://kb.vmware.com/kb/52370>.
8. On the **Select networks** page, select a source network and map it to a destination network. Click **Next**.
The Source Network column lists all networks that are defined in the OVF or OVA template.
9. On the **Customize template** page, customize the deployment properties of the OVF template and click **Next**.
10. On the **vService bindings** page, select a binding service provider and click **Next**.
11. On the **Ready to complete** page, review the page and click **Finish**.

As a result, a new task for creating the virtual machine appears in the Recent Task pane. After the task is complete, the new virtual machine is created on the selected resource.

Export an OVF Template

An OVF template captures the state of a virtual machine or vApp into a self-contained package. The disk files are stored in a compressed, sparse format.

Note: You cannot export encrypted virtual machines.

Prerequisites

- Power off the virtual machine or vApp.
- Required privilege: **vApp > Export**

Procedure

1. Navigate to a virtual machine or vApp and from the **Actions** menu, select **Template > Export OVF Template**.
2. In the **Name** folder, enter the name of the template.

Example: For example, enter `MyVm`.

Note: When you export an OVF template with a name that contains asterisk (*) characters, those characters turn into underscore (_) characters.

3. In the **Annotation** field, enter a description.
4. Select the **Enable advanced options** check box if you want to include additional information or configurations in the exported template.
The advanced settings include information about the BIOS UUID, MAC, boot order, PCI Slot numbers, and configuration settings used by other applications. These options limit portability.
5. To save each file associated with the template (.ovf, .vmdk, .mf), click **OK** and respond to the prompts.

Export an OVF Template with a vTPM Placeholder

With vSphere 8.0, you can manually export an OVF template of a virtual machine that contains a vTPM device. Because vTPMs are security devices with unique identities, a vTPM placeholder is added to the OVF template instead of a vTPM device and replaced with a new vTPM device upon instantiation. To deploy or export an OVF template of a virtual machine with a vTPM device, you can use the VMware OVF Tool. For more information about the VMware OVF Tool, see *OVF Tool User's Guide*.

Prerequisites

- Power off the virtual machine.
- Deactivate any applications in the guest operating system of the virtual machine that use the vTPM device, such as BitLocker.
- Verify that you have the required privilege: **Cryptographic operations > Decrypt**.

Procedure

1. Remove the vTPM device from the virtual machine.
 - a. Navigate to the virtual machine and click **Actions**.
 - b. Click **Edit Settings**.
 - c. On the **Virtual Hardware** tab, locate the Trusted Platform Module entry.
 - d. Move your pointer over the device and click the **Remove** icon.
This icon appears only for the virtual hardware that you can safely remove.
 - e. Click **Delete** to confirm you want to remove the device.
The vTPM device is marked for removal.
 - f. Click **OK**.
Verify that the Virtual Trusted Platform Module entry no longer appears in the virtual machine **Summary** tab in the **VM Hardware** pane.
2. Decrypt the virtual machine if it is encrypted.
For more information on how to decrypt a virtual machine, see [Decrypt an Encrypted Virtual Machine or Virtual Disk](#).
3. Export the virtual machine to an OVF template.
4. Open the OVF (.ovf) descriptor file and manually add the vTPM device placeholder:

```
<Item ovf:required="false">
  <rasd:AutomaticAllocation>false</rasd:AutomaticAllocation>
  <rasd:ElementName>Virtual TPM</rasd:ElementName>
  <rasd:InstanceID>13</rasd:InstanceID>
  <rasd:ResourceSubType>vmware.vtpm</rasd:ResourceSubType>
  <rasd:ResourceType>1</rasd:ResourceType>
```


</Item>

Note: The OVF descriptor contains a placeholder for a vTPM device which does not have any secrets. The virtual machine encryption keys are generated when you deploy a virtual machine from this OVF template.

5. Open the OVF manifest file (.mf) and change the SHA-256 checksum for the .ovf file.
6. Resign the .ovf file.
7. Deploy the manually exported OVF template.

Export an OVF Template of a Virtual Machine without a vTPM Placeholder

To deploy an OVF template of a virtual machine with a vTPM device, you can use the VMware OVF Tool. For more information about the VMware OVF Tool, see *OVF Tool User's Guide*.

Prerequisites

- Power off the virtual machine.
- Deactivate any applications in the guest operating system of the virtual machine that use the vTPM device, such as BitLocker.

Procedure

1. Remove the vTPM device from the virtual machine.
 - a. Navigate to the virtual machine and click **Actions**.
 - b. Click **Edit Settings**.
 - c. On the **Virtual Hardware** tab, locate the Trusted Platform Module entry.
 - d. Move your pointer over the device and click the **Remove** icon.
This icon appears only for the virtual hardware that you can safely remove.
 - e. Click **Delete** to confirm you want to remove the device.
The vTPM device is marked for removal.
 - f. Click **OK**.
Verify that the Virtual Trusted Platform Module entry no longer appears in the virtual machine **Summary** tab in the **VM Hardware** pane.
2. Decrypt the virtual machine if it is encrypted.
For more information on how to decrypt a virtual machine, see [Decrypt an Encrypted Virtual Machine or Virtual Disk](#).
3. Export the virtual machine to an OVF template.
4. Deploy the manually exported OVF template and add a vTPM device to the virtual machine.
For more information on how to add a vTPM device to a virtual machine, see [Add Virtual Trusted Platform Module to an Existing Virtual Machine](#).

Browse VMware Virtual Appliance Marketplace

The Virtual Appliance Marketplace contains a variety of virtual appliances packaged in OVF format that you can download and deploy in your vSphere environment.

1. Go to the [Virtual Appliance Marketplace](#), which is part of the VMware Solution Exchange.
2. Search the Marketplace to find a prepackaged appliance.
3. Log in and download the appliance.
4. Deploy the appliance in your vSphere environment.

Using Content Libraries

Content libraries are container objects for VM and vApp templates and other types of files, such as ISO images, text files, and so on.

To deploy virtual machines and vApps in the vSphere inventory, you can use the templates in the library. You can also use content libraries to share content across vCenter Server instances in the same or different locations. Sharing templates and files results in consistency, compliance, efficiency, and automation in deploying workloads at scale.

A content library stores and manages content in the form of library items. A single library item can consist of one file or multiple files. For example, the OVF template is a set of files (.ovf, .vmdk, and .mf). When you upload an OVF template to the library, you upload the entire set of files, but the result is a single library item of the OVF Template type.

Starting with vSphere 7.0 Update 3, you can protect the OVF items by applying default OVF security policy to a content library. The OVF security policy enforces strict validation on OVF items when you deploy or update the item, import items, or synchronize OVF and OVA templates. To make sure that the OVF and OVA templates are signed by a trusted certificate, you can add the OVF signing certificate from a trusted CA.

In previous versions of vSphere, content libraries supported only OVF templates. As a result, VM and vApp templates were both converted to OVF files when you uploaded them to a content library. Starting with vSphere 7.0 Update 1, content libraries also support VM templates. So, templates in the content library can either be of the OVF Template type, or the VM Template type. vApp templates are still converted to OVF files when you upload them to a content library.

You create and manage a content library from a single vCenter Server instance, but you can distribute the content to other vCenter Server instances if HTTP(S) traffic is allowed between the two systems. The distribution of VM templates additionally requires that the respective vCenter Server instances are in Enhanced Linked Mode or Hybrid Linked Mode and that the respective hosts are connected through a network.

You can create two types of content libraries.

- You can create a local content library to store and manage content in a single vCenter Server instance. If you want to share the contents of that library, you can enable publishing. When you enable publishing, other users can subscribe to the library and use its content. Alternatively, you can create subscriptions for the library, which gives you control over the distribution of content.
- You can create a subscribed content library to subscribe to a published library and use its contents. You cannot upload or import items into a subscribed library. Subscribers only use the content in the published library, but it is the administrator of the published library who manages the templates.

Content Library Administrator Role, Permissions Hierarchy, and Inheritance

vCenter Server provides a sample role that allows you to give users or groups privileges to manage content libraries.

Content Library Administrator Role

vCenter Server provides a sample role that allows you to give users or groups privileges to manage content libraries.

Content Library Administrator role is a predefined role that gives a user privileges to monitor and manage a library and its contents.

You can modify the role or use it as an example to create custom roles for specific tasks you want to allow other users to perform.

If a user has this role on a library, that user can perform the following tasks on that library.

- Create, edit, and delete local or subscribed libraries.
- Create and delete subscriptions to a local library with publishing enabled.
- Publish a library or a library item to a subscription.
- Synchronize a subscribed library and synchronize items in a subscribed library.
- View the item types supported by the library.
- Configure the global settings for the library.
- Import items to a library.
- Export library items.


Content Library Permissions Hierarchy and Inheritance

vSphere objects inherit permissions from a parent object in the hierarchy. Content libraries work in the context of a single vCenter Server instance. However, content libraries are not direct children of a vCenter Server system from an inventory perspective.

The direct parent for content libraries is the global root. This means that if you set a permission at a vCenter Server level and propagate it to the children objects, the permission applies to data centers, folders, clusters, hosts, virtual machines, and so on, but does not apply to the content libraries that you see and operate with in this vCenter Server instance. To assign a permission on a content library, an Administrator must grant the permission to the user as a global permission. Global permissions support assigning privileges across solutions from a global root object.

The figure illustrates the inventory hierarchy and the paths by which permissions can propagate.

Figure 65: vSphere Inventory Hierarchy



To let a user manage a content library and its items, an Administrator can assign the Content Library Administrator role to that user as a global permission. The Content Library Administrator role is a sample role in the vSphere Client.

Users who are Administrators can also manage libraries and their contents. If a user is an Administrator at a vCenter Server level, they have sufficient privileges to manage the libraries that belong to this vCenter Server instance, but cannot see the libraries unless they have a Read-Only role as a global permission.

For example, a user has an Administrator role that is defined at a vCenter Server level. When the Administrator navigates to Content Libraries in the object navigator, he sees 0 libraries despite there are existing libraries in the vSphere inventory of that vCenter Server instance. To see the libraries, the Administrator needs a Read-Only role assigned as a global permission.

Administrators whose role is defined as a global permission can see and manage the libraries in all vCenter Server instances that belong to the global root.

Because content libraries and their children items inherit permissions only from the global root object, when you navigate to a library or a library item and click **Configure** tab, you can see there is no **Permissions** tab. An Administrator cannot assign individual permissions on different libraries or different items within a library.

Create and Edit a Content Library

In the vSphere Client, you can create a local or a subscribed content library and change its configuration settings. By using content libraries, you can store and manage content in one vCenter Server instance. Alternatively, you can distribute content across vCenter Server instances to increase consistency and facilitate the deployment workloads at scale.

By using the OVF security policy, you can protect the items of your content library.

Create a Content Library

You can create a local content library and populate it with templates and other types of files. You can then use the content library templates to deploy virtual machines or vApps in your virtual environment. You can also share the contents of your local library with users from other vCenter Server instances.

Required privileges:

- **Content library > Create local library** or **Content library > Create subscribed library** on the vCenter Server instance where you want to create the library.
- **Datastore > Allocate space** on the destination datastore.

Note: VM Templates cannot be stored in a library that uses NFS or SMB storage. OVF templates can only be stored on the datastore that is associated with the content library.

You can also create a subscribed library to use at will the contents of an already existing published local library.

1. Navigate to **Menu > Content Libraries**.
2. Click **Create**.
The **New Content Library** wizard opens.
3. On the **Name and location** page, enter a name, select a vCenter Server instance for the content library and click **Next**.
4. On the **Configure content library** page, select the type of content library that you want to create.

Option	Description
Local content library	A local content library is accessible only in the vCenter Server instance where you create it by default. 1. To make the content of the library available to other vCenter Server instances, select Enable publishing.

Option	Description
	- If you want to require a password for accessing the content library, select Enable authentication and set a password. - Click Next.
Subscribed content library	A subscribed content library originates from a published content library. Use this option to take advantage of already existing content libraries. Note: If the published content library does not have a security policy enabled, you cannot apply the OVF default security policy to the subscribed content library. If you create a subscribed library, depending on the Subscription URL that you provide, the system performs a check to determine whether the provided URL can have a security policy enabled or not. To see up-to-date content, you can synchronize the subscribed library with the published library, but you cannot add or remove content from the subscribed library. Only an administrator of the published library can add, modify, and remove contents from the published library. To subscribe to a library, provide the following information: - In the Subscription URL text box, enter the URL address of the published library. Note: The system performs a check to determine whether the Subscription URL that you provide can have a security policy enabled or not. - If authentication is enabled on the published library, select Enable authentication and enter the publisher password. - Select a download method for the contents of the subscribed library. - To download a local copy of all the items in the published library immediately after subscribing to it, select immediately. - To save storage space, select when needed. You download only the metadata for the items in the published library. If you must use an item, synchronize the item or the entire library to download its content. - If prompted, accept the SSL certificate thumbprint. The SSL certificate thumbprint is stored on your system until you delete the subscribed content library from the inventory. - Click Next.

5. Optional: On the **Apply security policy** page, select **Apply Security Policy** and select **OVF default policy**.


Note: Currently, only the OVF default policy is supported.

When you apply the OVF default policy to a content library:

- All OVF items are validated by signing the OVF package with a valid certificate.
- All OVF or OVA items in the content library will have a **Yes** or **No** **Security Compliant** status based on the certificate validation.
- All items that are not OVF or OVA items will have a **Yes** **Security Compliant** status.

If the OVF default policy is not applied to a content library, all items in the content library will have a **Yes** **Security Compliant** status.

You can view the **Security Compliant** status under the **Templates** tab of a content library.



6. Click **Next**.
7. On the **Add storage** page, select a datastore as a storage location for the content library contents and click **Next**.
8. On the **Ready to complete** page, review the details and click **Finish**.

Edit a Content Library

In the vSphere Client, you can edit a local library by changing its name, description, or tags. You can also change the configuration settings of a local or subscribed content library.

The privileges that you need depend on the task that you want to perform.

Task	Privilege
Edit Local Content Library Settings	Content library > Update library and Content library > Update local library on the library.
Edit Subscribed Content Library Settings	Content library > Update subscribed library and Content library > Probe subscription information on the subscribed library.
Delete Content Library	Content library > Delete subscribed library or Content library > Delete local library depending on the type of library that you want to delete.

To share the contents of a local library across multiple vCenter Server instances, you must enable publishing for the library. From the **Edit Setting** dialog box, you can obtain the URL of your library and send it to other users to subscribe.

If a library is already published, you can change its password for authentication. To keep access to the published library, users who are subscribed to your library must update the password.

For subscribed libraries, you can change the download method or enable automatic synchronization with the published library.

1. Navigate to **Menu > Content Libraries**.
2. Right-click a content library and select the action that you want to perform.
 - To edit the content library description, select **Edit notes**.
 - To change the name of the content library, select **Rename**.
 - To assign a tag to the content library, select **Tags > Assign Tag**.
 - To remove a tag from the content library, select **Tags > Remove Tag**.

For information about tags and tag categories, see the *vCenter Server and Host Management* documentation.

- To delete the content library, select **Delete**.
3. To edit the configuration settings of a content library, right-click the content library and click **Edit Settings**.
The changes that you can make depend on the type of content library that you edit.

Content Library Type	Action
Local content library that is unpublished	To share the contents of a local library with other users, you can enable the publishing of the local library. 1. To publish the local library and share its contents with other users, select the Enable publishing check box. 2. To obtain the URL of your library and distribute it, click the Copy Link button. 3. Select Enable user authentication for access to this content library and set a password for the library. If you protect the library with a password, you must provide both the URL and the password to users who want to subscribe to your library. 4. To activate or deactivate the security policy, select or deselect the Apply Security Policy. When you deactivate the security policy of a content library, you cannot reuse the existing OVF items. 5. Click OK.
Local content library that is published	You can change the following settings of a local library that is published. • You can copy the subscription URL to your library and send it to other users to subscribe. • You can unpublish the library by deselecting the Enable publishing check box. Users who are currently subscribed to this library can no longer synchronize to the library, but they can still use the previously synchronized content. • You can activate or deactivate authentication for the library. • In the vSphere Client, you can change the password for authentication if the library is published and password-protected. a. Enter the current password and the new password. Confirm the new password. b. Click OK. • To activate or deactivate the security policy, select or deselect the Apply Security Policy. When you deactivate the security policy of a content library, you cannot reuse the existing OVF items.
Subscribed content library	You can change the following settings of a subscribed library: • Activate or deactivate the automatic synchronization with the published library. • Update the password for authentication to the published library. • Select a download method. You can either download all library content immediately or download the library content only when needed. If you switch from the option to download content only when needed to the option to download all library content immediately, a synchronization task starts and content

Content Library Type	Action
	starts downloading. The number and size of items in the published library determine the amount of time and network bandwidth that the task requires. To activate or deactivate the security policy, select or deselect the Apply Security Policy. When you deactivate the security policy of a content library, you cannot reuse the existing OVF items.

Configure Advanced Content Library Settings

You can monitor and configure the content library service settings across different vCenter Server instances. This advanced configuration ensures the consistency in your environment.

Verify that you are a member of the **SystemConfiguration.Administrators** group in the vCenter Single Sign-On domain.

1. Navigate to **Menu > Content Libraries**.
2. From the **Content Libraries** pane, select a library and click **Advanced**.
The **Advanced Configuration** dialog box opens.

3. Configure the parameters.

Configuration Parameter	Description
Library Auto Sync Enabled	Enables automatic synchronization of subscribed content libraries.
Library Auto Sync Refresh Interval (minutes)	Interval between two consequent automatic synchronizations of the subscribed content library. Measured in minutes.
Library Auto Sync Setting Refresh Interval (seconds)	Interval after which the refresh interval for the automatic synchronization settings of the subscribed library will be updated if it has been changed. Measured in seconds. If you change the refresh interval, you must restart the vCenter Server system.
Library Auto Sync Start Hour	The time of the day when the automatic synchronization of a subscribed content library starts.
Library Auto Sync Stop Hour	The time of the day when the automatic synchronization of a subscribed content library stops. Automatic synchronization stops until the start hour.
Library Maximum Concurrent Sync Items	Maximum number of concurrently synchronizing library items for each subscribed library.
Max concurrent NFC transfers per ESX host	Maximum concurrent NFC transfers limit per ESXi host.
Maximum Bandwidth Consumption	Bandwidth usage threshold across all transfers, measured in Mbps where 0 means unlimited bandwidth.
Maximum Number of Concurrent Priority Transfers	Concurrent transfer limit for priority files. If exceeded, transfers are queued. This threadpool is used only to transfer priority objects. For example, OVF. If you change the concurrent transfer limit for priority files, you must restart the vCenter Server system.
Maximum Number of Concurrent Transfers	Concurrent transfer limit. If exceeded, transfers are queued. If you change the concurrent transfer limit, you must restart the vCenter Server system.

4. Click **Save**.
5. From the **vCenter Server** drop-down menu, select the **vCenter Server** instance whose settings you want to configure.

Note: The **vCenter Server** drop-down menu appears only if your environment has more than one vCenter Server instance.

6. Edit the configuration parameters and click **Save**.
7. To apply the changes that require a content library service restart, click **Restart from VAMI**.

The **vCenter Server** Management Interface opens, and you can log in with your credentials.

Create and Delete a Subscription for a Local Library

To share the contents of a local library with users from other vCenter Server instances, enable publishing for the library. This way, other users can subscribe to the published library and use its contents.

As an administrator, you can also create subscriptions for the library to gain control over the distribution of content.

Note: You cannot create subscriptions to a local library that does not have publishing enabled. You also cannot deactivate publishing for a library that already has subscriptions. To deactivate publishing for a local library, you must first delete all of its subscriptions.

To avoid name collisions and failures when you publish content from the publisher library, the publisher and subscriber libraries must have unique folders dedicated to them.

Local Libraries

You use a local library to store and manage items in a single vCenter Server instance. When you enable publishing for the library, users from other vCenter Server instances can subscribe to it and gain access to the library items. In this case, only the subscriber initiates and controls the synchronization of content between the published library and the subscribed library.

If publishing is enabled, you can also create subscriptions to the local library. Creating a subscription creates a new subscribed library or links the publisher library to an existing subscriber. Unlike regular subscribed libraries, subscriptions give the administrator of the local library control over the distribution of content library items.

A publisher library is a local library with subscriptions.

Using Subscriptions

When you create a subscription for a local library, the result is a subscribed library. A publisher library is aware of its subscriptions. Subscriptions enable the administrator of the publisher library to control the content distribution. With subscriptions, content is distributed either when the subscriber initiates synchronization, or when the administrator of the local library publishes the library items to one or more of the existing subscriptions.

When you use subscriptions, you have the flexibility to decide how much of the library content you want to share with the subscribers. For example, you can publish some or all library items. You can also publish content to selected subscribers or to all subscribers.

The use of subscriptions allows content distribution between a publisher and a subscriber in the following scenarios.

- The publisher and subscriber are in the same vCenter Server instance.
- The publisher and subscriber are in vCenter Server instances that are in Enhanced Linked Mode.
- The publisher and subscriber are in vCenter Server instances that are in Hybrid Linked Mode. For more information about Hybrid Linked Mode, see the *VMware Cloud AWS* documentation.

Note: Publishing content is possible only from an on-premises publisher to a cloud subscriber, and not in the reverse scenario.

Limitations in Content Distribution

Content is distributed when a publisher library publishes content to its subscribers or when a subscriber synchronizes the content in the respective subscribed library with the published library. You can publish and synchronize a single content library item or an entire library. The following limitations exist in the distribution of content between a published and a subscribed library.

- You can publish only VM templates. If you publish an entire library that contains both VM templates and OVF templates, only the VM templates are replicated to the subscriber. To synchronize OVF templates and other types of files, the subscriber must initiate the synchronization.
- You can synchronize only OVF templates. If a subscriber initiates synchronization with a published library that contains both VM templates and OVF templates, only the OVF templates are synchronized in the subscribed library. VM templates are synchronized when a publisher library publishes them to its subscribers.

Create a Subscription for a Local Library

Subscriptions enable you to publish library items to a subscriber whenever you want. Create a subscription for a publisher library to control the distribution of templates to the subscriber.

- Verify that publishing is enabled for the library.
- Required privilege: **Content library > Create a subscription for a published library** on the content library for which you create a subscription.

When you create a subscription, you can link the publisher library to an existing subscribed library or create a new subscribed library. Creating a subscription to a new subscribed library triggers automatic synchronization. Creating a subscription to an existing library does not trigger automatic synchronization. If you create a subscription to an existing library, the synchronization happens when you publish an item or the entire library.

When you create a subscription to a new subscribed library and you select the option to download library content only when needed, only the metadata for the publisher library contents is downloaded to the associated storage. When the subscriber needs to use a library item, you either publish the item to the subscriber or the subscriber synchronizes the item to download its content to their local storage. For more information about synchronizing an entire library or a library item, see [Synchronize the Contents of a Subscribed Library](#).

1. Select **Menu > Content Libraries**.
2. Right-click a local library and select **New Subscription**.
The **Create Subscription** wizard starts.
3. Select your task.

Task	Steps
Create a new subscription to a new subscribed library	1. On the Select subscription type page, select the Create a new subscription to a new Subscriber library radio button and click Next. 2. On the Configure subscription page, configure the subscription and click Next. – Select a vCenter Server instance to create the subscribed library in. – Enter a name and description for the new subscribed library. – Select the download method for the new subscribed library. 3. On the Select folder page, select a location for the library contents and click Next. 4. On the Select compute resource page, select the compute resource for the subscribed library and click Next. 5. On the Select storage page, select a storage location for the subscribed library items and click Next. 6. On the Select network page, select a network for the synchronization of the subscribed library items and click Next.

Task	Steps
	7. On the Review page, review the configuration and click Finish .
Create a new subscription to an existing subscribed library	1. On the Select subscription type page, select the Create a new subscription to an existing Subscriber library radio button and click Next. 2. On the Configure subscription page, configure the subscription and click Next. a. From the vCenter Server drop-down menu, select the vCenter Server instance to create the subscription in. A list of all libraries that are subscribed to the publisher and that are in the selected vCenter Server instance appears. b. Select an existing subscribed library from the list. 3. On the Select folder page, select a location for the subscription and click Next. 4. On the Select compute resource page, select the compute resource for the subscriber library. 5. On the Select storage page, select where to store the contents of the subscriber and click Next. 6. On the Select network page, select a network for the synchronization of the subscribed library items use and click Next. 7. On the Review page, review the configuration and click Finish.

A new subscription is created. When you go to the **Subscriptions** tab of the local library, you can view a list of all existing subscriptions for the library.

Delete a Subscription from a Local Library

Delete a subscription if you no longer want to share the contents of a publisher library with a subscriber. You might also need to delete a subscription if the corresponding subscribed library has been deleted.

Required privilege: **Content library > Delete subscription of a published library** on the library.

If you do not delete a subscription to a deleted subscribed library and you publish content to that subscriber, the task triggers an error. To avoid errors, always delete subscriptions that link to deleted subscribed libraries.

Deleting is a useful option when you need to change a subscription. Because editing is currently not supported, your only choice is to delete the subscription that you want to modify and create a new one.

When you delete a subscription, the respective subscribed library and its contents are not deleted. Deleting a subscription means that you can no longer publish templates to the subscriber. The subscriber can still initiate the synchronization of OVF templates, if any. The subscriber can also continue using the previously synchronized content.

1. Navigate to the **Content Libraries** list.
2. Open a local library by clicking its name.
3. On the **Subscriptions** tab, select one or multiple subscriber libraries from the list.
4. Click the **Delete subscription** button to delete the selected subscriber libraries.

The subscription is deleted and you cannot use the publish function to synchronize the content in the respective subscribed library .

Publish the Library Contents and a Single Template to a Subscriber

Publish a library to replicate all VM templates stored in the publisher library to one or multiple subscribers. If you want to replicate a single VM template from a publisher library to a subscriber, you can publish the item, and not the entire library.

Required privileges on the content library:

- **Content library > Publish a library to its subscribers**
- **Content library > Sync library item**

If the publisher library contains OVF templates and other types of files, publishing the entire library to a subscription only replicates or updates the VM templates.

1. Navigate to the **Content Libraries** list.
2. Open a local library by clicking its name.
3. Publish a content library contents or a VM template to a subscriber.

Action	Description
Publish a content library contents to a subscriber	1. On the Subscriptions tab, select the subscriber libraries to update. 2. Click the Publish button. 3. To confirm the process, in the Publish Library pop-up window, click Publish. A publish task appears in the Recent Tasks pane. All VM templates from the publisher library are published to the selected subscribers.
Publish a single template to a subscriber	1. On the Templates tab, right-click a template of the VM template type and select Publish. 2. In the Publish Template dialog box, select the subscribers to which to publish the selected template. 3. Click OK. After the publishing finishes, the item content and metadata are downloaded to the storage of the subscribed library. On the Templates tab for the subscription, the value for the item in the Stored Content Locally column changes to Yes.

Synchronize the Contents of a Subscribed Library

You create a subscribed library to subscribe to a published library. You cannot add content to a subscribed library. You can only synchronize the content of the subscribed library with the content of the published library.

You can create the subscribed library in the same vCenter Server instance where the published library is, or in a different vCenter Server system. When you create a subscribed library, you can download all the contents of the published library immediately after the subscribed library is created. Alternatively, you can download only the metadata for the items from the published library and later download the full content of only those items that you need.

To ensure that the contents of a subscribed library are up-to-date, the subscribed library automatically synchronizes to the source published library at regular intervals. You can also manually synchronize a single item or an entire subscribed library.

When you update VM templates on the publisher library through the check-in and check-out operations, the VM templates are available in the subscriber library after you publish the subscriber library from the **Subscriptions** tab of the publisher.

The vertical timeline view is not available in the subscriber library. You can monitor only the latest version of the VM template.

For information about the VM templates management and the vertical timeline view, see [Manage VM Templates in a Content Library](#).

Download Methods for Synchronization

When you create a subscribed library, you can use the option to download content from the source published library immediately or only when needed to manage your storage space.

- When you synchronize a subscribed library that is configured to download all the contents of the published library immediately, the process synchronizes both the item metadata and the item contents. During synchronization, the library items that are new for the subscribed library are fully downloaded to the subscribed library storage. If some items are deleted from the published library, their contents remain at the storage location of your subscribed library, and you have to manually delete them.
- When you synchronize a subscribed library that is configured to download contents only when needed, the process synchronizes only the metadata for the library items from the published library, and does not download the contents of the items, which saves storage space. If you must use a library item, you need to synchronize that item. Synchronizing a library item downloads the full content of that item to your storage. When you no longer need the item, you can delete the item contents to free storage space.

The supported download methods depend on the source library. For more information, see the following table.

Table 156: Source Objects to Which You Can Subscribe by Creating a Subscribed Library in the vSphere Client

Source Object	Download library content immediately	Download library content when needed
A library running in a vCenter Server 6.x instance	Supported	Supported
A catalog running in a vCloud Director 5.5 and later instance	Supported	Not supported
A third-party library	Supported for third-party libraries that require authentication, if the username of the third-party library is <code>vcsp</code> . If the username of the source third-party library is different than <code>vcsp</code> , you can subscribe to it by using the VMware vCloud Suite API.	Supported for third-party libraries that require authentication, if the username of the third-party library is <code>vcsp</code> . If the username of the source third-party library is different than <code>vcsp</code> , you can subscribe to it by using the VMware vCloud Suite API.

Synchronize a Subscribed Content Library

To ensure that your subscribed library displays the latest content of the published library, you can manually initiate a synchronization task.

Required privilege: **Content library > Sync subscribed library** on the library.

You can also have subscribed libraries automatically synchronize with the content of the published library. Automatic synchronization requires a lot of storage space, because you download full copies of all the items in a published library.

To enable automatic synchronization, you must change the library settings. For information about changing content library settings, see [Edit a Content Library](#).

1. Navigate to the **Content Libraries** list.
2. Right-click a subscribed library and select **Synchronize**.

A new task for synchronizing the subscribed library appears in the **Recent Tasks** pane. After the task is complete, you can see the updated list with library items on the **Templates** and **Other Types** tabs.

Synchronize a Library Item in a Subscribed Library

To update or download the content of a library item in a subscribed library, you can synchronize the library item.

Required privilege: **Content library > Sync library item** on the library item.

When you create a subscribed library, only the metadata for the library contents is downloaded to the associated storage if you selected the option to download library content only when needed. When you need to use a library item, you synchronize it to download its content to your local storage.

When you no longer need the item, you can delete the content of the item to free storage space. You continue to see the item in your subscribed library, but it no longer takes up space on your storage because only the item's metadata remains on the storage. For information about deleting an item, see [Delete a Content Library Item](#).

1. Navigate to the **Content Libraries** list.
2. Select a subscribed library from the list.
3. Synchronize the item that you need to use.
 - On the **Templates** tab, right-click a template, and select **Synchronize Item**.

Note: Synchronization is not available for VM templates. You can only synchronize OVF templates.
 - On the **Other Types** tab, right-click an item, and select **Synchronize Item**.

After the synchronization finishes, the item content and metadata are downloaded to the storage of the subscribed library. On the **Templates** tab for the subscribed library, the value for the item in the **Stored Content Locally** column changes to Yes.

How to Populate Libraries with Content

You can populate a content library with VM templates and OVF templates that you can use to provision new virtual machines. You can also add other files to a content library, such as ISO images, scripts, and text files.

You can populate a library with items in several ways.

Import Items to a Content Library

You can add items to a local content library by importing files from your local machine or from a Web server. You can import OVF and OVA templates and other types of files, such as ISO images, certificates, and so on. You can keep the items in the library and share them with other users across multiple vCenter Server instances. You can also use the templates in the content library to deploy new virtual machines and vApps.

Required privilege: **Content library > Add library item** and **Content library > Update files** on the library.

1. Navigate to the **Content Libraries** list.
2. Right-click a local content library and select **Import Item**.
The **Import Library Item** dialog box opens.
3. In the **Source** section, choose the source of the item.

Option	Description
Import from a URL	Enter the path to the Web server where the item is. Note: You can import either an <code>.ovf</code> or <code>.ova</code> file. The resulting content library item is of the OVF Template type.
Import from a Local File	Click Browse to navigate to the file that you want to import from your local system. You can use the drop-down menu to filter files in your local system.

Option	Description
	Note: You can import either an <code>.ovf</code> or <code>.ova</code> file. When you import an OVF template, first select the OVF descriptor file (<code>.ovf</code>). Next, you are prompted to select the other files in the OVF template, for example the <code>.vmdk</code> file. The resulting content library item is of the OVF Template type.

vCenter Server reads and validates the manifest and certificate files in the OVF package during importing. A warning is displayed in the **Import Library Item** wizard, if certificate issues exist, for example if vCenter Server detects an expired certificate.

Note: vCenter Server does not read signed content, if the OVF package is imported from an `.ovf` file from your local machine.

4. In the **Destination** section, enter a name and a description for the item.
5. Click **Import**.

In the **Recent Tasks** pane you see two tasks, one about creating a new item in the library, and the second about uploading the contents of the item to the library. After the task is complete, the item appears on the **Templates** tab or on the **Other Types** tab.

Clone a vApp to a Template in a Content Library

You can clone existing vApps to vApp templates in a content library. You can use the vApp templates later to provision new vApps on a cluster or a host in your vSphere inventory. The vApp is exported to a content library in the OVF format.

1. In the vSphere Client, select **Menu > VMs and Templates**.
2. Right-click a vApp and select **Clone > Clone to Template in Library**.
The **Clone to Template in Content Library** wizard opens.
3. On the **Basic information** page, configure the content library template and click **Next**.
 - a) Select the **New template** radio button.
 - b) Enter a name and, optionally, a description for the template.
 - c) Optional: Select the **Preserve MAC-addresses on network adapters** check box to preserve the MAC addresses of the network adapters.
 - d) Optional: Select **Include extra configuration** to include vApp-related configuration in the template that you clone.
4. On the **Location** page, select a content library to clone the vApp to.
5. On the **Review** page, review the configuration details and click **Finish**.

A new task for cloning to OVF package appears in the Recent Tasks pane. After the task finishes, the vApp template appears on the **Templates** tab for the content library.

Use the template to provision vApps on a host or a cluster in your vSphere inventory. See [Create New vApp From a Template in a Content Library](#).

Clone a Virtual Machine or a Virtual Machine Template to a Template in a Content Library

You can add new templates to a content library by cloning virtual machines or virtual machine templates from your vCenter Server inventory to templates in the content library. You can use the content library items later to provision

virtual machines on a cluster or a host. You can also update an existing template in the content library by cloning a virtual machine or virtual machine template from the vCenter Server inventory.

Templates are primary copies of virtual machines that you can use to create virtual machines that are ready for use. You can change the template, such as installing additional software in the guest operating system, while preserving the state of the original template.

When you clone a virtual machine from the vCenter Server inventory as a template to the content library, you can choose what type of content library item to create. You can choose to create a library item of either the VM Template type or OVF Template type.

Important: If you choose to create a VM Template library item, an identical VM template is created in the vCenter Server inventory.

For more information, see [Templates in Content Libraries](#).

1. Navigate to the virtual machine or template that you want to clone.
2. Select your task.

Option	Description
Clone a virtual machine	1. Right-click the virtual machine and select Clone > Clone as Template in Library. The Clone Virtual Machine To Template wizard opens. 2. On the Basic information page, enter a name and description for the template, select the template type, and select an inventory folder for the template. You can create an OVF Template or VM Template in the content library. 3. On the Location page, select a local content library in which you want to add the template. 4. On the Select a compute resource page, select the compute resource for the template. 5. On the Select storage page, select the storage for the template disk and configuration files. 6. On the Review page, review the details and click Finish to complete the cloning task.
Clone a virtual machine template	1. Right-click the virtual machine template and select Clone to Library. The Clone to Template in Library dialog box opens. 2. Select the Clone as option. You can create a template or you can choose an existing template to update. 3. From the content libraries list, select the library in which you want to add the template. 4. Enter a name and description for the template. 5. Select the configuration data that you want to include in the template. You can select to preserve the MAC-addresses on the network adapters and include extra configuration. 6. Click OK.

A new task for cloning appears in the **Recent Tasks** pane. After the task is complete, the template appears in the **Templates** tab for the content library. You can view the type of template in the **Type** column.

Use the template to create virtual machines on hosts or clusters in the vSphere inventory.

Clone Library Items from One Library to Another Library

You can clone a template from one content library to another in the same vCenter Server instance. The cloned template is an exact copy of the original template.

When cloning a template between libraries, you can select the source library to also be a destination library in the clone wizard.

A subscribed library can be the source of an item you want to clone, but you cannot clone items to a subscribed library. The subscribed libraries are filtered out from the list with destination libraries in the Clone Library Item dialog box. When the source library of an item you want to clone is a subscribed library with the setting to download items only when needed, the item is first downloaded to the source subscribed library and then cloned to the destination library.

1. Navigate to the **Content Libraries** list.
2. Click a content library and click the **Templates** tab.
3. Right-click a template and select **Clone Item**.
The **Clone Library Item** dialog box opens.
4. Optional: Change the name and notes for the item you clone.
5. From the list of content libraries, select the library in which you want to clone the template and click **OK**.
You can select the destination library to be the same as the source library if you want to have identical copy of the template in the same library.

A new task for cloning the template appears in the Recent Tasks pane. After the task is complete, a clone of the template appears on the **Templates** tab of the destination content library.

Deploy a virtual machine from template on a host or a cluster in your vSphere inventory.

Manage Content Library Items

You can perform various tasks with the items in a content library. For example, you can synchronize an item from a subscribed library to download all its contents and use the item to deploy a virtual machine. You can delete items you no longer need to use, and so on.

Each template or other type of file in a content library is a library item. An item can contain a single file or multiple files. For example, when you add an OVF template to the library, you actually upload all the files that are associated with the template, but in the vSphere Client you only see one library item of the OVF Template type.

Update a Content Library Item

Managing and keeping your virtual environment up-to-date might require you to update the content of a library item. For example, you can directly update a template when you want to add a patch to it, instead of deleting the existing template and creating a new one.

Verify that you have the Content Library Administrator role.

Note: You can update only templates from the **OVF and OVA Templates** type. Updating templates from the **VM Templates** type is not supported.

You cannot update the contents of a subscribed library. In local and published libraries, you can update only templates of the OVF Template type.

1. Navigate to the **Content Libraries** list.
2. Click a content library and select the file that you want to update.
 - From the **Templates** tab, right-click a template from the library, and select **Update Item**.
 - From the **Other Types** tab, right-click a file from the library that is not a template, and select **Update Item**.

The **Update Library Item** dialog box opens.

3. In the **Source** section, select a file to overwrite the item in your library with.

Option	Description
URL	Enter the URL to a web server where the item is stored .
Browse	Navigate to an item that is stored on your local system.

4. Optional: In the **Destination** section, change the name of the item, the description of the item, or both. Click **OK**.

The content of the item is updated. On the **Summary** tab for the item, you can view the time of the last update of the item.

Export an Item from a Content Library to Your Local Computer

You might need to export an item from a content library to your local system.

Required privilege: **Content library > Download files** on the library.

Note: You can export only items from the **OVF and OVA Templates** type. Exporting items from the **VM Templates** type is not supported.

1. Navigate to the **Content Libraries** list.
2. Select a content library.
3. Select the type of file you want to export.
 - From the **Templates** tab, right-click a template from the library, and select **Export Item**.
 - From the **Other Types** tab, right-click a file from the library that is not a template, and select **Export Item**.
4. In the **Export Library Item** dialog box click **OK**.
5. If you are exporting an OVF template, you are prompted to save each file associated with the template to the browser download location (for example, `.vmdk` and `.mf` files).

Clone Library Items from One Library to Another Library

You can clone a template from one content library to another in the same vCenter Server instance. The cloned template is an exact copy of the original template.

When cloning a template between libraries, you can select the source library to also be a destination library in the clone wizard.

A subscribed library can be the source of an item you want to clone, but you cannot clone items to a subscribed library. The subscribed libraries are filtered out from the list with destination libraries in the Clone Library Item dialog box. When

the source library of an item you want to clone is a subscribed library with the setting to download items only when needed, the item is first downloaded to the source subscribed library and then cloned to the destination library.

1. Navigate to the **Content Libraries** list.
2. Click a content library and click the **Templates** tab.
3. Right-click a template and select **Clone Item**.
The **Clone Library Item** dialog box opens.
4. Optional: Change the name and notes for the item you clone.
5. From the list of content libraries, select the library in which you want to clone the template and click **OK**.
You can select the destination library to be the same as the source library if you want to have identical copy of the template in the same library.

A new task for cloning the template appears in the Recent Tasks pane. After the task is complete, a clone of the template appears on the **Templates** tab of the destination content library.
Deploy a virtual machine from template on a host or a cluster in your vSphere inventory.

Edit a Content Library Item

Edit a library item to change its name, description or tag properties.

Required privileges on the library:

- **Content library > Update library item**
- **Content library > Update local library**

You can edit items only in a local library, regardless of whether it is published or not. Library items in subscribed libraries cannot be modified.

You can edit both VM templates and OVF templates.

1. Navigate to the **Content Libraries** list.
2. Open a local library by clicking its name.
3. Navigate the library item to edit.
 - To edit a template, click the **Templates** tab.
 - To edit another type of file, click the **Other Types** tab.
4. Right-click the item and select your task from the context menu.
 - To edit the description of the item, select **Edit Notes**.
 - To rename the item, select **Rename**.
 - To assign a tag to the item, select **Tags > Assign Tag**.
 - To remove a tag from the item, select **Tags > Remove Tag**.

Delete a Content Library Item

If you use a subscribed library, and you synchronize it, you can later delete the library from storage but keep the metadata. You can also completely delete a library item such as a template.

Required privileges

Task	Required Privileges
Delete the contents of a library item	Content library > Evict library item

Task	Required Privileges
Delete a library item	Content library > Delete library item

If a subscribed library is created with the option to download library content only when needed, only metadata for the library items is stored in the associated with the library storage. When you want to use a library item, for example use a VM template to deploy a virtual machine, you have to synchronize the item. Synchronization downloads the entire content to the associated storage.

In the vSphere Client you can delete an item altogether.

1. Select **Menu > Content Libraries**.
2. Click a content library, select the type of item, and select the task you want to perform with the item.
3. From the **Templates** tab, right-click a template from the library, and select **Delete**.
4. From the **Other Types** tab, right-click a file from the library that is not a template, and select **Delete**.

Creating Virtual Machines and vApps from Templates in a Content Library

You can deploy virtual machines and vApps from the VM or OVF templates that are stored in a content library.

The library can be a local library to the vCenter Server instance where you want to deploy the VM or the vApp template, or can be a subscribed library to that vCenter Server instance.

The VM Template type is only supported in the vSphere Client. You can deploy virtual machines from VM Templates in a content library only in the vSphere Client.

Note: You can also use the API calls to create and manage VM templates in a content library.

The use of templates results in consistency, compliance, and efficiency when you deploy virtual machines and vApps in your data center.

Deploy a Virtual Machine from an OVF Template in a Content Library

In content libraries, you can use the OVF template, which is either the template of a virtual machine or a vApp, to deploy a virtual machine to a host or a cluster in your vSphere inventory.

Verify that you have the following privileges:

- **Virtual machine > Inventory > Create from existing**
- **Virtual machine > Configuration > Add new disk**
- **Virtual machine > Provisioning > Deploy template**
- **Virtual machine > Assign virtual machine to resource pool**
- **vApp > Import**
- **vApp > Create**

1. Navigate to **Menu > Content Libraries**.
2. Select a content library and click the **Templates** tab.
3. Right-click an OVF template and select **New VM from This Template**.
The **New Virtual Machine from Content Library** wizard opens.

4. On the **Select a name and folder** page, enter a name, select a location for the virtual machine, and select additional customization options for the new virtual machine.
 - a) To customize the guest operating system of the virtual machine, select the **Customize the operating system** check box.
 - b) To change the hardware settings of the virtual machine, select the **Customize this virtual machine's hardware** check box.

A new **Customize hardware** page appears if you want to customize the virtual machine's hardware.

5. Click **Next**.
6. On the **Select a compute resource** page, select a host, a cluster, a resource pool, or a vApp where to run the deployed template, and click **Next**.

Important:

If the template that you deploy has an NVDIMM device and virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, you cannot proceed with the task.

If the template that you deploy does not have an NVDIMM device, but it has virtual PMem hard disks, the destination host or cluster must have available PMem resource. Otherwise, all the hard disks of the virtual machine will use the storage policy and datastore selected for the configuration files of the source template.

7. On the **Review details** page, verify the template details and click **Next**.
8. On the **Select storage** page, select the datastore or datastore cluster in which to store the virtual machine configuration files and all of the virtual disks. Click **Next**.

Option	Description
Deploy a virtual machine from a template that has vPMem hard disks	1. Choose the type of storage for the template by selecting the Standard, the PMem, or the Hybrid radio button. If you select the Standard mode, all virtual disks will be stored on a standard datastore. If you select the PMem mode, all virtual disks will be stored on the host-local PMem datastore. Configuration files cannot be stored on a PMem datastore and you must additionally select a regular datastore for the configuration files of the virtual machine. If you select the Hybrid mode, all PMem virtual disks will remain stored on a PMem datastore. Non-PMem disks are affected by your choice of a VM storage policy and datastore or datastore cluster. For more information about persistent memory and PMem storage, see the <i>vSphere Resource Management</i> guide. 2. From the VM Storage Policy drop-down menu, select a virtual machine storage policy or leave the default one. 3. Select a datastore or a datastore cluster. 4. Select the Disable Storage DRS for this virtual machine check box if you do not want to use storage DRS with the virtual machine. 5. Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.
Deploy a virtual machine from a template that does not have vPMem hard disks	1. Select the disk format for the virtual machine virtual disks. Same format as source uses the same disk format as the source virtual machine. The Thick Provision Lazy Zeroed format creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out later, on demand, on first write from the virtual machine.

Option	Description
	Thick Provision Eager Zeroed is a type of thick virtual disk that supports clustering features such as Fault tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types of disks. The Thin Provision format saves storage space. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can grow to the maximum capacity allocated to it. 2. Select a VM storage policy or leave the default one. 3. Select a datastore or a datastore cluster. 4. Turn on the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.

Note: If you want to use the API calls to deploy an OVF template that contains vPMem hard disks and that has been exported from a content library, consult <https://kb.vmware.com/kb/52370>.

9. On the **Select networks** page, select a network for each network adapter in the template and click **Next**.

10. On the **Customize template** page, change the template settings and click **Next**.

11. On the **Customize hardware** page, configure the virtual machine's hardware options and click **Next**.

Note: You cannot change the existing virtual hardware disks, CD/DVD drives settings, and you cannot add new virtual disks or CD/DVD drives. You can change these settings after you deploy the virtual machine.

12. On the **Ready to complete** page, review the page and click **Finish**.

A new task for creating the virtual machine appears in the Recent Tasks pane. After the task is complete, the new virtual machine is created on the selected resource.

Deploy a Virtual Machine from a VM Template in a Content Library

In the vSphere Client, you can use a content library item of the VM Template type to deploy a virtual machine to a host or cluster in your vSphere environment.

Verify that you have the following privileges:

- **Virtual machine > Inventory > Create from existing**
- **Virtual machine > Configuration > Add new disk**
- **Virtual machine > Provisioning > Deploy template**
- **Virtual machine > Assign virtual machine to resource pool**
- **vApp > Import**
- **vApp > Create**
- To access customization options for Windows guest operating systems, Microsoft Sysprep tools must be installed on the vCenter Server system. The Sysprep Tool is built into the Windows Vista and Windows 2008 and later operating systems. For details about this and other customization requirements, see [Guest Operating System Customization Requirements](#).

Important: If the template that you deploy has an NVDIMM device and virtual PMem disks, the destination host or cluster must have an available PMem resource. Otherwise, you cannot proceed with the task.

If the template that you deploy does not have an NVDIMM device, but has virtual PMem hard disks, the destination host or cluster must have an available PMem resource. Otherwise, all hard disks of the virtual machine will use the storage policy and datastore selected for the configuration files of the source template.

For information about persistent memory and PMem storage, see the *vSphere Resource Management* guide.

For information how to configure the virtual machine hardware options, see [Configuring Virtual Machine Hardware](#) and [Configuring Virtual Machine Options](#)

Note: If you want to use the API calls to deploy an OVF template that contains vPMem hard disks and that has been exported from a content library, consult <https://kb.vmware.com/s/article/52370>.

1. Navigate to **Menu > Content Libraries**.
2. To open a content library, click its name.
3. On the **Templates** tab, right-click a VM Template and select **New VM from This Template**.
The **Deploy From VM Template** wizard opens.
4. On the **Select a name and folder** page, enter a name and select a location for the virtual machine.
5. On the **Select a compute resource** page, select a host, a cluster, a resource pool, or a vApp where to run the deployed VM template, and click **Next**.
6. On the **Select storage** page, select the datastore or datastore cluster in which to store the virtual machine configuration files and all virtual disks.

Option	Action
Deploy a virtual machine from a template that has vPMem hard disks	1. Select the type of storage for the template by clicking the Standard, the PMem, or the Hybrid radio button. — If you select the Standard mode, all virtual disks are stored on a standard datastore. — If you select the PMem mode, all virtual disks are stored on the host-local PMem datastore. Configuration files cannot be stored on a PMem datastore and you must additionally select a regular datastore for the configuration files of the virtual machine. — If you select the Hybrid mode, all PMem virtual disks remain stored on a PMem datastore. Your choice of a VM storage policy and datastore or datastore cluster affects the Non-PMem disks. 2. From the VM Storage Policy drop-down menu, select a virtual machine storage policy or leave the default one. 3. Select a datastore or a datastore cluster. 4. Select the Disable Storage DRS for this virtual machine check box if you do not want to use storage DRS with the virtual machine. 5. To select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk, enable the Configure per disk option. Note: To convert a PMem hard disk to a regular one, you can use the Configure per disk option, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.
Deploy a virtual machine from a template that does not have vPMem hard disks	1. Select the disk format for the virtual machine virtual disks. — The Same format as source option uses the same disk format as the source virtual machine. — The Thick Provision Lazy Zeroed format creates a virtual disk in a default thick format. The space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out later, on demand, or on first write from the virtual machine. — The Thick Provision Eager Zeroed format is a type of thick virtual disk that supports clustering features such as Fault Tolerance. The space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data

Option	Action
	remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create other types of disks. — The Thin Provision format saves storage space. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can extend to the maximum capacity allocated to it. 2. Select a VM storage policy or leave the default one. 3. Select a datastore or a datastore cluster. 4. Enable the Configure per disk option to select a separate datastore or a datastore cluster for the template configuration file and for each virtual disk. Note: You can use the Configure per disk option to convert a PMem hard disk to a regular one, but that change might cause performance problems. You can also convert a standard hard disk to a PMem hard disk.

7. Click **Next**.
8. On the **Select deploy options** page, apply a customization specification to the virtual machine and click **Next**.
You can choose to customize the guest operating system or the virtual machine hardware. You can also choose to power on the virtual machine after its creation.

Option	Action
Select an existing specification	Select a customization specification from the list.
Override	To change the guest customization specification for this deployment only, click Override , complete the steps in the Override VM Customization Specification wizard, and click OK .

9. Optional: On the **Customize guest OS** page, select a customization specification to apply to the virtual machine.
Customizing the guest operating system helps prevent conflicts that can result if you or other users deploy virtual machines with identical settings, such as duplicate computer names.
10. Optional: On the **User settings** page, specify the required settings for the virtual machine.
This page of the wizard appears only if the selected specification requires additional customization.
11. Optional: On the **Customize hardware** page, configure the virtual machine hardware and options and click **Next**.
You can leave the defaults and configure the virtual machine hardware and options later.
Important: If you chose to use PMem storage for the virtual machine, its default hard disk, the new hard disks that you configure, and the NVDIMM devices that you add to the virtual machine all share the same PMem resources. You must adjust the size of the newly added devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.
12. On the **Ready to complete** page, review the information and click **Finish**.

A new task for creating the virtual machine appears in the **Recent Tasks** pane. After the task finishes, the new virtual machine is created on the selected resource.

Create New vApp From a Template in a Content Library

You can use an OVF template to create a new vApp on a host or a cluster in your vSphere inventory.

Verify that you have the following privileges:

- **vApp > Create**

- **vApp > Import**

1. Navigate to **Menu > Content Libraries**.
2. Open a content library by clicking its name, and click the **Templates** tab.
3. Right-click a template of a vApp and select **New vApp from This Template**.
The **New vApp from Content Library** wizard opens.
4. On the **Select a name and folder** page, enter a name and select a location for the vApp, and click **Next**.
5. On the **Select a compute resource** page, select a host, a cluster, a resource pool, or a vApp to deploy the vApp to and click **Next**.
6. On the **Review details** page, verify the template details and click **Next**.
7. On the **Select storage** page, select disk format and a storage resource for the vApp.
8. On the **Select networks** page, select a destination network for each source network.
9. On the **Ready to complete** page, review the configurations you made for the vApp, and click **Finish**.

A new task for creating the vApp appears in the **Recent Tasks** pane. After the task is complete, the new vApp is created.

Manage VM Templates in a Content Library

In vSphere 7.0 and later, you can edit the contents of the VM templates by checking them out, making the necessary changes, and checking them in.

You can track history of changes over time by using the vertical timeline view. The vertical timeline view provides you with detailed information about the different VM template versions, the updates that privileged users have made, and when the last change was made. By using the vertical timeline, you can revert VM templates back to their previous state or delete the previous version of a VM template.

In addition, you can deploy a virtual machine from the latest version of the VM template without any disruptions while it is checked out for update. You can update the virtual machine and check it back in into the same VM template.

Templates in Content Libraries

Templates are primary copies of virtual machines that you can use to deploy virtual machines that are customized and ready for use. Templates promote consistency throughout your vSphere environment. You can use the content library to store and manage templates of virtual machines and vApps. You can use VM templates and vApp templates to deploy virtual machines and vApps to a destination object, such as a host or a cluster.

Content libraries support two types of templates, the OVF Template type and the VM Template type.

In a content library, you can store and manage virtual machine templates as OVF templates or VM templates. vApps are always converted to OVF templates in the content library.

VM Templates in Content Libraries and in the vCenter Server Inventory

A VM template is a template of a virtual machine. You create a VM template by cloning a virtual machine into a template.

When you create a VM template in a content library, the library item is backed by a VM template in the vCenter Server inventory. The content library item and the corresponding inventory object are related in the following ways.

- If you convert the VM template in the vCenter Server inventory to a virtual machine, the corresponding VM template library item is also deleted.
- If you rename the VM template in the vCenter Server, the corresponding VM template library item is also renamed.
- If you rename the VM template library item the associated VM template in the vCenter Server inventory is also renamed.

- If you delete the VM template in the vCenter Server inventory, the corresponding VM template library item is also deleted.
- If you delete the VM template library item, the associated VM template in the vCenter Server inventory is also deleted.

OVF Templates in Content Libraries

In a content library, an OVF template is either a template of a virtual machine, or a template of a vApp. When you clone a virtual machine into a template in a content library, you choose whether to create an OVF template or a VM template. However, if you clone a vApp into a template in a content library, the resulting content library item is always an OVF template. Because the OVF format is actually a set of files, if you export the template, all the files in the OVF template library item (.ovf, .vmdk, .mf) are saved to your local system.

Differences Between VM Templates and OVF Templates in Content Libraries

You can use both VM templates and OVF templates to deploy new virtual machines in your vSphere environment. However, the two types of templates have different properties and support different deployment options.

See the following table for a detailed list of the differences between VM templates and OVF templates in a content library.

Table 157: VM Templates and OVF Templates Properties

Property	VM Templates in Content Library	OVF Templates in Content Library
Datastore	VM templates can be stored on any datastore that you have privileges to. Note: VM templates cannot be stored in a library that uses NFS or SMB storage.	OVF templates can only be stored on the datastore that is associated with the content library.
Footprint	The default one.	Compressed or Thin.
Host/Datastore Maintenance Mode	When the host becomes inaccessible, VM templates are automatically migrated to another host.	When either the host or the datastore becomes inaccessible, you must manually migrate the OVF templates to another host or datastore.
Associated with a Host	Yes.	No.
Storage DRS	Supported.	Not supported.
Cross-vendor Compatibility	Not supported.	Supported.
Software License Agreement	Not supported.	Supported.
Encryption	Supported. You can create encrypted VM templates.	Not supported. While OVF templates cannot be encrypted themselves, you can still deploy an encrypted virtual machine from an OVF template.
Deployment Options	During the deployment of a VM template, hardware customization and guest OS customization are both supported.	During the deployment of an OVF template, only guest OS customization is supported. Hardware customization is not supported.

The supported operations on a content library template are different depending on the template type. You can edit the settings for both OVF and VM templates. However, you can update, export, and clone a template only if it is an OVF template.

Check Out a Virtual Machine from a Template

In the vSphere Client, you can edit the VM templates and monitor the changes that have been made by other privileged users. You can perform the checkout operation to update a virtual machine from the VM template. During this process, the VM template is not available for checkout from other users, but they can deploy a virtual machine from the VM template without any disruptions.

Verify that you have the following privileges:

- **Content library > Check out a template**
- **Resource > Assign virtual machine to resource pool**
- **Datastore > Allocate space**
- **Virtual machine > Edit inventory > Create from existing**
- **Virtual machine > Change Configuration > Set annotation**
- If you want to power on the checked out virtual machine, verify that you have the **Virtual machine > Interaction > Power On** privilege.

When you check out a VM template, you cannot convert the virtual machine to a template or migrate the virtual machine to a different vCenter Server inventory.

1. To check out a VM template

Option	Action
From a content library	1. Navigate to Menu > Content Libraries. 2. To open a local library, click its name. 3. On the Templates tab, select a VM template and click the Check out VM from this template button.
From the vSphere Client inventory	1. Navigate to Menu > VMs and Templates and click the VM template. 2. Click the Versioning tab and in the vertical timeline view, click Check out VM from this template.

The **Check out VM from VM Template** dialog box opens.

2. On the **Name and location** page, enter a virtual machine name, select the virtual machine location, and click **Next**.
3. On the **Select compute resource** page, select the compute resource for the checked out virtual machine and click **Next**.
4. On the **Review** page, review the configuration.
5. Choose whether to power on the virtual machine after checkout by selecting the **Power on VM after checkout** check box.
6. Click **Finish**.

The checked out virtual machine appears in the selected location marked with a blue circle icon. You can perform the necessary configuration changes.

After you complete the virtual machine updates, you can check in the virtual machine back to the template.

Check In a Virtual Machine to a Template

After you check out a virtual machine from a template and update the virtual machine, you must check the virtual machine back into the VM template. When you check in the virtual machine to a template, you create a new version of the VM template containing the updated state of the virtual machine.

Verify that the virtual machine is powered off or suspended. You cannot check in a powered on virtual machine to a VM template.

Required privileges:

- **Content library > Check in a template**

When you check in the virtual machine to the VM template, you allow the deployment of the last changes that you make to the virtual machine.

1. To check in a virtual machine to a template:

Option	Action
From a content library	1. Navigate to Menu > Content Libraries. 2. To open a content library, click its name. 3. On the Templates tab, select a VM template and click Check in VM to template.
From the vSphere Client inventory	1. Navigate to Menu > VMs and Templates and click the VM template. 2. Click the Versioning tab and in the vertical timeline view, click Check in VM to template.

The **Check in VM** dialog box opens.

2. To describe the change, enter a comment in **Check in notes**.
3. Click **Check in**.

The updated version of the VM template appears in the vertical timeline. You can see the check-in comment, the name of the user who made the changes, and the date of the change.

Discard a Checked Out Virtual Machine

If you check out a VM template and make no updates to the virtual machine or perform an update that you do not want to keep, you can discard the checked out virtual machine. Each time you check in the virtual machine back to the template, you create a new version of the VM template. You can discard the checked out virtual machine to avoid creating new versions or to prevent other users from using a faulty version.

Required privileges:

- **Virtual machine > Inventory > Delete**

1. To discard a checked out virtual machine:

Option	Action
From a content library	1. Navigate to Menu > Content Libraries. 2. To open a local library, click its name. 3. On the Templates tab, select a VM template. 4. From the vertical timeline, click the horizontal ellipsis icon (⋮) that appears in the checked out VM template box and select Discard Checked Out VM.
From the vSphere Client inventory	1. Navigate to Menu > VMs and Templates and click the VM template. 2. Click the Versioning tab in the vertical timeline.

Option	Action
	3. Click the horizontal ellipsis icon (⋮) that appears in the checked out VM template box, and select Discard Checked Out VM .

The **Discard Checked Out VM** dialog box opens.

2. To delete the checked out virtual machine and discard all changes, click **Discard**.

You deleted the virtual machine from the inventory and discarded all changes.

Revert to a Previous Version of a Template

If the latest VM template contains changes that you no longer want to keep or you made a mistake during your last checkin, you can revert the VM template to the previous version.

Required privileges:

- **Content library > Check in a template**

1. To revert to a previous version of a template:

Option	Action
From a content library	1. Navigate to Menu > Content Libraries. 2. To open a local library, click its name. 3. On the Templates tab, select a VM template.
From the vSphere Client inventory	1. Navigate to Menu > VMs and Templates and click the VM template. 2. Click the Versioning tab.

2. From the vertical timeline, navigate to the previous state of the VM template, click the horizontal ellipsis icon (⋮), and select **Revert to This Version**.

The **Revert to Version** dialog box opens.

3. Enter a reason for the revert operation and click **Revert**.

The VM template that you revert to becomes the current VM template.

Delete a Previous Version of a VM Template

Delete a previous version of a VM template if you no longer want to allow the use of the template. Deleting a VM template removes the template and its content from the inventory.

Required privileges:

- **Content Library > Delete library item**

1. To delete a previous version of a template:

Option	Action
From a content library	1. Navigate to Menu > Content Libraries. 2. To open a local library, click its name. 3. On the Templates tab, select a VM template.
From the vSphere Client inventory	1. Navigate to Menu > VMs and Templates and click the VM template.

Option	Action
	2. Click the Versioning tab.

2. From the vertical timeline, navigate to the previous state of the VM template, click the horizontal ellipsis icon (⋮), and select **Delete Version**.
The **Confirm Delete** dialog box opens.
3. To delete permanently the VM template and its contents, click **Yes**.

Configuring Virtual Machine Hardware

You can add or configure most virtual machine hardware settings during virtual machine creation or after you create the virtual machine and install the guest operating system.

When you configure the virtual machine hardware, you can view the existing hardware configuration and add or remove hardware. You can change nearly every setting that was selected during virtual machine creation.

Virtual Machine Compatibility

When you create a virtual machine or upgrade an existing virtual machine, you use the virtual machine compatibility setting to select the ESXi host versions that the virtual machine can run on.

The compatibility setting determines the virtual hardware available to the virtual machine, which corresponds to the physical hardware available on the host. Virtual hardware includes BIOS and EFI, available virtual PCI slots, maximum number of CPUs, maximum memory configuration, and other characteristics. New virtual hardware capabilities are typically released once a year with major or minor releases of vSphere.

Each virtual machine compatibility level supports at least five major or minor vSphere releases. For example, a virtual machine with ESXi 6.5 and later compatibility can run on ESXi 6.7, ESXi 6.7 Update 2, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, ESXi 8.0 Update 3.

Table 158: Virtual Machine Compatibility Options

Compatibility	Description
ESXi 8.0 Update 2 and later	This virtual machine (hardware version 21) is compatible with ESXi 8.0 Update 2 and later.
ESXi 8.0 and ESXi 8.0 Update 1	This virtual machine (hardware version 20) is compatible with ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3.
ESXi 7.0 Update 3 and later	This virtual machine (hardware version 19) is compatible with ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, ESXi 8.0 Update 3.
ESXi 7.0 Update 2 and later	This virtual machine (hardware version 19) is compatible with ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3.
ESXi 7.0 Update 1 and later	This virtual machine (hardware version 18) is compatible with ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3.
ESXi 7.0 and later	This virtual machine (hardware version 17) is compatible with ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3.
ESXi 6.7 Update 2 and later	This virtual machine (hardware version 15) is compatible with ESXi 6.7 Update 2, ESXi 6.7 Update 3, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3.
ESXi 6.7 and later	This virtual machine (hardware version 14) is compatible with ESXi 6.7, ESXi 6.7 Update 2, ESXi 6.7 Update 3, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3.

Compatibility	Description
ESXi 6.5 and later	This virtual machine (hardware version 13) is compatible with ESXi 6.5, ESXi 6.7, ESXi 6.7 Update 2, ESXi 6.7 Update 3, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, ESXi 8.0 Update 3.

The compatibility setting that appears in the **Compatible with** drop-down menu is the default for the virtual machine that you are creating. The following factors determine the default virtual machine compatibility:

- The ESXi host version on which the virtual machine is created.
- The inventory object that the default virtual machine compatibility is set on, including a host, cluster, or data center.

You can accept the default compatibility or select a different setting. It is not always necessary to select the latest ESXi host version. Selecting an earlier version can provide greater flexibility and is useful in the following situations:

- To standardize testing and deployment in your virtual environment.
- If you do not need the capabilities of the latest host version.
- To maintain compatibility with older hosts.

When you create a virtual machine, consider the environment that the virtual machine runs in and weigh the benefits of different compatibility strategies. Consider your options for these scenarios, which demonstrate the flexibility inherent with each virtual machine compatibility selection.

Objects in Environment	Compatibility	Results
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 6.5 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 6.7 and later.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 6.7 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 7.0 and later. A virtual machine with such compatibility cannot run on ESXi 6.5.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 7.0 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 7.0 Update 1 and later. A virtual machine with such compatibility cannot run on ESXi 6.7 and earlier.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 7.0 Update 1 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 7.0 Update 2 and later. A virtual machine with such compatibility cannot run on ESXi 7.0 and earlier.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 7.0 Update 2 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 7.0 Update 3 and later. A virtual machine with such compatibility cannot run on ESXi 7.0 Update 1 and earlier.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 7.0 Update 3 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 8.0 and later. A virtual machine with such compatibility cannot run on ESXi 7.0 Update 2 and earlier.

Objects in Environment	Compatibility	Results
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 8.0 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 8.0 Update 1 and later. A virtual machine with such compatibility cannot run on ESXi 7.0 Update 3 and earlier.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 8.0 Update 1 and later	This virtual machine does not have all the capabilities available to virtual machines that run on ESXi 8.0 Update 2 and later. A virtual machine with such compatibility cannot run on ESXi 8.0 and earlier.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 8.0 Update 2 and later	This provides access to the latest virtual hardware features and ensures best performance. However, a virtual machine with such compatibility cannot run on ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, or ESXi 8.0 Update 1.
Cluster with ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, ESXi 8.0 Update 2, and ESXi 8.0 Update 3 hosts	ESXi 8.0 Update 3 and later	This provides access to the latest virtual hardware features and ensures best performance. However, a virtual machine with such compatibility cannot run on ESXi 6.5, ESXi 6.7, ESXi 7.0, ESXi 7.0 Update 1, ESXi 7.0 Update 2, ESXi 7.0 Update 3, ESXi 8.0, ESXi 8.0 Update 1, or ESXi 8.0 Update 2.

Set the Default Compatibility for Virtual Machine Creation

To ensure that virtual machines are compatible with the ESXi hosts that reside in the vSphere environment, you can set the default compatibility for virtual machine creation on the host, cluster, or data center.

Required privileges:

- On the host or cluster: **Host > Inventory > Modify cluster**
- On the data center: **Datacenter > Reconfigure datacenter**

The following conditions apply:

- To set the default compatibility on the cluster, the cluster must contain hosts that are connected and not in maintenance mode.
- A default compatibility setting on the host overrides a default cluster or data center setting.
- A default compatibility setting on the cluster overrides a default data center setting.

Select a host, cluster, or data center in the inventory.

Option	Action
Host	- Click the Configure tab, and under Virtual Machines, select Default VM Compatibility. - Click Edit. The Edit Default VM Compatibility dialog box opens. - From the Compatible with drop-down menu, select the compatibility and click OK. Note: You can set the compatibility only on hosts that are not part of a cluster.
Cluster	- Click the Configure tab and under Configuration, select General. - In the Default VM Compatibility section, click Edit. - From the Compatible with drop-down menu, select the compatibility and click OK.

Option	Action
	When you change the compatibility for a cluster, the compatibility for all hosts in the cluster changes as well.
Data Center	1. Right-click the data center and select Edit Default VM Compatibility. 2. From the Compatible with drop-down menu, select the compatibility and click OK.

When you create a virtual machine on one of these objects, the default compatibility setting is used.

Schedule a Compatibility Upgrade for a Single Virtual Machine

To make the virtual machine compatible with the latest version of the host, you can schedule a compatibility upgrade. The compatibility level determines the virtual hardware available to a virtual machine, which corresponds to the physical hardware on the host.

- Create a backup or snapshot of the virtual machines.
- Upgrade to the latest version of VMware Tools. On Microsoft Windows virtual machines, if you upgrade the compatibility level before you upgrade VMware Tools, the virtual machine might lose its network settings.
- Verify that all `.vmdk` files are available to the ESX/ESXi host on a VMFS5, or NFS datastore.
- Verify that the virtual machine is stored on VMFS5 or NFS datastores.

To schedule an upgrade for multiple virtual machines, see [Schedule a Compatibility Upgrade for a Virtual Machine](#).

1. Right-click a virtual machine and select **Compatibility > Schedule VM Compatibility Upgrade**.
2. In the **Schedule VM Compatibility Upgrade** dialog box, confirm that you want to schedule a compatibility upgrade by clicking **Yes**.
3. From the **Compatible with** drop-down menu, select the compatibility to upgrade to.
The virtual machine compatibility is upgraded the next time you restart the virtual machine.
4. Optional: To upgrade the compatibility when you do regularly scheduled guest maintenance, select **Only upgrade after normal guest OS shutdown**.

The virtual machine compatibility is upgraded and the new version appears on the virtual machine Summary tab.

Change the Default Virtual Machine Compatibility Setting

You can change the compatibility of a virtual machine by upgrading its compatibility or scheduling a compatibility upgrade.

- Create a backup or snapshot of the virtual machines. See [Manage Virtual Machines With Snapshots](#).
- Upgrade to the latest version of VMware Tools. If you upgrade the compatibility before you upgrade VMware Tools, the virtual machine might lose its network settings.
- Verify that all `.vmdk` files are available to the ESXi host on a VMFS3, VMFS5, or NFS datastore.
- Verify that the virtual machines are stored on VMFS3, VMFS5 or NFS datastores.
- Verify that the compatibility settings for the virtual machines are not the latest supported version.
- Determine the ESXi versions that you want the virtual machines to be compatible with. See [Virtual Machine Compatibility](#).

You can also change the default compatibility setting for the host, cluster, or data center.

The virtual machine compatibility determines the virtual hardware available to the virtual machine.

- Optional: To determine the compatibility setting of a virtual machine, select the virtual machine in the inventory and click the **Summary** tab.
- Select your task.

Client	Tasks
vSphere Client	- Change the default compatibility setting of a virtual machine. - Right-click a virtual machine and click Compatibility > Upgrade VM Compatibility. - Right-click a virtual machine and click Compatibility > Schedule VM Compatibility Upgrade.

Hardware Features Available with Virtual Machine Compatibility Settings

The virtual machine compatibility setting determines the virtual hardware available to the virtual machine, which corresponds to the physical hardware available on the host. To determine whether to upgrade the virtual machines in your environment, you can review and compare the hardware available for different compatibility levels.

With vSphere 8.0 Update 2, you can add up to 128 PCI passthrough devices to a virtual machine.

For more information about the virtual machine compatibility and the supported hardware versions, see the VMware KB article at: <https://kb.vmware.com/s/article/2007240>.

For more information about the UEFI version used in virtual machines, see the VMware KB article at: <https://kb.vmware.com/s/article/90088>.

Table 159: Supported Features for Virtual Machine Compatibility

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
		21	20	19	18	17	15
ESXi compatibility	See KB 2007240	ESXi 8.0 Update 2 and later	ESXi 8.0 and later	ESXi 7.0 Update 2 and later	ESXi 7.0 Update 1 and later	ESXi 7.0 and later	ESXi 6.7 Update 2 and later
BIOS firmware (1)	• HWV 4 or later	Y	Y	Y	Y	Y	Y
UEFI firmware (1)	• HWV 8 or later	Y	Y	Y	Y	Y	Y
Secure Boot	• HWV 13 or later	Y	Y	Y	Y	Y	Y
Maximum number of logical processors (2)	768 for HWV 18 or later 256 for HWV 15 - 17 128 for HWV 11 - 14 64 for HWV 9 - 10	768	768	768	768	256	256

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
	• 32 for HWV 8 • 8 for HWV 7 • 4 for HWV 4						
Maximum number of cores (virtual CPUs) per socket (3)	• 256 for HWV 20 • 64 for HWV 9 - 19 • 32 for HWV 8 • 8 for HWV 7 • 1 for HWV 4	256	256	64	64	64	64
CPU Hot Add support	• HWV 11 or later	Y	Y	Y	Y	Y	Y
Maximum memory (GB) (4)	• 24560 for HWV 20 or later • 6128 for HWV 13 - 17 • 4080 for HWV 11 • 1011 for HWV 8 - 10 • 256 for HWV 7 • 64 for HWV 4	24560	24560	24560	24560	6128	6128
Memory Hot Add Support	• HWV 7 or later	Y	Y	Y	Y	Y	Y
Virtual NUMA Topology	• HWV 20 or later	Y	Y	N	N	N	N
Synthetic Timers (18)	• HWV 20 or later	Y	Y	N	N	N	N
Maximum PCI passthrough devices	• HWV 20: — Max 32 for ESXi 8.0 — Max 64 for ESXi 8.0 Update 1	128	32 for ESXi 8.0 64 for ESXi 8.0 Update 1	16	16	16	16

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
	16 for HWV 13 - 19						
PCI Hot Add support	HWV 13 or later	Y	Y	Y	Y	Y	Y
Virtual PMC support	HWV 9 or later	Y	Y	Y	Y	Y	Y
Dynamic DirectPath	HWV 17 or later	Y	Y	Y	Y	Y	N
Enhanced DirectPath I/O	HWV 20 or later	Y	Y	N	N	N	N
Vendor Device Groups	HWV 20 or later	Y	Y	N	N	N	N
Intel Advanced Matrix eXtensions (AMX)	HWV 20 or later	Y	Y	N	N	N	N
AMD SEV-ES	HWV 18 or later	Y	Y	Y	Y	N	N
Virtual Hyperthreading	HWV 20 or later	Y	Y	N	N	N	N
Virtual I/O MMU for Intel hosts	HWV 14 or later	Y	Y	Y	Y	Y	Y
Virtual I/O MMU for AMD hosts	HWV 18 or later	Y	Y	Y	Y	N	N
Virtual SGX device	HWV 17 or later	Y	Y	Y	Y	Y	N
Nested virtualization (5)	HWV 9 or later	Y	Y	Y	Y	Y	Y
Microsoft VBS for Intel hosts	HWV 14 or later	Y	Y	Y	Y	Y	Y
Microsoft VBS for AMD hosts	HWV 19 or later	Y	Y	Y	N	N	N
Communication Interfaces							
Maximum Serial ports	32 for HWV 11 or later 4 for HWV 4 - 10	32	32	32	32	32	32
Maximum Parallel ports	3 for all HWV	3	3	3	3	3	3
USB 2.0 (13)	HWV 7 or later	Y	Y	Y	Y	Y	Y

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
USB 3.1 SuperSpeed (14)	• HWV 8 or later	Y	Y	Y	Y	Y	Y
USB 3.1 SuperSpeedPlus (14)	• HWV 17 or later	Y	Y	Y	Y	Y	N
USB 3.2 (14)	• HWV 21 or later	Y	N	N	N	N	N
Maximum USB devices	• HWV 7 or later	20	20	20	20	20	20
VMCI	• HWV 7 or later	Y	Y	Y	Y	Y	Y
DataSets Service	• HWV 20 or later	Y	Y	N	N	N	N
vMotion App Notification	• HWV 20 or later	Y	Y	N	N	N	N
Networking							
Maximum NICs	• 10 for HWV 7 or later • 4 for HWV 4	10	10	10	10	10	10
VMXNet3 version	• 8 for HWV 20 or later • 7 for HWV 19 • 6 for HWV 18 • 5 for HWV 15 and 17 • 4 for HWV 14 • 3 for HWV 13 • 2 for HWV 11 • 1 for HWV 7 - 10	8	8	7	6	5	5
VMXNet2 version	• HWV 4 or later	Y	Y	Y	Y	Y	Y
VMXNet version	• HWV 4 or later	Y	Y	Y	Y	Y	Y
PCNet32 (6)	• HWV 4 or later	Y	Y	Y	Y	Y	Y
E1000e	• HWV 8 or later	Y	Y	Y	Y	Y	Y

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
E1000	HWV 4 or later	Y	Y	Y	Y	Y	Y
PVRDMA RoCE versions	- RoCEv2 for HWV 14-20 - RoCEv1 for HWV 13-19	RoCEv2	RoCEv2	RoCEv1, RoCEv2	RoCEv1, RoCEv2	RoCEv1, RoCEv2	RoCEv1, RoCEv2
Maximum PVRDMA NICs	HWV 13 or later	10	10	10	1	1	1
PVRDMA native endpoints (7)	HWV 18 or later	Y	Y	Y	Y (7)	N	N
ATA Storage							
Maximum Floppy Devices	2 for all HWV	2	2	2	2	2	2
Maximum IDE controllers	HWV 4 or later	1	1	1	1	1	1
Maximum Virtual IDE disks and Virtual IDE CD-ROMs in any combination	HWV 4 or later	4	4	4	4	4	4
Maximum Virtual IDE CD-ROMs	HWV 4 or later	Y	Y	Y	Y	Y	Y
Virtual IDE Disks	HWV 7 or later	Y	Y	Y	Y	Y	Y
IDE hot add support	Not supported for all HWV	N	N	N	N	N	N
Maximum disks per Virtual SATA controller	HWV 10 or later	30	30	30	30	30	30
Maximum SATA controllers	4 for HWV 10 or later	4	4	4	4	4	4
SCSI Storage							
Maximum Virtual SCSI disks	HWV 4 or later	256	256	256	256	256	60
Maximum SCSI controllers	4 for HWV 4 or later (12)	4	4	4	4	4	4

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
		64	64	64	64	64	15
Maximum disks per VMware Paravirtual SCSI controllers	• HWV 7 or later						
Maximum disks per LSI Logic SAS SCSI controllers	• HWV 7 or later	15	15	15	15	15	15
Maximum disks per LSI Logic SCSI controllers	• HWV 4 or later	15	15	15	15	15	15
Maximum disks per Bus Logic SCSI controllers	• HWV 4 or later	15	15	15	15	15	15
SCSI passthrough support	• HWV 4 or later	Y	Y	Y	Y	Y	Y
SCSI hot add support	• HWV 4 or later	Y	Y	Y	Y	Y	Y
SCSI Read only support	• HWV 21 or later	Y	N	N	N	N	N
NVMe Storage							
Maximum Virtual NVMe disks	• 256 for HWV 21 or later • 60 for HWV 13 - 20	256	60	60	60	60	60
Maximum NVMe controllers	• 4 for HWV 13 or later	4	4	4	4	4	4
Maximum NVMe disks per controller	• 64 for HWV 21 or later • 15 for HWV 13 - 20	64	15	15	15	15	15
NVMe version (8)	• 1.3c for HWV 19 or later • 1.0e for HWV 13 - 18 • See notes (10) and (11)	1.3c (11)	1.3c (10)	1.3c (10)	1.0e	1.0e	1.0e
Virtual Disk							

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
		62 TB	62 TB	62 TB	62 TB	62 TB	62 TB
Maximum Virtual disk size	HWV 13 or later	62 TB	62 TB	62 TB	62 TB	62 TB	62 TB
Maximum redo log depth	all HWV	255	255	255	255	255	255
Snapshots							
Maximum snapshot levels	all HWV	32 (16)	32 (16)	32 (16)	32 (16)	32 (16)	32 (16)
Maximum total snapshots	all HWV	optional (17)	optional (17)	optional (17)	optional (17)	optional (17)	optional (17)
Persistent Memory							
Virtual PMEM disks support	HWV 14 or later	Y	Y	Y	Y	Y	Y
Maximum NVDIMM controllers	1 for HWV 14 or later	1	1	1	1	1	1
Maximum NVDIMM devices	64 for HWV 14 or later	64	64	64	64	64	64
Graphics							
Maximum vGPUs	16 for HWV 21 8 for HWV 20 4 for HWV 13 - 19	16	8	4	4	4	4
Maximum SVGA displays	10 for HWV 7 or later 1 for HWV 4	10	10	10	10	10	10
Maximum video memory (MB)	256 for HWV 18 or later 128 for HWV 11 - 17 512 for HWV 9 - 10 128 for HWV 4 - 8	256	256	256	256	128	128
Maximum 3D graphics memory (GB)	8 for HWV 18 or later 4 for HWV 17 2 for HWV 11 - 15	8	8	8	8	4	2

Feature	Feature Availability by Virtual Hardware Version (HWV)	Recent Hardware Versions					
SVGA 3D hardware acceleration	• HWV 8 or later	Y	Y	Y	Y	Y	Y
WDDM version	1.2 for HWV 20 or later 1.1 for HWV 17 - 19 1.0 for HWV 7 - 15	1.2	1.2	1.1	1.1	1.1	1.0
Direct3D version (9)	11.0 for HWV 19 or later 10.1 for HWV 17 - 18 10.0 for HWV 13 - 15 9.x for HWV 7 - 12	11.0 (15)	11.0 (15)	11.0 (15)	10.1	10.1	10.0
OpenGL version	4.3 for HWV 20 or later 4.0 for HWV 19 3.3 for HWV 13 - 18	4.3 (15)	4.3 (15)	4.0 (15)	3.3	3.3	3.3
Other Virtual Devices							
Memory ballooning	• HWV 4 or later	Y	Y	Y	Y	Y	Y
Virtual TPM 2.0	• HWV 14 or later	Y	Y	Y	Y	Y	Y
Virtual Precision Clock device	• HWV 17 or later	Y	Y	Y	Y	Y	N
Virtual Watchdog Timer device	• HWV 17 or later	Y	Y	Y	Y	Y	N

- (1) - By default, for each combination of hardware version and guest operating system, vSphere will use either UEFI or BIOS. Some guest operating systems require either UEFI or BIOS firmware, or support both firmware types. If the guest operating system supports BIOS and UEFI, change the default firmware type of a virtual machine before you install the guest operating system. Otherwise, you must reinstall the guest operating system.
- (2) - UEFI firmware is required for more than 128 logical processors.
- (3) - UEFI firmware is required for more than 64 virtual CPUs per sockets.

- (4) - UEFI firmware is required for more than 6 TB of virtual memory.
- (5) - For more information about the nested virtualization support limits, see the VMware KB article at <https://kb.vmware.com/s/article/2009916>.
- (6) - 32-bit virtual machines only.
- (7) - PVRDMA native endpoints for virtual machines of hardware version 18 do not support vSphere vMotion.
- (8) - By default, Windows virtual machines of hardware version 20 use NVMe version 1.0e.
- (9) - Direct3D requires VMware WDDM driver in the Windows guest operating system.
- (10) - NVMe version 1.0 is the default NVMe controller for Windows virtual machines of hardware versions 19 and 20.
- (11) - NVMe 1.3c is the default NVMe controller for virtual machines of hardware version 21 with Windows Server 2022 or later and Windows 11 or later guest operating systems.
- (12) - The maximum number of SCSI controllers is 4 for any combination of PVSCI, LSI, or BusLogic controllers.
- (13) - When you select USB 2.0 from the vSphere user interface, two controllers are added to the virtual machine - UHCI for USB 1.x and EHCI for USB 2.0.
- (14) - Virtual machines can have only one USB 3.x controller and an additional USB 2.0 controller.
- (15) - To power on a virtual machine, configured with Direct3D 11.0/OpenGL 4.0 and later, verify that the ESXi host graphics hardware is available. Virtual machines, configured with Direct3D 10.1/OpenGL 3.3 do not require an ESXi host graphics adapter.
- (16) - Depends on the number of virtual disks of a virtual machine:
 - 32 if the number is either less than or equal to 184.
 - 17 if the number is more than 184 but either less than or equal to 256.
 - 3 if the number is more than 256.
- (17) - Maximum total number of snapshots for a virtual machine that can be set with the `snapshot.maxSnapshots` configuration file option. The default value is `-1` or `unenforced`.
- (18) - Synthetic Timers are only available for Windows 10 and later guest operating systems.

Virtual CPU Configuration and Limitations

You can set most of the CPU parameters when you create virtual machines or after you install the guest operating system. Some actions require that you power off the virtual machine before you change the settings.

VMware uses the following terminology. Understanding these terms can help you plan your strategy for CPU resource allocation.

CPU	The CPU, or processor, is the component of a computer system that performs the tasks required for computer applications to run. The CPU is the primary element that performs the computer functions. CPUs contain cores.
CPU Socket	A CPU socket is a physical connector on a computer motherboard that connects to a single physical CPU. Some motherboards have multiple sockets and can connect multiple multicore processors (CPUs).
Core	A core contains a unit containing an L1 cache and functional units needed to run applications. Cores can independently run applications or threads. One or more cores can exist on a single CPU.
Resource sharing	Shares specify the relative priority or importance of a virtual machine or resource pool. If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when the two virtual machines are competing for resources.
Resource allocation	You can change CPU resource allocation settings, such as shares, reservation, and limit, when available resource capacity does not meet demands. For example, if at year end, the workload on

vSphere Virtual Symmetric Multiprocessing (Virtual SMP)

accounting increases, you can increase the accounting resource pool reserve.

Virtual SMP or vSphere Virtual Symmetric Multiprocessing is a feature that enables a single virtual machine to have multiple processors.

Virtual CPU Limitations

The maximum number of virtual CPUs that you can assign to a virtual machine is 768. The number of virtual CPUs depends on the number of logical CPUs on the host, and the type of guest operating system that is installed on the virtual machine.

Be aware of the following limitations:

- A virtual machine cannot have more virtual CPUs than the number of logical cores on the host. The number of logical cores is equal to the number of physical cores if hyperthreading is deactivated or two times that number if hyperthreading is activated.
- If a running virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a running virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.
- The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.
- Not every guest operating system supports Virtual SMP, and guest operating systems that support this functionality might support fewer processors than are available on the host. For information about Virtual SMP support, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Configuring Multicore Virtual CPUs

VMware multicore virtual CPU support lets you control the number of cores per virtual socket in a virtual machine. This capability lets operating systems with socket restrictions use more of the host CPU cores, which increases overall performance.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

Using multicore virtual CPUs can be useful when you run operating systems or applications that can take advantage of only a limited number of CPU sockets.

You can configure a virtual machine with ESXi 7.0 Update 1 and later compatibility to have up to 768 virtual CPUs. A virtual machine cannot have more virtual CPUs than the actual number of logical CPUs on the host. The number of logical CPUs means the number of physical processor cores or two times that number if hyperthreading is enabled. For example, if a host has 128 logical CPUs, you can configure the virtual machine for 128 virtual CPUs.

You configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single-core CPU, dual-core CPU, tri-core CPU, and so on. Your selection determines the number of sockets that the virtual machine has.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

For more information about multicore CPUs, see the *vSphere Resource Management* documentation.

Configure Virtual Machine CPU Resources

To improve the virtual machine performance, you can add, change, or configure CPU resources.

Enable CPU Hot Add

By default, you cannot add CPU resources to a machine when the virtual machine is powered on. The CPU hot add option lets you add CPU resources to a running virtual machine. The following conditions apply.

- For best results, use virtual machines that are compatible with ESXi 5.0 or later.
- Hot adding multicore virtual CPUs is supported only with virtual machines that are compatible with ESXi 5.0 or later.
- If a virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs to be over 128. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.
- Not all guest operating systems support CPU hot add. You can deactivate these settings if the guest operating system is not supported.
- To use the CPU hot add feature with virtual machines that are compatible with ESXi 4.x and later, set the Number of cores per socket to 1.
- Adding CPU resources to a running virtual machine with CPU hot add enabled disconnects and reconnects all USB passthrough devices that are connected to that virtual machine.

Note: If the ESXi host is of version 7.0 Update 2 and earlier, hot-adding virtual CPUs to a virtual machine with NVIDIA vGPU requires that the ESXi host have a free vGPU slot. Starting with vSphere 7.0 Update 3, the source host does not require to have a free vGPU slot.

Prerequisites

- Verify that the virtual machine is configured as follows.
 - Latest version of VMware Tools installed.
 - Guest operating system that supports CPU hot add.
 - Virtual machine compatibility is ESX/ESXi 4.x or later.
 - Virtual machine is powered off.
- Required privileges: **Virtual Machine > Configuration > Settings**

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**, and select **Enable CPU Hot Add**.
3. Click **OK**.

Results

You can now add CPUs even if the virtual machine is powered on.

Change the Number of Virtual CPUs

A virtual machine with ESXi 7.0 Update 1 and later compatibility can have up to 768 virtual CPUs. You can change the number of virtual CPUs while your virtual machine is powered off. If virtual CPU hot add is enabled, you can increase the number of virtual CPUs while the virtual machine is running. Virtual CPU hot add is supported for virtual machines with multicore CPU support and ESXi 5.0 and later compatibility. When the virtual machine is powered on and CPU hot add is enabled, you can hot add virtual CPUs to the running virtual machine. You can add only multiples of the number of cores per socket. If a virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768. The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

Prerequisites

- If CPU hot add is not enabled, power off the virtual machine before adding virtual CPUs.
- To hot add multicore CPUs, verify that the virtual machine is compatible with ESXi 5.0 and later.
- Verify that you have the **Virtual Machine > Configuration > Change CPU Count** privilege.

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**.
3. From the **CPU** drop-down menu, select the number of cores per socket and click **OK**.

Expose VMware Hardware Assisted Virtualization

You can expose full CPU virtualization to the guest operating system so that applications that require hardware virtualization can run on virtual machines without binary translation or paravirtualization.

Prerequisites

Note: VMware does not support running third-party hypervisors on ESXi with the exception of specifically defined and documented functionality. Currently, this is limited to Microsoft Hyper-V, strictly for VBS only and not for virtualizing multiple VMs.

For more information, see the VMware KB article at <https://kb.vmware.com/s/article/2009916>.

- Verify that the virtual machine compatibility is ESXi 5.1 and later.
- Intel Nehalem Generation (Xeon Core i7) or later processors or AMD Opteron Generation 3 (Greyhound) or later processors.
- Verify that Intel VT-x or AMD-V is enabled in the BIOS so that hardware assisted virtualization is possible.
- Verify that you have set the **Virtual Machine > Change Configuration > Change Settings** privilege on the vCenter Server system.
-

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**, and select **Expose hardware assisted virtualization to the guest OS**.
3. Click **OK**.

The **Configure** tab refreshes, and the Nested Hypervisor CPU option shows **Enabled**.

Activate Virtual CPU Performance Counters

You can use performance tuning tools in the guest operating system for software profiling. You can identify and improve processor performance problems. This capability is useful for software developers who optimize or debug software that runs in the virtual machine.

The following conditions apply:

- If virtual CPU performance counters are enabled, you can migrate the virtual machine only to hosts that have compatible CPU performance counters.
- If an ESXi host's BIOS uses a performance counter or if Fault Tolerance is enabled, some virtual performance counters might not be available for the virtual machine to use.

Note: If a virtual machine resides on an ESXi host in an EVC cluster, CPU counters are not supported for virtual machine creation or editing. You must deactivate CPU performance counters.

For a list of virtualized Model-Specific Registers (MSRs), see the VMware knowledge base article at <http://kb.vmware.com/kb/2030221>.

Prerequisites

- Verify that the virtual machine compatibility is ESXi 5.1 and later.
- Verify that the virtual machine is turned off.
- Verify that Intel Nehalem Generation (Xeon Core i7) or later processors or AMD Opteron Generation 3 ("Greyhound") or later processors are installed.
- Verify that Intel VT-x or AMD-V is enabled in the BIOS so that hardware-assisted virtualization is possible.
- Verify that you have set the **Virtual Machine > Change Configuration > Change Settings** privilege on the vCenter Server system.

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**, and select **Enable virtualized CPU performance counters** check box.
3. Click **OK**.

Configure Processor Scheduling Affinity

The **Scheduling Affinity** option gives you detailed control over how virtual machine CPUs are distributed across the host's physical cores. The option supports hyperthreading if hyperthreading is enabled. ESXi generally manages processor scheduling well, even when hyperthreading is enabled. These settings are useful only for fine-tuning critical virtual machines.

Using CPU affinity, you can assign a virtual machine to a specific processor. This assignment allows you to restrict the assignment of virtual machines to a specific available processor in multiprocessor systems.

This setting does not appear for virtual machines in a DRS cluster or when the host has only one processor core and no hyperthreading.

For potential issues with CPU affinity, see the *vSphere Resource Management* documentation.

Prerequisites

- Verify that the virtual machine is turned off.
- Verify that the virtual machine does not reside in a DRS cluster.
- Verify that the host has more than one physical processor core.
- Verify that you have the **Virtual Machine > Change Configuration > Change resource** privilege.

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**, and enter a comma-separated list of hyphenated processor ranges in the **Scheduling Affinity**
For example, "0,4-7" would indicate affinity with CPUs 0,4,5,6, and 7. Selecting all processors is identical to selecting no affinity. You must provide at least as many processor affinities as you have virtual CPUs.
3. Click **OK**.

Activate IO/MMU Virtualization Settings

ESXi can determine whether a virtual machine needs hardware support for virtualization. ESXi makes this determination based on the processor type and the virtual machine. Overriding the automatic selection can provide better performance for some use cases.

Important: Modern x86 processors can fully support virtualized workloads without software assistance. So, the CPU/MMU Virtualization setting has been deprecated in ESXi 6.7 and later. The CPU/MMU Virtualization setting is available only for virtual machines with ESXi 6.5 and earlier compatibility.

You can use software MMU when your virtual machine runs heavy workloads, such as Translation Lookaside Buffers (TLBs) intensive workloads that have significant impact on the overall system performance. However, software MMU has a higher overhead memory requirement than hardware MMU. So, to support software MMU, the maximum overhead supported for virtual machine limit in the VMkernel must be increased.

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**, and select **Enabled** check box to activate the I/O MMU options for a virtual machine.

Note: To take advantage of all features that virtual hardware version 13 provides, use the default MMU setting.

3. Click **OK**.

Enable CPU Hot Add

By default, you cannot add CPU resources to a virtual machine when the virtual machine is powered on. The CPU hot add option lets you add CPU resources to a running virtual machine.

- Verify that the virtual machine is configured as follows.
 - Latest version of VMware Tools installed.
 - Guest operating system that supports CPU hot add.
 - Virtual machine compatibility is ESX/ ESXi 4.x or later.
 - Virtual machine is powered off.
- Required privileges: **Virtual Machine > Configuration > Settings**

The following conditions apply.

- For best results, use virtual machines that are compatible with ESXi 5.0 or later.
- Hot adding multicore virtual CPUs is supported only with virtual machines that are compatible with ESXi 5.0 or later.
- If a virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs to be over 128. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.
- Not all guest operating systems support CPU hot add. You can deactivate these settings if the guest operating system is not supported.
- To use the CPU hot add feature with virtual machines that are compatible with ESXi 4.x and later, set the **Number of cores per socket** to 1.
- Adding CPU resources to a running virtual machine with CPU hot add enabled disconnects and reconnects all USB passthrough devices that are connected to that virtual machine.

Note: If the ESXi host is of version 7.0 Update 2 and earlier, hot-adding virtual CPUs to a virtual machine with NVIDIA vGPU requires that the ESXi host have a free vGPU slot. Starting with vSphere 7.0 Update 3, the source host does not require to have a free vGPU slot.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **VM Options** tab, expand **CPU Topology**, and select **Enable CPU Hot Add**.
3. Click **OK**.

You can now add CPUs even if the virtual machine is powered on.

Change the Number of Virtual CPUs

A virtual machine with ESXi 7.0 Update 1 and later compatibility can have up to 768 virtual CPUs. You can change the number of virtual CPUs while your virtual machine is powered off. If virtual CPU hot add is enabled, you can increase the number of virtual CPUs while the virtual machine is running.

- If CPU hot add is not enabled, power off the virtual machine before adding virtual CPUs.
- To hot add multicore CPUs, verify that the virtual machine is compatible with ESXi 5.0 and later.
- Verify that you have the **Virtual Machine > Configuration > Change CPU Count** privilege.

Virtual CPU hot add is supported for virtual machines with multicore CPU support and ESXi 5.0 and later compatibility. When the virtual machine is powered on and CPU hot add is enabled, you can hot add virtual CPUs to the running virtual machine. You can add only multiples of the number of cores per socket.

If a virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**.
3. From the **CPU** drop-down menu, select the number of cores.
4. From the **Cores Per Socket** drop-down menu, select the number of cores per socket and click **OK**.

Expose VMware Hardware Assisted Virtualization

You can expose full CPU virtualization to the guest operating system so that applications that require hardware virtualization can run on virtual machines without binary translation or paravirtualization.

Note: VMware does not support running third-party hypervisors on ESXi with the exception of specifically defined and documented functionality. Currently, this is limited to Microsoft Hyper-V, strictly for VBS only and not for virtualizing multiple VMs.

For more information, see the VMware KB article at <https://kb.vmware.com/s/article/2009916>.

- Verify that the virtual machine compatibility is ESXi 5.1 and later.
 - Intel Nehalem Generation (Xeon Core i7) or later processors or AMD Opteron Generation 3 (Greyhound) or later processors.
 - Verify that Intel VT-x or AMD-V is enabled in the BIOS so that hardware assisted virtualization is possible.
 - Required Privileges: **Virtual machine > Change Configuration > Change Settings** set on the vCenter Server system.
1. Right-click a virtual machine in the inventory and select **Edit Settings**.
 2. On the **Virtual Hardware** tab, expand **CPU**, and select **Expose hardware assisted virtualization to the guest OS**.
 3. Click **OK**.
The **Configure** tab refreshes, and the Nested Hypervisor CPU option shows **Enabled**.

Activate Virtual CPU Performance Counters

You can use performance tuning tools in the guest operating system for software profiling. You can identify and improve processor performance problems. This capability is useful for software developers who optimize or debug software that runs in the virtual machine.

- Verify that the virtual machine compatibility is ESXi 5.1 and later.
- Verify that the virtual machine is turned off.
- Verify that Intel Nehalem Generation (Xeon Core i7) or later processors or AMD Opteron Generation 3 ("Greyhound") or later processors are installed.
- Verify that Intel VT-x or AMD-V is enabled in the BIOS so that hardware-assisted virtualization is possible.
- Required Privileges: **Virtual machine > Change Configuration > Change Settings** is set on the vCenter Server system.

The following conditions apply:

- If virtual CPU performance counters are enabled, you can migrate the virtual machine only to hosts that have compatible CPU performance counters.
- If an ESXi host's BIOS uses a performance counter or if Fault Tolerance is enabled, some virtual performance counters might not be available for the virtual machine to use.

Note: If a virtual machine resides on an ESXi host in an EVC cluster, CPU counters are not supported for virtual machine creation or editing. You must deactivate CPU performance counters.

For a list of virtualized Model-Specific Registers (MSRs), see the VMware knowledge base article at <http://kb.vmware.com/kb/2030221>.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU** and select the **Enable virtualized CPU performance counters** checkbox.
3. Click **OK**.

Configure Processor Scheduling Affinity

The **Scheduling Affinity** option gives you detailed control over how virtual machine CPUs are distributed across the host's physical cores. The option supports hyperthreading if hyperthreading is enabled. ESXi generally manages processor scheduling well, even when hyperthreading is enabled. These settings are useful only for fine-tuning critical virtual machines.

- Verify that the virtual machine is turned off.
- Verify that the virtual machine does not reside in a DRS cluster.
- Verify that the host has more than one physical processor core.
- Privileges: **Virtual machine > Change Configuration > Change resource**

Using CPU affinity, you can assign a virtual machine to a specific processor. This assignment allows you to restrict the assignment of virtual machines to a specific available processor in multiprocessor systems.

This setting does not appear for virtual machines in a DRS cluster or when the host has only one processor core and no hyperthreading.

For potential issues with CPU affinity, see the *vSphere Resource Management* documentation.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**, and enter a comma-separated list of hyphenated processor ranges in the **Scheduling Affinity** text box.
For example, "0,4-7" would indicate affinity with CPUs 0,4,5,6, and 7. Selecting all processors is identical to selecting no affinity. You must provide at least as many processor affinities as you have virtual CPUs.
3. Click **OK**.

Activate IO/MMU Virtualization Settings

ESXi can determine whether a virtual machine needs hardware support for virtualization. ESXi makes this determination based on the processor type and the virtual machine. Overriding the automatic selection can provide better performance for some use cases.

Important: Modern x86 processors can fully support virtualized workloads without software assistance. So, the CPU/MMU Virtualization setting has been deprecated in ESXi 6.7 and later. The CPU/MMU Virtualization setting is available only for virtual machines with ESXi 6.5 and earlier compatibility.

You can use software MMU when your virtual machine runs heavy workloads, such as Translation Lookaside Buffers (TLBs) intensive workloads that have significant impact on the overall system performance. However, software MMU has a higher overhead memory requirement than hardware MMU. So, to support software MMU, the maximum overhead supported for virtual machine limit in the VMkernel must be increased.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**, and select the **Enabled** check box to activate the I/O MMU options for a virtual machine.
Note: To take advantage of all features that virtual hardware version 13 provides, use the default hardware MMU setting.
3. Click **OK**.

Virtual Memory Configuration

You can add, change, or configure virtual machine memory resources or options to improve virtual machine performance. You can set most of the memory parameters during virtual machine creation or after the guest operating system is installed.

Some actions require that you power off the virtual machine before changing the settings.

The memory resource settings for a virtual machine determine how much of the host's memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size. ESXi hosts limit the memory resource use to the maximum amount useful for the virtual machine, so that you can accept the default of Unlimited memory resources.

Change the Memory Configuration

You can reconfigure the amount of memory allocated to a virtual machine to improve performance.

Verify that you have the **Virtual machine > Change Configuration > Change Memory** privilege on the virtual machine.

Minimum memory size is 4 MB for virtual machines that use BIOS firmware. Virtual machines that use EFI firmware require at least 96 MB of RAM or they cannot power on.

Maximum memory size for virtual machines that use BIOS firmware is 24560 GB. You must use EFI firmware for virtual machines with memory size greater than 6128 GB.

Maximum memory size for a virtual machine depends on the physical memory of the host and the virtual machine compatibility settings.

If the virtual machine memory is greater than the host memory size, swapping occurs, which can have a severe effect on virtual machine performance. The maximum for best performance represents the threshold above which the physical memory of the host is insufficient to run the virtual machine at full speed. This value fluctuates as conditions on the host change, for example, as virtual machines are powered on or off.

The memory size must be a multiple of 4 MB.

Table 160: Maximum Virtual Machine Memory

Introduced in Host Version	Virtual Machine Compatibility	Maximum Memory Size
9.0	9.0 and later	24560 GB
8.0 Update 3	8.0 Update 3 and later	24560 GB
8.0 Update 2	8.0 Update 2 and later	24560 GB
8.0 Update 1	8.0 Update 1 and later	24560 GB
8.0	8.0 and later	24560 GB
7.0 Update 3	7.0 Update 3 and later	24560 GB
7.0 Update 2	7.0 Update 2 and later	24560 GB
7.0 Update 1	7.0 Update 1 and later	24560 GB
7.0	7.0 and later	6128 GB
6.7 Update 2	6.7 Update 2 and later	6128 GB
6.7	6.7 and later	6128 GB
6.5	6.5 and later	6128 GB

The ESXi host version indicates when support began for the increased memory size. For example, the memory size of a virtual machine with 6.0 and later compatibility running on 6.5 is restricted to 4080 GB.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Memory** and change the memory configuration.
 - a) In the **Memory** text box, enter the amount of RAM to assign to the virtual machine.
 - b) Select whether the memory is specified in MB, GB or TB.
3. Click **OK**.

Allocate Memory Resources to a Virtual Machine in the VMware Host Client

You can change the amount of memory resources allocated to a virtual machine by using the shares, reservations, and limits settings. The host determines the appropriate amount of physical RAM to allocate to virtual machines based on these settings. You can assign a high or low shares value to a virtual machine, depending on its load and status.

Power off the virtual machine.

The following user-defined settings affect the memory resource allocation of a virtual machine.

Limit	Places a limit on the consumption of memory for a virtual machine. This value is expressed in megabytes.
--------------	--

Reservation

Specifies the guaranteed minimum allocation for a virtual machine. The reservation is expressed in megabytes. If the reservation cannot be met, the virtual machine will not power on.

Shares

Each virtual machine is granted a number of memory shares. The more shares a virtual machine has, the greater share of host memory it receives. Shares represent a relative metric for allocating memory capacity. For more information about share values, see the *vSphere Resource Management* documentation.

You cannot assign a reservation to a virtual machine that is larger than the virtual machine configured memory. If you give a virtual machine a large reservation and reduce the virtual machine configured memory size, the reservation is reduced to match the new configured memory size.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **Memory**, and allocate the memory capacity for the virtual machine.

Option	Description
Reservation	Guaranteed memory allocation for this virtual machine.
Limit	Upper limit for this virtual machine's memory allocation. Select Unlimited to specify no upper limit.
Shares	The values Low , Normal , High , and Custom are compared to the sum of all shares of all virtual machines on the server.

4. Click **Save**.

Manage Persistent Memory

ESXi 6.7 and later provides support for the latest computer memory technology, which is called non-volatile memory (NVM) or persistent memory (PMem).

PMem combines the high data transfer rate of volatile computer memory with the persistence and resiliency of traditional storage. PMem devices have low access latency and can retain stored data through reboots or power outages.

Modes of Consumption of the Persistent Memory Resources of the Host

When you add a physical PMem device to a host, ESXi detects the PMem resource and exposes it as a host-local PMem datastore to the virtual machines that run on the host. Depending on the guest operating system, virtual machines can have direct access to the PMem resources.

Each host can have only one local PMem datastore that pools and represents all PMem resources of the host.

Persistent memory combines the properties of both memory and storage. So, virtual machines can consume the PMem resources of the ESXi host as memory (through virtual NVDIMM devices) or as storage (through virtual PMem hard disks).

The host-local PMem datastore stores all direct-accessed NVDIMM devices and virtual PMem hard disks.

Virtual PMem (vPMem)

In this mode, if the guest operating system is PMem-aware, the virtual machine can have direct access to the physical PMem resources of the host and use them as standard, byte-addressable memory.

Virtual machines use virtual non-volatile dual in-line memory modules (NVDIMMs) for direct access to PMem. The NVDIMM is a memory device that sits on an ordinary memory channel, but contains non-volatile memory. In vSphere

7.0, the virtual NVDIMM is a new type of device that represents the physical PMem regions of the host. A single virtual machine can have up to 64 virtual NVDIMM devices. Each NVDIMM device is stored on the host-local PMem datastore.

Note: To add an NVDIMM device to a virtual machine, the virtual machine must be of hardware version 14 and the guest operating system must support persistent memory. If the guest operating system is not PMem-aware, you can still use PMem, but you cannot add an NVDIMM device to the virtual machine.

Virtual PMem Disks (vPMemDisk)


In this mode, the virtual machine does not have direct access to the PMem resources of the host.

You must add a virtual PMem hard disk to the virtual machine. A virtual PMem hard disk is a traditional SCSI disk to which the PMem Storage Policy is applied. The policy automatically places the hard disk on the host-local PMem datastore.

In this mode of usage, there are no requirements for the hardware version of the virtual machine and the guest operating system.

Note: If the guest operating system is not PMem-aware, virtual machines can use PMem only through vPMemDisks.

The following diagram illustrates how the persistent memory components interact.



For information about how to configure and manage VMs with NVDIMMs or virtual persistent memory disks, see the *vSphere Resource Management* documentation.

Add an NVDIMM Device to a Virtual Machine

Add a virtual NVDIMM device to a virtual machine to enable it to use non-volatile, or persistent, computer memory. Non-volatile memory (NVM), or persistent memory (PMem), combines the high data transfer rates of the volatile memory with the persistence and resiliency of traditional storage. The virtual NVDIMM device is a virtual NVM device that can retain stored data through reboots or power source failures.

- Power off the virtual machine.
- Verify that the guest operating system of the virtual machine supports PMem.
- Verify that you have the **Datastore > Allocate space** privilege on the virtual machine.
- Verify that the host or the cluster on which the virtual machine resides has available PMem resources.
- To add an NVDIMM device to a virtual machine, verify that the virtual machine hardware is of version 14 or higher.
- To configure vSphere HA for PMem virtual machines, verify that the virtual machine is of version 19 or later, and vSphere HA is enabled on the cluster.

If an ESXi host fails or the datastore is no longer accessible, when you add an NVDIMM device to a virtual machine, you can configure vSphere HA to failover all PMem virtual machines on another host.

Note: If a host fails, NVDIMM PMem data cannot be restored. vSphere HA will restart the virtual machine on another host with a new and empty NVDIMM of the same size.

When you add an NVDIMM device to a virtual machine, you can take a snapshot of the virtual machine by configuring the NVDIMM snapshot mode of the virtual machine in the independent erase on revert mode or in the independent persistent mode.

NVDIMM Snapshot Conditions and Limitations

- You cannot configure the NVDIMM snapshot mode when the virtual machine has existing snapshots.
- You cannot add another NVDIMM device to the virtual machine if the NVDIMM snapshot mode is independent persistent or independent erase on revert, and the virtual machine has existing snapshots.
- You cannot add an NVDIMM device to a virtual machine, remove an NVDIMM device from a virtual machine, or change the NVDIMM device parameters if the virtual machine has snapshots. To allow these operations, remove all virtual machine snapshots.
- You can use vMotion to migrate virtual machines that have snapshots containing NVDIMM devices and revert these snapshots.

For more information, see the *vSphere Resource Management* guide.

1. Add an NVDIMM device to a virtual machine when you deploy a virtual machine or edit an existing virtual machine.

Option	Action
Create a virtual machine	1. Right-click any inventory object that is a valid parent object of a virtual machine and select New Virtual Machine. 2. On the Select a creation type page, select Create a new virtual machine, and click Next. 3. Navigate through the pages of the wizard. 4. On the Customize hardware page, click the Virtual Hardware tab. 5. On the Virtual Hardware tab, click the Add New Device button. 6. From the drop-down menu, select NVDIMM.
Edit a virtual machine	1. Right-click a virtual machine in the inventory and select Edit Settings.

Option	Action
	- Click the Virtual Hardware tab. - On the Virtual Hardware tab, click the Add New Device button. - From the drop-down menu, select NVDIMM.

The NVDIMM device appears in the Virtual Hardware devices list together with the virtual NVDIMM controller. Each virtual machine can have a maximum of one virtual NVDIMM controller and each NVDIMM controller can have up to 64 virtual NVDIMM devices.

Note: You can change the size of the NVDIMM device at a later time. The virtual machine must be powered off.

- In the **New NVDIMM** text box, enter the size of the NVDIMM device and select the units from the drop-down menu.

Note: Adjust the size of the newly added devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.

- Expand the **New NVDIMM device** section and select the **Allow failover on another host for all NVDIMM devices** check box.

Note: When you allow the failover process, if you add another NVDIMM device to the virtual machine, the NVDIMM device will have **PMem HA** enabled. If you want to preserve the NVDIMM content of the virtual machine during a host failure, make sure to deselect the **Allow failover on another host for all NVDIMM devices** check box.

- From the **Snapshot Mode**, configure the NVDIMM snapshot mode.

Option	Description
Persistent	Persistent mode is the default NVDIMM snapshot mode. Taking snapshots of virtual machines with an NVDIMM device in persistent mode is not supported.
Independent - Persistent	Taking snapshots of virtual machines with an NVDIMM device in independent-persistent mode is supported. When you revert to the previous snapshot, the current contents of the NVDIMM device is preserved.
Independent - Erase on revert	Taking snapshots of virtual machines with an NVDIMM device in independent-erase on revert mode is supported. When you revert to the previous snapshot, the content of the NVDIMM device is erased.

- If you deploy a virtual machine, click **Next**.
- If you edit an existing virtual machine, click **OK**.

When you power on the virtual machine, you can view the vSphere HA Protection status of the virtual machine in the **vSphere HA** panel on the **Summary** tab.

Change Memory Hot Add Settings

Memory hot add lets you add memory resources for a virtual machine while that virtual machine is powered on.

- Power off the virtual machine.
- Verify that the virtual machine has a guest operating system that supports memory hot add functionality.
- Verify that the virtual machine compatibility is ESXi 4.x and later.
- Verify that VMware Tools is installed.

Enabling memory hot add produces some memory overhead on the ESXi host for the virtual machine.

Note: If the ESXi host is of version 7.0 Update 2 and earlier, hot-adding memory to a virtual machine with NVIDIA vGPU requires that the ESXi host has a free vGPU slot. Starting with vSphere 7.0 Update 3, the source host does not require to have a free vGPU slot.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Memory**, and select **Enable** to enable adding memory to the virtual machine while it is powered on.
3. Click **OK**.

You can now add memory to a virtual machine, even if the virtual machine is turned on.

Virtual Disk Configuration

You can add large-capacity virtual disks to virtual machines and add more space to existing disks, even when the virtual machine is running.

You can set most of the virtual disk parameters during virtual machine creation or after you install the guest operating system.

You can store virtual machine data in a new virtual disk, an existing virtual disk, or a mapped SAN LUN. A virtual disk appears as a single hard disk to the guest operating system. The virtual disk is composed of one or more files on the host file system. You can copy or move virtual disks on the same hosts or between hosts.

For virtual machines running on an ESXi host, you can store virtual machine data directly on a SAN LUN instead of using a virtual disk file. This option is useful if in your virtual machines you run applications that must detect the physical characteristics of the storage device. Mapping a SAN LUN also allows you to use existing SAN commands to manage storage for the disk.

When you map a LUN to a VMFS volume, vCenter Server or the ESXi host creates a raw device mapping (RDM) file that points to the raw LUN. Encapsulating disk information in a file allows vCenter Server or the ESXi host to lock the LUN so that only one virtual machine can write to it. This file has a `.vmdk` extension, but the file contains only disk information that describes the mapping to the LUN on the ESXi system. The actual data is stored on the LUN. You cannot deploy a virtual machine from a template and store its data on a LUN. You can store its data only in a virtual disk file.

The amount of free space in the datastore is always changing. Ensure that you leave sufficient space for virtual machine creation and other virtual machine operations, such as growth of sparse files, snapshots, and so on. To review space utilization for the datastore by file type, see the *vSphere Monitoring and Performance* documentation.

Thin provisioning lets you create sparse files with blocks that are allocated upon first access, which allows the datastore to be over-provisioned. The sparse files can continue growing and fill the datastore. If the datastore runs out of disk space while the virtual machine is running, it can cause the virtual machine to stop functioning.

About Virtual Disk Provisioning Policies

When you perform certain virtual machine management operations, you can specify a provisioning policy for the virtual disk file. The operations include creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine.

NFS datastores with Hardware Acceleration and VMFS datastores support the following disk provisioning policies. On NFS datastores that do not support Hardware Acceleration, only thin format is available.

You can use Storage vMotion or cross-host Storage vMotion to transform virtual disks from one format to another.

Thick Provision Lazy Zeroed

Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand later on first write from the virtual

Thick Provision Eager Zeroed

machine. Virtual machines do not read stale data from the physical device.

A type of thick virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take longer to create virtual disks in this format than to create other types of disks. Increasing the size of an Eager Zeroed Thick virtual disk causes a significant stun time for the virtual machine.

Thin Provision

Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the virtual disk size.

However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations. If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire datastore space provisioned to it.

Thin provisioning is the fastest method to create a virtual disk because it creates a disk with just the header information. It does not allocate or zero out storage blocks. Storage blocks are allocated and zeroed out when they are first accessed.

Note: If a virtual disk supports clustering solutions such as Fault Tolerance, do not make the disk thin.

Large Capacity Virtual Disk Conditions and Limitations

Virtual machines with large capacity virtual hard disks, or disks greater than 2 TB, must meet resource and configuration requirements for optimal virtual machine performance.

The maximum value for large capacity hard disks is 62 TB. When you add or configure virtual disks, always leave a small amount of overhead. Some virtual machine tasks can quickly consume large amounts of disk space, which can prevent successful completion of the task if the maximum disk space is assigned to the disk. Such events might include taking snapshots or using linked clones. These operations cannot finish when the maximum amount of disk space is allocated. Also, operations such as snapshot quiesce, cloning, Storage vMotion, or vMotion in environments without shared storage, can take significantly longer to finish.

Virtual machines with large capacity disks have the following conditions and limitations:

- The guest operating system must support large capacity virtual hard disks.
- You can move or clone disks that are greater than 2 TB to ESXi 6.0 or later hosts or to clusters that have such hosts available.
- The datastore format must be one of the following:
 - VMFS5 or later
 - An NFS volume on a Network Attached Storage (NAS) server
 - vSAN
- Fault Tolerance is not supported.
- BusLogic Parallel controllers are not supported.

Change the Virtual Disk Configuration

If you run out of disk space, you can increase the size of the disk. You can change the virtual device node and the persistence mode for virtual disk configuration for a virtual machine.

If you change the disk mode to **Independent - Nonpersistent** or from **Independent - Nonpersistent**, power off the virtual machine.

Verify that you have the following privileges:

- **Virtual machine > Change Configuration > Modify device settings** on the virtual machine.
- **Virtual machine > Configuration > Extend virtual disk** on the virtual machine.
- **Datastore > Allocate space** on the datastore.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Hard disk** to view or change the disk settings, and click **OK**.

Option	Description
Maximum Size	Shows the maximum size of this hard disk on this VM. Note: Extending the size of a virtual hard disk causes stun time for the virtual machine. The stun time is longer if the virtual disk is of the Eager Zeroed Thick type.
VM storage policy	Select one of the available storage policies. See the <i>vSphere Storage</i> documentation for details. Note: You cannot change the VM storage policy of an existing PMem hard disk. You also cannot change the storage policy of an existing non-PMem disk to Host-local PMem Default Storage Policy.
Type	Shows the storage type. You cannot change this setting for an existing hard disk. You choose the storage type of a hard disk when you add the hard disk to the virtual machine. For more information about storage types and available disk formats, see the <i>vSphere Storage</i> documentation.
Sharing	Specifies sharing information.
Disk File	Lists disk files on the datastore.
Shares	Shares is a value that represents the relative metric for controlling disk bandwidth. The values Low, Normal, High, and Custom are compared to the sum of all shares of all virtual machines on the host.
Limit - IOPs	Allows you to customize IOPs. This value is the upper limit of I/O operations per second allocated to the virtual disk.
Disk mode	Disk mode determines how a virtual disk is affected by snapshots.
Virtual Device Node	Displays the virtual device node.

Use Disk Shares to Prioritize Virtual Machines

If multiple virtual machines access the same VMFS datastore and the same logical unit number (LUN), use disk shares to prioritize the disk accesses from the virtual machines. Disk shares distinguish high-priority from low-priority virtual machines.

You can allocate the host disk's I/O bandwidth to the virtual hard disks of a virtual machine. Disk I/O is a host-centric resource so you cannot pool it across a cluster.

Shares is a value that represents the relative metric for controlling disk bandwidth to all virtual machines. The values are compared to the sum of all shares of all virtual machines on the server.

Disk shares are relevant only within a given host. The shares assigned to virtual machines on one host have no effect on virtual machines on other hosts.

You can select an IOP limit, which sets an upper bound for storage resources that are allocated to a virtual machine. IOPs are the number of I/O operations per second.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Hard disk** to view the disk options.
3. From the **Shares** drop-down menu, select a value for the shares to allocate to the virtual machine. Alternatively, you can select **Custom** and you can enter a number of shares in the text box manually.
4. In the **Limit - IOPs** box, enter the upper limit of storage resources to allocate to the virtual machine, or select **Unlimited**.
5. Click **OK**.

Determine the Virtual Disk Format and Convert a Virtual Disk from Thin Provision to a Thick Provision Format

When the disk space is exhausted and a thin-provisioned disk cannot expand, the virtual machine cannot boot. If you created a virtual disk in the thin provision format, you can convert it to the thick provision format.

The thin provisioned disk starts small and at first, uses just as much storage space as it needs for its initial operations. After you convert the disk, it grows to its full capacity and occupies the entire datastore space provisioned to it during the disk's creation.

For more information about thin provisioning and available disk formats, see the *vSphere Storage* documentation.

Note: You cannot create thick provisioned disks on NFS datastores. For more information, see the VMware KB article at: <https://kb.vmware.com/s/article/2147607>.

1. Verify that the disk format of a virtual hard disk is Thin Provision.
 - a) Right-click a virtual machine and click **Edit Settings**.
 - b) On the **Virtual Hardware** tab, expand **Hard disk** and check the **Type** field.
 - c) To exit the wizard, click **Cancel**.
2. To open the datastore management panel, click the **Datastores** tab, and click a datastore from the list.
The datastore that stores the virtual machine files is listed.
3. Click the **Files** tab, and open the virtual machine folder.
4. Browse to the virtual disk file that you want to convert.
The file has the `.vmdk` extension.
5. To convert the virtual disk to a thick provision format, select the virtual disk file and click the **Inflate**.
The **Inflate** window appears.
6. Click **OK**.

The inflated virtual disk occupies the entire datastore space originally provisioned to it.

Note: The **Inflate** option might not be available if the virtual disk is thick or when the virtual machine is powered on.

How do I Add a Hard Disk to a Virtual Machine

When you create a virtual machine, a default virtual hard disk is added. You can add another hard disk if you run out of disk space, if you want to add a boot disk, or for other file management purposes. When you add a hard disk to a virtual machine, you can create a virtual disk, add an existing virtual disk, or add a mapped SAN LUN.

When you add a hard disk to a virtual machine, you can create a virtual disk, add an existing virtual disk, or add a mapped SAN LUN.

You can add a virtual hard disk to a virtual machine before or after you add a SCSI or SATA storage controller. The new disk is assigned to the first available virtual device node on the default controller, for example (0:1). Only device nodes for the default controller are available unless you add additional controllers.

The following ways to add disks can help you plan your disk configuration. These approaches show how you can optimize controller and virtual device nodes for different disks. For storage controller limitations, maximums, and virtual device node behavior, see [SCSI, SATA, and NVMe Storage Controller Conditions, Limitations, and Compatibility](#).

Add an existing hard disk that is configured as a boot disk during virtual machine creation.

To ensure that the virtual machine can boot, remove the existing disk before you add the boot disk. After you add a new hard disk to the virtual machine, you might need to go into the BIOS setup to ensure that the disk you were using to boot the virtual machine is still selected as the boot disk. You can avoid this problem by not mixing adapter types, and by using device node 0 on the first adapter as the boot disk.

Keep the default boot disk and add a new disk during virtual machine creation.

The new disk is assigned to the next available virtual device node, for example (0:1). You can add a new controller and assign the disk to a virtual device node on that controller, for example (1:0) or (1:1).

Add multiple hard disks to an existing virtual machine.

If you add multiple hard disks to a virtual machine, you can assign them to several SCSI or SATA controllers to improve performance. The controller must be available before you can select a virtual device node. For example, if you add controllers 1, 2, and 3, and add four hard disks, you might assign the fourth disk to virtual device node (3:1).

Add a New Hard Disk to a Virtual Machine

You can add a virtual hard disk to an existing virtual machine, or you can add a hard disk when you customize the virtual machine hardware during the virtual machine creation process. For example, you might need to provide additional disk space for an existing virtual machine with a heavy work load. During virtual machine creation, you might want to add a hard disk that is preconfigured as a boot disk.

- Ensure that you are familiar with configuration options and caveats for adding virtual hard disks. See [Virtual Disk Configuration](#).
- Before you add disks greater than 2 TB to a virtual machine, see [Large Capacity Virtual Disk Conditions and Limitations](#).
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the destination folder or datastore.

During virtual machine creation, a hard disk and a SCSI or SATA controller are added to the virtual machine by default, based on the guest operating system that you select. If this disk does not meet your needs, you can remove it and add a new hard disk at the end of the creation process.

If you add multiple hard disks to a virtual machine, you can assign them to several controllers to improve performance. For controller and bus node behavior, see [SCSI, SATA, and NVMe Storage Controller Conditions, Limitations, and Compatibility](#).

1. Right-click a virtual machine in the inventory and select **Edit Settings**.

2. On the **Virtual Hardware** tab, click the **Add New Device** button.

3. Select **Hard Disk** from the drop-down menu.

The hard disk appears in the Virtual Hardware devices list.

Note: If the host where the virtual machine resides has available PMem resources, you can place the new hard drive on the host-local PMem datastore.

4. Expand **New hard disk** and customize the settings of the new hard disk.

a) Enter a size for the hard disk and select the unit from the drop-down menu.

b) From the **VM storage policy**, select a storage policy or leave the default one.

c) From the **Location** drop-down menu, select the datastore location where you want to store virtual machine files.

d) From the **Disk Provisioning** drop-down menu, select the format for the hard disk.

Option	Action
Same format as source	Use the same format as the source virtual machine.
Thick Provision Lazy Zeroed	Create a virtual disk in a default thick format. Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine.
Thick Provision Eager Zeroed	Create a thick disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out during creation. It might take longer to create disks in this format than to create other types of disks.
Thin Provision	Use the thin provisioned format. At first, a thin provisioned disk uses only as much datastore space as the disk initially needs. If the thin disk needs more space later, it can grow to the maximum capacity allocated to it.

e) From the **Shares** drop-down menu, select a value for the shares to allocate to the virtual disk. Alternatively, you can select **Custom** and enter a value in the text box.

Shares is a value that represents the relative metric for controlling disk bandwidth. The values Low, Normal, High, and Custom are compared to the sum of all shares of all virtual machines on the host.

f) From the **Limit - IOPs** drop-down menu, customize the upper limit of storage resources to allocate to the virtual machine, or select **Unlimited**.

This value is the upper limit of I/O operations per second allocated to the virtual disk.

g) From the **Disk Mode** drop-down menu, select a disk mode.

Option	Description
Dependent	Dependent disks are included in snapshots.

Option	Description
Independent - Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk even if you revert a snapshot. When you power off or reset a virtual machine, the disk and all its snapshots are preserved.
Independent - Nonpersistent	Disks in nonpersistent mode behave like read-only disks. Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset the virtual machine, or when you delete a snapshot.

h) From the **Virtual Device Node**, select a virtual device node or leave the default one.

In most cases, you can accept the default device node. For a hard disk, a nondefault device node is useful to control the boot order or to have different SCSI controller types. For example, you might want to boot from an LSI Logic controller and share a data disk with another virtual machine that is using a BusLogic controller with bus sharing turned on.

Add an Existing Hard Disk to a Virtual Machine

You can add an existing virtual hard disk to a virtual machine when you customize the virtual machine hardware during the virtual machine creation process or after the virtual machine is created. For example, you might want to add an existing hard disk that is preconfigured as a boot disk.

- Make sure that you are familiar with controller and virtual device node behavior for different virtual hard disk configurations. See [How do I Add a Hard Disk to a Virtual Machine](#).
- Before you add disks greater than 2 TB to a virtual machine, see [Large Capacity Virtual Disk Conditions and Limitations](#).
- Verify that you have the **Virtual machine > Change Configuration > Add existing disk** privilege on the destination folder or datastore.

During virtual machine creation, a hard disk and a SCSI or SATA controller are added to the virtual machine by default, based on the guest operating system that you select. If this disk does not meet your needs, you can remove it and add an existing hard disk at the end of the creation process.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Optional: To delete the existing hard disk, move your pointer over the disk and click the **Remove** icon.
The disk is removed from the virtual machine. If other virtual machines share the disk, the disk files are not deleted.
3. On the **Virtual Hardware** tab, click the **Add New Device** button.
4. Select **Existing Hard Disk** from the drop-down menu.
The **Select File** dialog box opens.
5. In the **Select File**, expand a datastore, select a virtual machine folder, and select the disk to add.
6. Click **OK**.
The disk file appears in the **Contents** column. The **File Type** drop-down menu shows the compatibility file types for this disk.

7. Optional: Expand **New Hard disk** and make further customizations for the hard disk.
8. Click **OK**.

Add an RDM Disk to a Virtual Machine

You can use a raw device mapping (RDM) to store virtual machine data directly on a SAN LUN, instead of storing it in a virtual disk file. You can add an RDM disk to an existing virtual machine, or you can add the disk when you customize the virtual machine hardware during the virtual machine creation process.

- Ensure that you are familiar with SCSI controller and virtual device node behavior for different virtual hard disk configurations. See [How do I Add a Hard Disk to a Virtual Machine](#).
- Before you add disks greater than 2TB to a virtual machine, see [Large Capacity Virtual Disk Conditions and Limitations](#).
- Required privilege: **Virtual machine > Configuration > Configure Raw device**

When you give a virtual machine direct access to an RDM disk, you create a mapping file that resides on a VMFS datastore and points to the LUN. Although the mapping file has the same `.vmdk` extension as a regular virtual disk file, the mapping file contains only mapping information. The virtual disk data is stored directly on the LUN.

During virtual machine creation, a hard disk and a SCSI or SATA controller are added to the virtual machine by default, based on the guest operating system that you select. If this disk does not meet your needs, you can remove it and add an RDM disk at the end of the creation process.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click the **Add New Device** button and select **RDM Disk** from the drop-down menu. The **Select Target LUN** dialog box opens.
3. In the **Select Target LUN** dialog box, select the target LUN for the raw device mapping and click **OK**. The disk appears in the virtual device list.
4. Select the location for the mapping file.
 - To store the mapping file with the virtual machine configuration file, select **Store with the virtual machine**.
 - To select a location for the mapping file, select **Browse** and select the datastore location for the disk.
5. Select a compatibility mode.

Option	Description
Physical	Allows the guest operating system to access the hardware directly. Physical compatibility is useful if you are using SAN-aware applications on the virtual machine. However, a virtual machine with a physical compatibility RDM cannot be cloned, made into a template, or migrated if the migration involves copying the disk.
Virtual	Allows the RDM to behave as if it were a virtual disk, so that you can use such features as taking snapshots, cloning, and so on. When you clone the disk or make a template out of it, the contents of the LUN are copied into a <code>.vmdk</code> virtual disk file. When you migrate a virtual compatibility mode RDM, you can migrate the mapping file or copy the contents of the LUN into a virtual disk.

6. Accept the default or select a different virtual device node.

In most cases, you can accept the default device node. For a hard disk, a nondefault device node is useful to control the boot order or to have different SCSI controller types. For example, you might want to boot from an LSI Logic controller and share a data disk with another virtual machine using a BusLogic controller with bus sharing turned on.

7. Optional: If you selected virtual compatibility mode, select a disk mode to change the way that disks are affected by snapshots.

Disk modes are not available for RDM disks using physical compatibility mode.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent - Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk even if you revert a snapshot. When you power off or reset a virtual machine, the disk and all its snapshots are preserved.
Independent - Nonpersistent	Disks in nonpersistent mode behave like read-only disks. Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset the virtual machine, or when you delete a snapshot.

8. Click **OK**.

SCSI, SATA, and NVMe Storage Controller Conditions, Limitations, and Compatibility

To access virtual disks, CD/DVD-ROM, and SCSI devices, a virtual machine uses storage controllers, which are added by default when you create the virtual machine. You can add controllers or change the controller type after virtual machine creation or while you are in the creation wizard.

If you know about node behavior, controller limitations, and compatibility of different types of controllers before you change or add a controller, you can avoid potential boot problems.

How Storage Controller Technology Works

Storage controllers appear to a virtual machine as different types of SCSI controllers, including BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual SCSI. AHCI, SATA, and NVMe Express (NVMe) controllers are also available.

NVMe is a standardized protocol designed specifically for high-performance multi-queue communication with NVM devices. ESXi supports the NVMe protocol to connect to local and networked storage devices. For more information, see the *vSphere Storage* documentation.

When you create a virtual machine, the default controller is optimized for best performance. The controller type depends on the guest operating system, the device type, and sometimes, the virtual machine's compatibility. For example, when you create virtual machines with Apple Mac OS X guests and ESXi 5.5 and later compatibility, the default controller type for both the hard disk and the CD/DVD drive is SATA. When you create virtual machines with Windows Vista and later guests, a SCSI controller is the default for the hard disk and a SATA controller is the default for the CD/DVD drive.

In high performance storage environments you can benefit from using VMware Paravirtual SCSI controllers. The VMware Paravirtual SCSI controller ensures greater throughput and lower CPU use, which boosts performance as compared to the other SCSI controller options. For platform support for VMware Paravirtual SCSI controllers, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Each virtual machine can have a maximum of four SCSI controllers and four SATA controllers. The default SCSI or SATA controller is 0. When you create a virtual machine, the default hard disk is assigned to the default controller 0 at bus node (0:0).

When you add storage controllers, they are numbered sequentially 1, 2, and 3. If you add a hard disk, SCSI, or CD/DVD-ROM device to a virtual machine after virtual machine creation, the device is assigned to the first available virtual device node on the default controller, for example (0:1).

If you add a SCSI controller, you can reassign an existing or new hard disk or device to that controller. For example, you can assign the device to (1:*z*), where 1 is SCSI controller 1 and *z* is a virtual device node from 0 to 15. For SCSI controllers, *z* cannot be 7. By default, the virtual SCSI controller is assigned to virtual device node (*z*:7), so that device node is unavailable for hard disks or other devices.

If you add a SATA controller, you can reassign an existing or new hard disk or device to that controller. For example, you can assign the device to (1:*z*), where 1 is SATA controller 1 and *z* is a virtual device node from 0 to 29. For SATA controllers, you can use device nodes 0 through 29, including 0:7.

Alternatively, each virtual machine can have a maximum of four NVMe controllers. You can reassign an existing or new hard disk or device to that controller. For example, you can assign the hard disk to (*x*:*z*), where *x* is NVMe controller and *z* is a virtual device node. *x* has values from 0 to 3, and *z* has values from 0 to 14.

Storage Controller Limitations

Storage controllers have the following requirements and limitations:

- LSI Logic SAS and VMware Paravirtual SCSI are available for virtual machines with ESXi 4.x and later compatibility.
- AHCI SATA is available only for virtual machines with ESXi 5.5 and later compatibility.
- NVMe is available only for virtual machines with ESXi 6.5 and later compatibility.
- BusLogic Parallel controllers do not support virtual machines with disks larger than 2TB.
- Disks on VMware Paravirtual SCSI controllers might not experience optimal performance if they have snapshots or if the hosts memory is overcommitted.

CAUTION: Changing the controller type after the guest operating system is installed will make the disk and any other devices connected to the adapter inaccessible. Before you change the controller type or add a new controller, make sure that the guest operating system installation media contains the necessary drivers. On Windows guest operating systems, the driver must be installed and configured as the boot driver.

Storage Controller Compatibility

Adding different types of storage controllers to virtual machines that use BIOS firmware can cause operating system boot problems. In the following cases, the virtual machine might fail to boot correctly and you might have to enter the BIOS setup and select the correct boot device:

- If the virtual machine boots from LSI Logic SAS or VMware Paravirtual SCSI, and you add a disk that uses BusLogic, LSI Logic, or AHCI SATA controllers.
- If the virtual machine boots from AHCI SATA, and you add BusLogic Parallel or LSI Logic controllers.

Adding additional disks to virtual machines that use EFI firmware does not cause boot problems.

Table 161: VMware Storage Controller Compatibility

Existing Controller	Added Controller						
	BusLogic Parallel	LSI Logic	LSI Logic SAS	VMware Paravirtual SCSI	AHCI SATA	IDE	NVMe
BusLogic Parallel	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LSI Logic	Yes	Yes	Yes	Yes	Yes	Yes	Yes
LSI Logic SAS	Requires BIOS setup	Requires BIOS setup	Usually Works	Usually Works	Requires BIOS setup	Yes	Usually Works

Existing Controller	Added Controller						
	BusLogic Parallel	LSI Logic	LSI Logic SAS	VMware Paravirtual SCSI	AHCI SATA	IDE	NVMe
VMware Paravirtual SCSI	Requires BIOS setup	Requires BIOS setup	Usually Works	Usually Works	Requires BIOS setup	Yes	Usually Works
AHCI SATA	Requires BIOS setup	Requires BIOS setup	Yes	Yes	Yes	Yes	Yes
IDE	Yes	Yes	Yes	Yes	Yes	N/A	Yes
NVMe	Requires BIOS setup	Requires BIOS setup	Usually Works	Usually Works	Requires BIOS setup	Yes	Usually Works

Add a SATA Controller

If a virtual machine has multiple hard disks or CD/DVD-ROM devices, you can add up to three additional SATA controllers to assign the devices to. When you spread the devices among several controllers, you can improve performance and avoid data traffic congestion. You can also add additional controllers if you exceed the thirty-device limit for a single controller.

- Verify that the virtual machine compatibility is ESXi 5.5 and later.
- Verify that you are familiar with storage controller behavior and limitations.
- Verify that you have the **Virtual machine > Change Configuration > Add or remove device** privilege on the virtual machine.

You can boot virtual machines from SATA controllers and use them for large-capacity virtual hard disks.

Not all guest operating systems support AHCI SATA controllers. Typically, when you create virtual machines with ESXi 5.5 and later compatibility and Mac OS X guest operating systems, a SATA controller is added by default for the virtual hard disk and CD/DVD-ROM devices. Most guest operating systems, including Windows Vista and later have a default SATA controller for CD/DVD-ROM devices. To verify support, see the *VMware Compatibility Guides* at <http://www.vmware.com/resources/compatibility>.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click the **Add New Device** button.
3. Select **SATA Controller** from the drop-down menu.
The controller appears in the Virtual Hardware devices list.
4. Click **OK**.

You can add a hard disk or CD/DVD drive to the virtual machine and assign it to the new controller.

Add a SCSI Controller to a Virtual Machine

Many virtual machines have a SCSI controller by default, depending on the guest operating system. If you have a heavily loaded virtual machine with multiple hard disks, you can add up to three additional SCSI controllers to assign the disks to. When you spread the disks among several controllers, you can improve performance and avoid data traffic congestion. You can also add additional controllers if you exceed the 15-device limit for a single controller.

Verify that you have the **Virtual machine > Change Configuration > Add or remove device** privilege on the virtual machine.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click the **Add New Device** button. add a new SCSI controller.
3. Select **SCSI Controller** from the drop-down menu.
The controller appears in the Virtual Hardware device list.
4. Click **OK**.

You can now add a hard disk or other SCSI device to the virtual machine and assign it to the new SCSI controller.

How to Activate and Deactivate the Automatic Replacement of the LSI SAS Controller

vSphere 8.0 can automatically and safely replace the LSI SAS controller with the native VMware PVSCSI controller for existing virtual machines.

- Verify that the guest operating system of the virtual machines is Windows 10 and later or Windows Server 2016 and later.
- Verify that you have VMware Tools version 12.1.0 or later.
- The automatic replacement is possible for virtual machines that have VMware Paravirtual SCSI (PVSCSI) driver installed with VMware Tools.
- Verify that the VMware Tools Device Helper policy is enabled.

Starting with vSphere 8.0, you can activate and deactivate the automatic replacement of the LSI SAS controller by configuring the VMware Tools Device Helper policy, which is *unconfigured* by default. The default unconfigured policy setting deactivates the automatic replacement of the LSI SAS controller.

You can configure the VMware Tools policy settings by using the Local Group Policy Editor (*gpedit.msc*). You can also place the VMware Tools policy files to a domain controller and assign the policy to all virtual machines in the domain.

To enable the VMware Tools policy and activate the replacement of the LSI SAS controller with the VMware PVSCSI controller, perform the following steps.

1. Open a command-line window with the **Run as administrator** setting.
For example, click **Start**, type `cmd`, right-click the `cmd.exe` icon, and select **Run as administrator**.
2. Enter `xcopy "c:\ProgramData\Vmware\Vmware Tools\deviceHelperPolicy.admx" "c:\Windows\PolicyDefinitions"`.
3. Enter `xcopy "c:\ProgramData\Vmware\Vmware Tools\en-US\deviceHelperPolicy.adml" "c:\Windows\PolicyDefinitions\en-US"`.
4. Start the Local Group Policy Editor (*gpedit.msc*) and navigate to **Computer Configuration > Administrative Templates > VMware Tools Configuration**.
5. Select the **Device Helper Enable** setting and select **Enabled**.
6. Click **OK** and exit the Local Group Policy Editor.
7. Stop and start the VMTools service.
 - a) To stop the VMTools service, enter `sc stop vmtools`.
 - b) To start the VMTools service, enter `sc start vmtools`.

When you restart the VMTools service, the configuration is read and the Device Helper installs the PVSCSI driver for the next system restart.

8. Restart the virtual machine.

The LSI SAS controller is replaced with the VMware PVSCSI controller.

You can view the changes in the virtual machine settings or in the Windows virtual machine by using the Device Manager utility.

You can deactivate the automatic replacement of the LSI SAS controller by using the Local Group Policy Editor.

Option	Action
Deactivate by using the Local Group Policy Editor	In the Windows virtual machine, set the VMware Tools policy to <code>Disabled</code> or <code>Not configured</code> .

Add a Paravirtualized SCSI Adapter

You can add a VMware Paravirtual SCSI high performance storage controller to a virtual machine to provide greater throughput and lower CPU use.

- Verify that the virtual machine has a guest operating system with VMware Tools installed.
- Verify that the virtual machine compatibility is ESXi 4.x and later.
- Ensure that you are familiar with the VMware Paravirtual SCSI controller type.
- To access boot disk devices attached to a VMware Paravirtual SCSI controller, verify that the virtual machine has a Windows 2003 or Windows 2008 guest operating system.
- In some operating systems, before you change the controller type, create a virtual machine with an LSI Logic controller, install VMware Tools, and then change to paravirtual mode.

VMware Paravirtual SCSI controllers are best suited for environments, especially SAN environments, running I/O-intensive applications.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click the **Add New Device** button.
3. Select **SCSI Controller** from the drop-down menu.
4. Expand **New SCSI controller** and from the **Change Type** menu, select **VMware Paravirtual**.
The controller appears at the bottom of the Virtual Hardware device list.
5. Click **OK**.

Add an NVMe Controller

If a virtual machine has multiple hard disks, you can add up to four virtual NVMe controllers to which to assign the virtual disks. Using an NVMe controller significantly reduces the software overhead for processing guest OS I/O, as compared to AHCI SATA or SCSI controllers.

- Verify that the virtual machine has a guest operating system that supports NVMe.
- Verify that the virtual machine compatibility is ESXi 6.5 or later.
- Verify that you are familiar with storage controller behavior and limitations.
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the virtual machine.

NVMe controllers perform best with virtual disks on an all-flash disk array, local NVMe SSD, and PMem storage.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click the **Add New Device** button.
3. Select **NVMe Controller** from the drop-down menu.
The controller appears in the Virtual Hardware devices list.

4. Click **OK**.

You can add a hard disk to the virtual machine and assign it to the NVMe controller.

Change the SCSI Controller Configuration

You can specify the SCSI controller type and you can set the type of SCSI bus sharing for a virtual machine.

- Verify that you are familiar with the limitations and conditions for configuring SCSI controllers.
- Verify that you have the **Virtual machine > Change Configuration > Modify device settings** privilege on the virtual machine.

The choice of a SCSI controller type does not affect whether your virtual disk is an IDE or SCSI disk. The IDE adapter is always ATAPI. The default for your guest operating system is already selected.

The choice of a SCSI bus sharing option determines whether virtual machines on different hosts can access the same virtual disk.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **SCSI controller**, and select a SCSI controller type from the **Change Type** drop-down menu.

CAUTION: Changing the SCSI controller type might result in a virtual machine boot failure.

Do not select a BusLogic Parallel controller for virtual machines with disks larger than 2 TB. This controller does not support large capacity hard disks.

The vSphere Client displays information about what happens if you change the SCSI controller type. If you select a controller type that is not recommended for the virtual machine's guest operating system, a warning is displayed.

3. Expand **SCSI controller** and select the type of sharing in the **SCSI Bus Sharing** drop-down menu.

Option	Description
None	Virtual disks cannot be shared by other virtual machines.
Physical	Virtual disks can be shared by virtual machines on any ESXi host.
Virtual	Virtual disks can be shared by virtual machines on the same ESXi host.

For virtual or physical bus sharing, select **Thick provision eager zeroed** when you create the disk.

4. Click **OK**.

Virtual Machine Network Configuration

vSphere networking features provide communication between virtual machines on the same host, between virtual machines on different hosts, and between other virtual and physical machines.

When you configure networking for a virtual machine, you select or change an adapter type, a network connection, and whether to connect the network when the virtual machine powers on.

Network Adapter Basics

When you configure a virtual machine, you can add network adapters (NICs) and specify the adapter type.

Network Adapter Types

The type of network adapters that are available depend on the following factors:

- The virtual machine compatibility, which depends on the host that created or most recently updated it.

- Whether the virtual machine compatibility has been updated to the latest version for the current host.
- The guest operating system.

Supported NICs currently differ between an on-premises environment and VMware Cloud on AWS. The following NIC types are supported in an on-premises deployment:

E1000E	Emulated version of the Intel 82574 Gigabit Ethernet NIC. E1000E is the default adapter for Windows 8 and Windows Server 2012.
E1000	Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in most newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later.
Flexible	Identifies itself as a Vlan adapter when a virtual machine boots, but initializes itself and functions as either a Vlan or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlan adapter to the higher performance VMXNET adapter.
Vlan	Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in 32-bit legacy guest operating systems. A virtual machine configured with this network adapter can use its network immediately.
VMXNET	Optimized for performance in a virtual machine and has no physical counterpart. Because operating system vendors do not provide built-in drivers for this card, you must install VMware Tools to have a driver for the VMXNET network adapter available.
VMXNET 2 (Enhanced)	Based on the VMXNET adapter but provides high-performance features commonly used on modern networks, such as jumbo frames and hardware offloads. VMXNET 2 (Enhanced) is available only for some guest operating systems on ESX/ ESXi 3.5 and later.
VMXNET 3	A paravirtualized NIC designed for performance. VMXNET 3 offers all the features available in VMXNET 2 and adds several new features, such as multiqueue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery. VMXNET 3 is not related to VMXNET or VMXNET 2.
PVRDMA	A paravirtualized NIC that supports remote direct memory access (RDMA) between virtual machines through the OFED verbs API. All virtual machines must have a PVRDMA device and should be connected to a distributed switch. PVRDMA supports VMware vSphere vMotion and snapshot technology. It is available in virtual machines with hardware version 13 and guest operating system Linux kernel 4.6 and later. For information about assigning an PVRDMA network adapter to a virtual machine, see the <i>vSphere Networking</i> documentation.
SR-IOV passthrough	Representation of a virtual function (VF) on a physical NIC with SR-IOV support. The virtual machine and the physical adapter exchange data without using the VMkernel as an intermediary. This adapter type is suitable for virtual machines where latency might cause failure or that require more CPU resources. SR-IOV passthrough is available in ESXi 6.0 and later for guest operating systems Red Hat Enterprise Linux 6 and later, and Windows Server 2008 R2 with SP2. An operating system release might contain a default VF driver for certain NICs, while for others

you must download and install it from a location provided by the vendor of the NIC or of the host.

For information about assigning an SR-IOV passthrough network adapter to a virtual machine, see the *vSphere Networking* documentation.

For network adapter compatibility considerations, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Legacy Network Adapters and ESXi Virtual Hardware Versions

The default network adapter types for all legacy virtual machines depend on the adapters available and compatible to the guest operating system and the version of virtual hardware on which the virtual machine was created.

If you do not upgrade a virtual machine to use a virtual hardware version, your adapter settings remain unchanged. If you upgrade your virtual machine to take advantage of newer virtual hardware, your default adapter settings will likely change to be compatible with the guest operating system and upgraded host hardware.

To verify the network adapters that are available to your supported guest operating system for a particular version of vSphere ESXi, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Network Adapters and Legacy Virtual Machines

Legacy virtual machines are virtual machines that are supported by the product in use, but are not current for that product.

The default network adapter types for all legacy virtual machines depend on the adapters available and compatible to the guest operating system and the version of virtual hardware on which the virtual machine was created.

If you do not upgrade a virtual machine to correspond with an upgrade to a newer version of an ESXi host, your adapter settings remain unchanged. If you upgrade your virtual machine to take advantage of newer virtual hardware, your default adapter settings will likely change to be compatible with the guest operating system and upgraded host hardware.

To verify the network adapters that are available to your supported guest operating system for a particular version of vSphere ESXi, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Change the Virtual Machine Network Adapter Configuration

You can change the virtual machine network configuration, including its power on behavior and resource allocation.

- Required privilege: **Network > Assign network** on a network if you are changing the network the virtual machine connects to.
- To activate the UPT support, make sure that the virtual machine is of hardware version 20 or later.

Starting with vSphere 8.0, you can enable the Uniform Passthrough (UPT) compatibility on a VMXNET3 adapter.

For details about configuring the networking for virtual machine network adapters, see the *vSphere Networking* documentation.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Network adapter**, and select the port group to connect to from the drop-down menu.

The menu lists all standard and distributed port groups that are available for virtual machine use on the host.

If you want to provision bandwidth to the network adapter from a reserved quota by using vSphere Network I/O Control version 3, select a port group that is associated with the network resource pool that provides the quota.

3. Optional: Change the **Status** settings.

Option	Description
Connected	To connect or disconnect the virtual network adapter, select or deselect this option while the virtual machine is running. This check box is not available when the virtual machine is powered off.
Connect at power on	To connect to the network when the virtual machine powers on, select this option for the virtual network adapter. If you do not select this option, you must manually connect the adapter in order for the virtual machine to access the network.

4. To allow changes in the maximum transmissions unit (MTU) of packets from the guest operating system, select an option from the **Allow Guest MTU Change** drop-down menu.5. To enable the UPT, select the **Use UPT Support** check box.

Note: When you activate the UPT support, the full memory size of the virtual machine is automatically reserved.

6. If the network adapter is connected to a distributed port group of a distributed switch that has vSphere Network I/O Control version 3 enabled, allocate bandwidth to the adapter.

Note: You cannot allocate bandwidth to **SR-IOV passthrough** network adapters.

- From the **Shares** drop-down menu, set the relative priority of the traffic from this virtual machine as shares from the capacity of the connected physical adapter.
- In the **Reservation** text box, reserve a minimum bandwidth that must be available to the VM network adapter when the virtual machine is powered on.
- In the **Limit** text box, set a limit on the bandwidth that the VM network adapter can consume.

7. Click **OK**.

How to Add a Network Interface Controller to a Virtual Machine

You can add a network interface controller (NIC), or network adapter to a virtual machine to connect to a network, to enhance communications, or to replace an older adapter.

When you add a NIC to a virtual machine, you select the adapter type, network connection, whether the device must connect when the virtual machine is powered on, and the bandwidth allocation.

Starting with vSphere 8.0, you can enable the Uniform Passthrough (UPT) compatibility on a VMXNET3 adapter.

For details about configuring the networking for virtual machine network adapters, see the *vSphere Networking* documentation.

- Right-click a virtual machine in the inventory and select **Edit Settings**.
- On the **Virtual Hardware** tab, click **Add New Device** and select **Network Adapter** from the drop-down menu. The new network adapter appears at the bottom of the network adapters list.

3. Expand **New Network** and select the standard or distributed port group to connect to.

The menu lists all standard and distributed port groups that are available for virtual machine use on the host.

If you want to provision bandwidth to the network adapter from a reserved quota by using vSphere Network I/O Control version 3, select a port group that is associated with the network resource pool that provides the quota.

4. Optional: Review and optionally change the **Status** settings.

Option	Description
Connected	To connect or disconnect the virtual network adapter, select this option while the virtual machine is running. This check box is not available when the virtual machine is powered off.

Option	Description
Connect at power on	To connect to the network when the virtual machine powers on, select this option for the virtual network adapter. If you do not select this option, you must manually connect the adapter for the virtual machine to access the network.

5. From the **Adapter Type** drop-down menu, select the network adapter type.
6. To allow changes in the maximum transmissions unit (MTU) of packets from the guest operating system, select an option from the **Allow Guest MTU Change** drop-down menu.
7. To enable the UPT, select the **Use UPT Support** check box.
Note: When you activate the UPT support, the full memory size of the virtual machine is automatically reserved.
8. If the network adapter is connected to a distributed port group of a distributed switch that has vSphere Network I/O Control version 3 enabled, allocate bandwidth to the adapter.
Note: You cannot allocate bandwidth to **SR-IOV passthrough** network adapters.
 - a) From the **Shares** drop-down menu, set the relative priority of the traffic from this virtual machine as shares from the capacity of the connected physical adapter.
 - b) In the **Reservation** text box, reserve a minimum bandwidth that must be available to the VM network adapter when the virtual machine is powered on.
 - c) In the **Limit** text box, set a limit on the bandwidth that the VM network adapter can consume.
9. Click **OK**.

Virtual Machine Parallel and Serial Port Configuration

Parallel and serial ports are interfaces for connecting peripherals to the virtual machine. The virtual serial port can connect to a physical serial port or to a file on the host computer.

You can also use it to establish a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. You can add parallel and serial ports and change the parallel and serial port configuration. Hardware version 11 and later versions allow you to configure virtual machines in such a way that serial and parallel ports are absent from the virtual chipset altogether.

Starting with vSphere 8.0, you cannot add, remove, and configure parallel ports. For information, see <https://kb.vmware.com/s/article/78978>.

Using Serial Ports with vSphere Virtual Machines

You can set up virtual serial port connections for vSphere virtual machines in several ways. The connection method that you select depends on the task that you need to accomplish.

You can set up virtual serial ports to send data in the following ways.

Physical serial port on the host

Sets the virtual machine to use a physical serial port on the host computer. This method lets you use an external modem or a hand-held device in a virtual machine.

Output to file

Sends output from the virtual serial port to a file on the host computer. This method lets you capture the data that a program running in the virtual machine sends to the virtual serial port.

Connect to a named pipe

Sets a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. With this method, two virtual machines or a virtual machine and a process on the host can communicate as if they

Connect over the network

were physical machines connected by a serial cable. For example, use this option for remote debugging of a virtual machine.

Enables a serial connection to and from a virtual machine's serial port over the network. The Virtual Serial Port Concentrator (vSPC) aggregates traffic from multiple serial ports onto one management console. vSPC behavior is similar to physical serial port concentrators. Using a vSPC also allows network connections to a virtual machine's serial ports to migrate seamlessly when you use vMotion to migrate the virtual machine. For requirements and steps to configure the Avocent ACS v6000 virtual serial port concentrator, see <http://kb.vmware.com/kb/1022303>.

Server and Client Connections for Named Pipe and Network Serial Ports

You can select a client or server connection for serial ports. Your selection determines whether the system waits for a connection or initiates it. Typically, to control a virtual machine over a serial port, you select a server connection. This selection lets you control the connections, which is useful if you connect to the virtual machine only occasionally. To use a serial port for logging, select a client connection. This selection lets the virtual machine connect to the logging server when the virtual machine starts and to disconnect when it stops.

Supported Serial Ports

When you use a physical serial port for serial port passthrough from an ESXi host to a virtual machine, serial ports that are integrated into the motherboard are supported. A virtual machine can use up to 32 serial ports.

Unsupported Serial Ports

When you use a physical serial port for serial port passthrough from an ESXi host to a virtual machine, the serial ports connected through USB are not supported for serial port passthrough. They might be supported by USB passthrough from an ESXi host to a virtual machine. See [USB Configuration from an ESXi Host to a Virtual Machine](#).

In addition, you cannot use Migration with vMotion when you use a physical serial port for serial passthrough.

Adding a Firewall Rule Set for Serial Port Network Connections

If you add or configure a serial port that is backed by a remote network connection, ESXi firewall settings can prevent transmissions.

Before you connect network-backed virtual serial ports, you must add one of the following firewall rule sets to prevent the firewall from blocking communication:

- **VM serial port connected to vSPC.** Use to connect the serial port output through a network with the **Use virtual serial port concentrator** option enabled to allow only outgoing communication from the host.
- **VM serial port connected over network.** Use to connect the serial port output through a network without the virtual serial port concentrator.

Important: Do not change the allowed IP list for either rule set. Updates to the IP list can affect other network services that might be blocked by the firewall.

For details about allowing access to an ESXi service through the firewall, see the *vSphere Security* documentation.

Configure Virtual Machine Communication Interface Firewall

You can configure the virtual machine Communication Interface firewall (VMCI) to restrict virtual machines accessing the hypervisor-based services and VMCI-based services.

You can restrict VMCI usage to a subset of VMCI-based services on each virtual machine. For example, you can allow certain virtual machines to access VMCI services and deny access to others for security reasons.

Currently, VMCI devices support guest to host communication. A virtual machine can communicate with VMCI services through the following means:

- ESXi hypervisor
- Services installed on the host operating system in the form of a vmkernel module
- Applications installed by a verified vSphere Installation Bundle

Change the Serial Port Configuration

You can connect the virtual serial port to a physical serial port or to a file on the host computer. You can also use a host-side named pipe to set up a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. In addition, you can use a port or vSPC URI to connect a serial port over the network. You can add up to 32 serial ports to a virtual machine.

- Check that you know the correct media types for the port to access, vSPC connections, and any conditions that might apply. See [Using Serial Ports with vSphere Virtual Machines](#).
- To connect a serial port over a network, add a Firewall rule set. See [Adding a Firewall Rule Set for Serial Port Network Connections](#).
- To use authentication parameters with network serial port connections, see [Authentication Parameters for Virtual Serial Port Network Connections](#).
- Required privileges:
 - **Virtual machine > Change Configuration > Modify device settings** on the virtual machine.
 - **Virtual machine > Interaction > Device connection** on the virtual machine to change the device connection status.

Virtual machines can be in a powered-on state during configuration.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Serial port**, and select a connection type.

Option	Action
Use physical serial port	Select this option to have the virtual machine use a physical serial port on the host computer. Select the serial port from the drop-down menu.
Use output file	Select this option to send output from the virtual serial port to a file on the host computer. Browse to select an output file to connect the serial port to.
Use named pipe	Select this option to set a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. 1. Type a name for the pipe in the Pipe Name field. 2. Select the Near end and Far end of the pipe from the drop-down menus.
Use Network	Select Use network to connect through a remote network. 1. Select the network backing. – Select Server to have the virtual machine monitor incoming connections from other hosts. – Select Client to have the virtual machine initiate a connection to another host. 2. Enter a Port URI. The URI is the remote end of the serial port to which the virtual machine's serial port should connect.

Option	Action
	3. If vSPC is used as an intermediate step to access all virtual machines through a single IP address, select Use Virtual Serial Port Concentrator and enter the vSPC URI location.
Printer	Select Printer to connect to a remote printer.

3. Optional: Select **Yield CPU on poll**.

Select this option only for guest operating systems that use serial ports in polled mode. This option prevents the guest from consuming excessive CPUs.

4. Optional: Select **Connect at power on** to connect the serial port when the virtual machine powers on.

5. Click **OK**.

Establishing Serial Port Network Connections to a Client or Server Without Authentication Parameters

If you do not use vSPC and you configure your virtual machine with a serial port connected as a server with a `telnet://:12345` URI, you can connect to your virtual machine's serial port from your Linux or Windows operating system.

```
telnet:yourESXiServerIPAddress12345
```

Similarly, if you run the Telnet Server on your Linux system on port 23 (`telnet://yourLinuxBox:23`), you configure the virtual machine as a client URI.

```
telnet://yourLinuxBox:23
```

The virtual machine initiates the connection to your Linux system on port 23.

Authentication Parameters for Virtual Serial Port Network Connections

When you establish serial port connections over the network, you can use authentication parameters to secure the network. These parameters can support an encrypted connection with a remote system using SSL over Telnet or Telnets, or an encrypted connection with a concentrator using SSL over Telnet or Telnets.

URI Forms

If you do not use virtual serial port network connection (vSPC) and you configure your virtual machine with a serial port connected as a server with a `telnet://:12345` URI, you can connect to your virtual machine's serial port from your Linux or Windows operating system. You use one of the following formats:

- Telnet over TCP.

```
telnet://host:port
```

The virtual machine and remote system can negotiate and use SSL if the remote system supports the Telnet authentication option. If not, the connection uses unencrypted text (plain text).

- Telnets over SSL over TCP.

```
telnets://host:port
```

SSL negotiation begins immediately, and you cannot use the Telnet authentication option.

Authentication Parameters

For an encrypted connection, the URI includes a set of authentication parameters. Enter the parameters as key words or key/value pairs. You can enter authentication parameters for secure Telnet (`telnets`), or for Telnet (`telnet`) as shown in the following syntax:

```
telnet://host:port#key[=value] [&key[=value] ...]
```

The first parameter must have a number sign (#) prefix. Additional parameters must have an ampersand (&) prefix. The following parameters are supported.

thumbprint=*value*

Specifies a certificate thumbprint against which the peer certificate thumbprint is compared. When you specify a thumbprint, certificate verification is enabled.

peerName=*value*

Specifies the peer name that is used to validate the peer certificate. When you specify a peer name, certificate verification is enabled.

verify

Forces certificate verification. The virtual machine will verify that the peer certificate subject matches the specified peerName and that it was signed by a certificate authority known to the ESXi host. Verification is enabled if you specify a thumbprint or peerName

cipherList=*value*

Specifies a list of SSL ciphers. The ciphers are specified as a list separated by colons, spaces, or commas.

Establishing Serial Port Network Connections to a Client or Server**Simple Server Connection**

To connect to a virtual machine's serial port from a Linux or Windows operating system if you do not use vSPC, configure the virtual machine with a serial port connected as a server with a `telnet://:12345` URI. To access a virtual serial port from a client, use `telnet yourESXiServerIPAddress 12345`.

Secure Server Connection

To enforce an encrypted connection to the virtual machine's serial port from a Linux operating system, you can configure Telnet to enforce encryption by configuring the virtual machine with a serial port connected as a server with a `telnet://:12345#verify` URI. To access a virtual serial port from a client, use `telnet-ssl yourESXServerName 12345`. This connection will fail if the Telnet program you are using does not support SSL encryption.

Simple Client Connection

If you are running a Telnet server on your system and you want the virtual machine to automatically connect to it, you can configure the virtual machine as a client using `telnet://yourLinuxBox:23`.

The Virtual machine keeps initiating the Telnet connection to port 23 on *yourLinuxBox*.

Secure Client Connection

Additional URI options allow you to enforce a specific server certificate and restrict the ciphers being used. Virtual machines with a serial port configured as a client with `telnet://ipOfYourLinuxBox:23#cipherList=DHE-RSA-AES256-SHA256:DHE-RSA-AES256-SHA&peerName=myLinuxBoxName.withDomain` will connect to *ipOfYourLinuxBox* only if the system supports one of two listed ciphers, and if it presents a trusted certificate issued to *myLinuxBoxName.withDomain*. Replace *.withDomain* with the full domain name, for example, `example.org`.

Add a Serial Port to a Virtual Machine

You can connect the virtual serial port to a physical serial port or to a file on the host computer. You can also use a host-side named pipe to set up a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. In addition, you can use a port or vSPC URI to connect a serial port over the network. A virtual machine can use up to 32 serial ports.

- Verify that the virtual machine is powered off.
- Check that you know the correct media types for the port to access, vSPC connections, and any conditions that might apply. See [Using Serial Ports with vSphere Virtual Machines](#).
- To connect a serial port over a network, add a Firewall rule set. See [Adding a Firewall Rule Set for Serial Port Network Connections](#).
- To use authentication parameter with network serial port connections, see [Authentication Parameters for Virtual Serial Port Network Connections](#).
- Required privilege: **Virtual Machine > Configuration > Add or Remove Device**

Important: With virtual hardware version 11 and later, if you configure a virtual machine without serial ports, they are entirely removed from the virtual chipset and they are not visible to the virtual machine OS.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click the **Add New Device** button.
3. Select **Serial Port** from the drop-down menu.
The new serial port appears at the bottom of the device list.
4. From the **New Serial port** drop-down menu, select a connection type.

Option	Action
Use output file	Select this option to send output from the virtual serial port to a file on the host computer. To select an output file to connect to the serial port, click Browse .
Use physical serial port	Select this option to have the virtual machine use a physical serial port on the host computer. Select the serial port from the drop-down menu.
Use named pipe	Select this option to set a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. 1. Enter a name for the pipe in the Pipe Name field. 2. Select the Near end and Far end of the pipe from the drop-down menus.
Use Network	To connect through a remote network, select Use network . 1. Select the network backing. — To have the virtual machine monitor incoming connections from other hosts, select Server. — To have the virtual machine initiate a connection to another host, select Client. 2. Enter a Port URI. The URI is the remote end of the serial port to which the virtual machine's serial port should connect. 3. If vSPC is used as an intermediate step to access all virtual machines through a single IP address, select Use Virtual Serial Port Concentrator and enter the vSPC URI location.

5. Optional: Select **Yield CPU on poll**.

Select this option only for guest operating systems that use serial ports in polled mode. This option prevents the guest from consuming excessive CPUs.

6. Click **OK**.

Establishing Serial Port Network Connections to a Client or Server Without Authentication Parameters

If you do not use vSPC and you configure your virtual machine with a serial port connected as a server with a `telnet://:12345` URI, you can connect to your virtual machine's serial port from your Linux or Windows operating system.

```
telnetyourESXiServerIPAddress12345
```

Similarly, if you run the Telnet Server on your Linux system on port 23 (`telnet://yourLinuxBox:23`), you configure the virtual machine as a client URI.

```
telnet://yourLinuxBox:23
```

The virtual machine initiates the connection to your Linux system on port 23.

Other Virtual Machine Device Configuration

You can add and configure virtual hardware, such as DVD/CD-ROM drives to virtual machines.

Not all devices are available to add and configure. For example, you cannot add a video card, but you can configure available video cards and PCI devices.

For information about how to add, remove, and configure floppy drives or SCSI devices, see the VMware KB article at <https://kb.vmware.com/s/article/78978>.

Change the CD/DVD Drive Configuration

You can configure DVD or CD devices to connect to client devices, host devices, datastore ISO files, or content library ISO files.

Configure a Datastore ISO File for the CD/DVD Drive

To install a guest operating system and its applications on a new virtual machine, you can connect the CD/DVD device to an ISO file that is stored on a datastore accessible to the host.

Verify that you have the following privileges:

- **Virtual machine** > **Interaction** > **Configure CD media** on the virtual machine.
- **Datastore** > **Browse datastore** on the datastore to which you upload the installation media ISO image.
- **Datastore** > **Low level file operations** on the datastore to which you upload the installation media ISO image.

If an ISO image is not available on a local or shared datastore, upload the file to a datastore from your local system by using the datastore file browser. See [Upload ISO Image Installation Media for a Guest Operating System](#).

To avoid performance issues and possible conflicts between virtual machines that might try to simultaneously access the ISO image, unmount and disconnect the ISO file when the installation finishes.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Expand **CD/DVD drive**, and select **Datastore ISO File** from the drop-down menu.
The **Select File** dialog box opens
3. Browse to select the file and click **OK**.
4. From the **Virtual Device Node** drop-down menu, select the node that the drive uses in the virtual machine.
5. Optional: Select **Connect At Power On** to connect the device when the virtual machine powers on.

Note: If you install a guest operating system, select **Connect At Power On**.

6. Click **OK**.
7. Power on the virtual machine and click the **Summary** tab.
8. Expand the **VM Hardware** panel and click the **Connected** icon next to the datastore ISO file to connect the device

Configure a Content Library ISO File for the CD/DVD Drive

To install a guest operating system and its applications on a new virtual machine, you can connect the CD/DVD device to an ISO file that is stored in a content library.

Verify that you have the **Virtual machine** > **Interaction** > **Configure CD media** privilege on the virtual machine.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Expand **CD/DVD drive** and select **Content Library ISO File** from the drop-down menu.
The **Choose an ISO image to mount** dialog box opens.
3. Select the ISO file and click **OK**.
4. Optional: Select **Connect At Power On** to connect the device when the virtual machine powers on.
5. Click **OK**.
6. Power on the virtual machine and click the **Summary** tab.
7. Expand the **VM Hardware** panel and from the drop-down menu next to the **Connected** icon, select to connect the CD/DVD drive to a content library ISO file.

Configure a Host Device Type for the CD/DVD Drive

You can configure the CD/DVD drive of a virtual machine to connect to a physical CD or DVD device on the host so that you can install a guest operating system, VMware Tools, or other applications.

- Verify that the virtual machine is powered off.
- You cannot use vMotion to migrate virtual machines that have CD drives that are backed by the physical CD drive on the host. Disconnect these devices before you migrate the virtual machine.
- Verify that you have the **Virtual machine** > **Interaction** > **Configure CD media** privilege on the virtual machine.

When you create a virtual machine, a controller is added by default and the CD/DVD drive is attached to that controller. The controller and driver type depend on the guest operating system. Typically, virtual machines with newer guest operating systems have a SATA controller and CD/DVD drive. Other guests use an IDE controller and CD/DVD drive.

If you connect to media that does not require you to power off the virtual machine, you can select the media to connect to from the CD/DVD drive connection icon on the virtual machine **Summary** tab.

When you add a CD/DVD drive that is backed by a USB CD/DVD drive on the host, you must add the drive as a SCSI device.

For information how to add, remove, and configure SCSI devices, see <https://kb.vmware.com/s/article/78978>.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CD/DVD** and select **Host Device** from the drop-down menu.
3. Optional: Select **Connect At Power On** to connect the device when the virtual machine powers on.
4. If more than one type of CD/DVD media is available on the host, select the media.
5. In the **Virtual Device Node** drop-down menu, select the node the drive uses in the virtual machine.
The first available node is selected by default. You do not typically need to change the default.

6. Click **OK**.
7. Power on the virtual machine and click the **Summary** tab.

The connected CD/DVD device appears in the **VM Hardware** list.

Configure a Client Device Type for the CD/DVD Drive

To install a guest operating system and its applications or other media on a virtual machine, you can connect the CD/DVD device to a physical DVD or CD device on the system from which you access the vSphere Client.

Verify that the virtual machine is turned on.

By default, passthrough IDE mode is used for remote client device access. You can write or burn a remote CD only through passthrough mode access.

1. Navigate to a virtual machine in the inventory and click the **Summary** tab.
2. In the **VM Hardware** panel, click the **CD/DVD drive** connection icon, select an available drive to connect to, and browse for the CD/DVD media.
An Access Control dialog box opens. Click **allow** to proceed. To change your selection, click the connection icon, select **Disconnect**, and select a different option.

How do I Add or Modify a Virtual Machine CD or DVD Drive

CD/DVD drives are necessary for installing a guest operating system and VMware Tools.

- Verify that the virtual machine is powered off.
- If an ISO image file is not available on a local or shared datastore, upload an ISO image to a datastore from your local system by using the datastore file browser. See [Upload ISO Image Installation Media for a Guest Operating System](#).
- Verify that you have the **Virtual machine > Change Configuration > Add or remove device** privilege on the virtual machine.

When you power on the virtual machine, you can select the media to connect to from the **VM Hardware** panel on the virtual machine **Summary** tab.

The following conditions exist.

- If you add a CD/DVD drive that is backed by a USB CD/DVD drive on the host, you must add the drive as a SCSI device.
 - You must disconnect virtual machines that have CD drives that are backed by the physical CD drive on the host, before you migrate the virtual machine.
 - You access the host CD-ROM device through emulation mode. Passthrough mode is not functional for local host CD-ROM access. You can write or burn a remote CD only through passthrough mode access, but in emulation mode you can only read a CD-ROM from a host CD-ROM device.
1. Right-click a virtual machine in the inventory and select **Edit Settings**.
 2. Select your task.

Option	Description
Add a CD/DVD drive	On the Virtual Hardware tab, click the Add New Device button and select CD/DVD Drive .
Modify CD/DVD settings	On the Virtual Hardware tab, expand CD/DVD drive and change the configuration settings.

3. To change CD/DVD settings, select the device type from the **CD/DVD drive** drop-down menu.

Option	Action
Client Device	Select this option to connect the CD/DVD device to a physical DVD or CD device on the system from which you access the vSphere Client.

Option	Action
	From the Device Mode drop-down menu, select Passthrough CD-ROM .
Host Device	1. Select to connect the CD/DVD device to a physical DVD or CD device on the host. 2. From the CD/DVD Media drop-down menu, select the media to connect to. 3. From the Device Mode drop-down menu, select Emulate CD-ROM.
Datastore ISO File	Select this option to connect the CD/DVD device to an ISO file that is stored on a datastore accessible to the host. The Select File dialog box opens. 1. In the Select File dialog box, browse to the file containing the ISO image to connect to. 2. Click OK.
Content Library ISO File	Select this option to connect the CD/DVD device to an ISO file that is stored in a content library. The Choose an ISO image to mount dialog box opens 1. In the Choose an ISO image to mount, select the ISO image to connect to. 2. Click OK.

4. Optional: Specify additional settings for the CD/DVD drive.

Option	Description
Connect At Power On	Select this option to connect to the device when the virtual machine turns on.
Device Mode	Select Passthrough CD-ROM for a CD/DVD drive that is connected to the physical client machine. Select Emulate CD-ROM otherwise.
Virtual Device Node	Specify the location of the ISO that you are mounting. To change the device node from the default, select a new mode from the Virtual Device Node drop-down menu. For example, if you select SATA(0:2), the guest operating system detects the drive as ID 2 on controller 0.

5. Power on the virtual machine and click the **Summary** tab.
6. Expand the **VM Hardware** panel and click **Connected** next to select to.

You can now install the guest operating system or other applications.

How do I Add a PCI Device to a Virtual Machine

vSphere allows you to easily assign multiple PCI passthrough devices to a virtual machine without specifying an exact physical device on a particular ESXi host.

Verify that you have the privileges that you need for the task that you plan to perform.

- If you plan to add a PCI device when you edit a virtual machine, verify that you have the **Virtual machine > Change Configuration > Add or remove device** privilege.
- If you plan to increase the memory reservation when you edit a virtual machine, verify that you have the **Virtual machine > Change Configuration > Change resource** privilege.
- If you plan to reduce the virtual machine memory when you edit a virtual machine, verify that you have the **Virtual machine > Change Configuration > Change Memory** privilege.
- Power off the virtual machine.

- To use Dynamic DirectPath I/O, verify that the virtual machine is compatible with ESXi 7.0 or later.
- To use DirectPath, verify that Intel Virtualization Technology for Directed I/O (VT-d) or AMD I/O Virtualization Technology (IOMMU) is enabled in the host BIOS.
- Verify that the PCI devices are connected to the host and marked as available for passthrough. If your ESXi host is configured to boot from a USB device, or if the active coredump partition is configured to be on a USB device or SD cards connected through USB channels, deactivate the USB controller for passthrough. VMware does not support USB controller passthrough for ESXi hosts that boot from USB devices or SD cards connected through USB channels. A configuration in which the active coredump partition is configured to be on a USB device or SD card connected through USB channels is also not supported. For information, see <http://kb.vmware.com/kb/1021345>.
- To use Enhanced DirectPath I/O, verify that:
 - The virtual machine is compatible with ESXi 8.0 or later.
 - You download and install special drivers provided by the hardware vendor.
- To use NVIDIA GRID vGPU graphic devices:
 - Verify that an NVIDIA GRID vGPU graphic device with an appropriate driver is installed on the host. See the *VMware ESXi Upgrade* documentation.
 - Verify that the virtual machine is compatible with ESXi 6.0 and later.
- To add multiple NVIDIA GRID vGPUs to a virtual machine:
 - Verify that the virtual machine is compatible with ESXi 6.7 Update 2 and later.
 - Use only NVIDIA vGPU profiles with a maximum frame buffer.
 - Only Q-series and C-series vGPU types are supported.
- To use Vendor Device Groups, verify that:
 - The virtual machine is compatible with ESXi 8.0 or later.

You can connect to the guest operating system of a virtual machine all PCI devices that are configured on an ESXi host and made available for passthrough.

You can also take advantage of the vSphere Distributed Resource Scheduler (DRS) feature. vSphere DRS can move a virtual machine to a compatible ESXi host when the virtual machine powers on.

PCI vSphere DirectPath I/O devices

vSphere DirectPath I/O devices allow you to specify the physical location of the devices that you want to add to a virtual machine. vSphere DirectPath I/O allows a virtual machine to access directly the physical PCI and PCIe devices connected to a specific host. This way you can directly access devices, such as high-performance graphics or sound cards. You can connect each virtual machine to up to sixteen PCI devices.

You configure PCI devices on an ESXi host to make them available for passthrough to a virtual machine. See the *vSphere Networking* documentation. However, you must not enable PCI passthrough for ESXi hosts that are configured to boot from USB devices.

When PCI vSphere DirectPath I/O devices are made available to a virtual machine, you cannot perform certain operations on the virtual machine. These operations include suspending, migration with vMotion, and taking or restoring snapshots of the virtual machine.

vSphere Enhanced DirectPath I/O devices

vSphere Enhanced DirectPath I/O devices are an improvement that allows hardware device vendors to extend the capabilities of their devices when used in virtual machines. For example, if implemented by the hardware vendor, an Enhanced DirectPath I/O device might support a virtual machine migration with vMotion or suspend and resume operations. Enhanced DirectPath I/O devices might also have configuration parameters associated

with them and you can view and change these configuration parameters when you add such PCI devices to a virtual machine. The configuration parameters, if implemented by the vendor, can help you further specify and select particular hardware devices in accordance with the amount of the resources available on the device or their type.

vSphere Enhanced DirectPath I/O devices allow vSphere DRS to identify a host within the cluster that has an available device with the required resources.

PCI vSphere Dynamic DirectPath I/O devices

vSphere Dynamic DirectPath I/O devices allow you to specify the device type that you want to add to a virtual machine. vSphere Dynamic DirectPath I/O provides you with the ability to assign multiple PCI passthrough devices to a virtual machine. vSphere Dynamic DirectPath I/O allows vSphere DRS to identify a host within the cluster that has an available device with the same vendor and model name.

Vendor Device Groups

Vendor Device Groups is a collection of two or more hardware devices that you can allocate together to a virtual machine. By using Vendor Device Groups, other properties can be considered when allocating devices to a virtual machine. For example, one could ensure that a pair of devices are always allocated on the same PCI switch. Without Vendor Device Groups, each device will be allocated and assigned to a virtual machine independently of one another.

Note: You can only add and remove the entire Vendor Device Groups entity and not an individual device of that group.

If you select a Vendor Device Groups that has a NIC, you must add an appropriate network for the NIC in that Vendor Device Groups. The network that you add is one of the networks that you configure for the SR-IOV NICs to be attached to.

NVIDIA GRID GPU devices

If an ESXi host has an NVIDIA GRID GPU graphics device, you can configure a virtual machine to use the NVIDIA GRID virtual GPU (vGPU) technology.

NVIDIA GRID vGPU devices optimize complex graphics operations and make them run at high performance without overloading the CPU. By using NVIDIA GRID vGPU, you can share a single physical GPU among multiple virtual machines as separate vGPU-enabled passthrough devices.

Starting with vSphere 7.0 Update 2, you can configure a virtual machine to use the NVIDIA Multi-Instance GPU (MIG) feature. By using NVIDIA MIG, you can securely partition applicable GPUs into separate GPU instances. Each GPU instance has dedicated resources, such as memory, memory caches, and compute cores. If a GPU is in MIG mode, you can assign unique vGPU profile names to a virtual machine. VMware will create GPU and compute instances automatically, so you should not create them manually.

Note: When you add a PCI device to a virtual machine, the full memory size of the virtual machine is automatically reserved.

1. Add a PCI device to a virtual machine when you deploy a virtual machine or edit an existing virtual machine.

Option	Action
Create a new virtual machine	1. Select any inventory object that is a valid parent object of a virtual machine and click Actions .

Option	Action
	- Click New Virtual Machine. - On the Select a creation type page, select Create a new virtual machine, and click Next. - Navigate through the pages of the wizard. - On the Customize hardware page, click the Virtual Hardware tab.
Edit a virtual machine	- Select any inventory object that is a valid parent object of a virtual machine and click Actions. - Click Edit Settings. - Click the Virtual Hardware tab.

- On the **Virtual Hardware** tab, click the **Add New Device** button.
- From the drop-down menu, under **Other Devices**, select **PCI Device**.
The **Device Selection** window that contains a list of devices appears.
- Select the PCI device to connect to the virtual machine and click **Select**.

Option	Action
DirectPath I/O Dynamic DirectPath I/O Enhanced DirectPath I/O	- Select a PCI device from the list and click Select. The PCI device appears in the Virtual Hardware device list. - Expand the New PCI device section and if present, enter the configuration parameters for the PCI device. - To remove the PCI device, click Remove. The device is removed from the list. - To add another PCI device, click the Add Device button or undo the operation by clicking Undo.
NVIDIA GRID vGPU	- Select a PCI device from the list and click Select. The PCI device appears in the Virtual Hardware device list. - Expand the New PCI device section. - From the vGPU Profile drop-down menu, select the NVIDIA GRID vGPU passthrough device to connect to the virtual machine. - To remove the NVIDIA GRID vGPU passthrough device, click Remove. The device is removed from the list. - To add another NVIDIA GRID vGPU passthrough device, click the Add Device button or undo the operation by clicking Undo.
Group	- Select a Vendor Device Groups entry from the list and click Select. The Vendor Device Group devices appear in the Virtual Hardware device list. - Expand the New Vendor Device Group section. - If you add a Vendor Device Group that has a NIC, add an appropriate network. - From the Network drop-down menu click Browse, and select an appropriate network. - To remove the Vendor Device Group, click the vertical ellipsis icon and click Remove device. The Vendor Device Group is removed from the list. - To undo the operation, click Undo.

5. Click **Next** and complete the wizard.

6. Power on the virtual machine.

The connected PCI devices appear:

- On the **Hardware** tab of the **Edit Settings** wizard.
- On the **Summary** tab in the **PCI Devices** panel.

You can clone the virtual machine with the PCI passthrough devices.

How do I Configure 3D Graphics

When you create or edit a virtual machine, you can configure 3D graphics to take advantage of Windows AERO, CAD, Google Earth, and other 3D design, modeling, and multimedia applications.

Before you enable 3D graphics, become familiar with the available options and requirements.

How Enabling 3D Graphics Affects the Virtual Machine

You can use vMotion to migrate virtual machines that have 3D graphics enabled. If the 3D Renderer is set to Automatic, virtual machines use either the GPU on the destination host or a software renderer, depending on GPU availability. To migrate virtual machines with the 3D Renderer set to Hardware, the destination host must have a GPU.

You can set a group of virtual machines to use only Hardware rendering. For example, if you have virtual machines that run CAD applications or have other complex engineering capabilities, you might require that those virtual machines have persistent high-quality 3D capability present. When you migrate such virtual machines, the destination host must also have GPU capability. If the host does not have GPU, the migration cannot proceed. To migrate such virtual machines, you must turn them off and change the renderer setting to Automatic.

You can enable 3D on virtual machines that have Windows desktop or Linux guest operating systems. Not all guests support 3D graphics. To verify 3D support for a guest operating system, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Prerequisites

VMware supports AMD and NVIDIA graphics cards. See the vendor website for supported cards. To use the graphics card or GPU hardware, download the appropriate VMware graphics driver from the vendor website.

- Go to the NVIDIA website for information about the VMware graphics driver for your NVIDIA graphics card.
- Go to the AMD website for information about the VMware graphics driver for your AMD graphics card.

Linux distributions must have a 3.2 or later kernel. If 3D is not available on a Linux guest, verify that the driver is available in the Linux kernel. If it is not available, upgrade to a more recent Linux distribution. The location of the kernel depends on whether the distribution is based on `deb` or `rpm`.

Table 162: Linux Driver Location

VMware Linux Guest Kernel Drivers	Debian Format	RPM Format
vmwgfx.ko	<code>dpkg -S vmwgfx.ko</code>	<code>rpm -qf vmwgfx.ko</code>
vmwgfx_dri.so	<code>dpkg -S vmwgfx_dri</code>	<code>rpm -qf vmwgfx_dri</code>
vmware_drv.so	<code>dpkg -S vmware_drv</code>	<code>rpm -qf vmware_drv</code>

VMware Linux Guest Kernel Drivers	Debian Format	RPM Format
libxatracker.so.1	dpkg -S libxatracker	rpm -qf libxatracker

3D Rendering Options

You can select the 3D rendering options for each virtual machine to be Hardware, Software, or Automatic.

Table 163: 3D Rendering Options

Rendering Option	Description
Hardware	The virtual machine must have access to a physical GPU. If the GPU is not available, the virtual machine cannot power on.
Software	The virtual machine's virtual device uses a software renderer and will not attempt to use a GPU, even if one is present.
Automatic	The default setting. The virtual device selects whether to use a physical GPU or software-based rendering. If a GPU is available on the system and has the resources required by the virtual machine, the virtual machine uses the GPU. Otherwise software rendering is used.

Configure 3D Graphics and Video Cards

When you enable 3D graphics, you can select a hardware or software graphics renderer and optimize the graphics memory allocated to the virtual machine. You can increase the number of displays in multi-monitor configurations and change the video card settings to meet your graphics requirements.

- Verify that the virtual machine is powered off.
- Verify that the virtual machine compatibility is ESXi 5.0 and later.
- To enable 3D graphics in virtual machines with Windows 8 guest operating systems, the virtual machine compatibility must be ESXi 5.1 or later.
- To use a Hardware 3D renderer, ensure that graphics hardware is available. See [How do I Configure 3D Graphics](#).
- If you update the virtual machine compatibility from ESXi 5.1 and later to ESXi 5.5 and later, reinstall VMware Tools to get the latest SVGA virtual graphics driver and Windows Display Driver Model driver.
- Verify that you have the **Virtual machine > Change Configuration > Modify device settings** privilege on the virtual machine.

The default setting for total video RAM is adequate for minimal desktop resolution. For more complex situations, you can change the default memory. Typically, 3D applications require a video memory of 64–512 MB.

Fault Tolerance is not supported for virtual machines that have 3D graphics enabled.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Video Card**.
3. Select custom or automatic settings for your displays from the drop-down menu.

Option	Description
Auto-detect settings	Applies common video settings to the guest operating system.
Specify custom settings	Lets you select the number of displays and the total video memory.

4. Select the number of displays from the drop-down menu.
You can set the number of displays and extend the screen across them.
5. Enter the required video memory.
6. Optional: Select **Enable 3D support**.
This check box is active only for guest operating systems on which VMware supports 3D.
7. Optional: Select a 3D Renderer.

Option	Description
Automatic	Selects the appropriate option (software or hardware) for this virtual machine.
Software	Uses normal CPU processing for 3D calculations.
Hardware	Requires graphics hardware (GPU) for faster 3D calculations. Note: The virtual machine will not power on if graphics hardware is not available.

8. Click **OK**.

Sufficient memory allocation is set for this virtual machine's graphics.

Reduce Memory Overhead for Virtual Machines with 3D Graphics Option

Virtual machines with the 3D graphics option enabled can have higher memory consumption than other virtual machines. You can reduce the memory overhead by editing the configuration file (.vmx file) of your virtual machines and disabling certain memory-related settings. Reducing the memory overhead of virtual machines can help you increase the number of virtual machines per host.

Verify that your virtual machines are using hardware version 10 or later.

1. Shut down the virtual machine on which the 3D graphics option is enabled.
2. Deactivate the **Accelerate 3D Graphics** option.
3. Upgrade your ESXi host to use the features available in hardware version 10 or later.
4. Set the maximum size of your display to the size you need.
5. Locate the configuration file (.vmx) of your virtual machine.
6. Open the virtual machine configuration file in a text editor and add the line, `svga.vgaOnly=TRUE`.
This option removes all graphics and SVGA functionality from your SVGA device, but does not remove the settings that allow BIOS to enter VGA mode.
7. Save the changes and exit the text editor.
8. Power on your virtual machine and check the display console.
9. Verify the memory reservation settings in the `vmware.log` file.

How do I Add a Virtual Watchdog Timer Device to a Virtual Machine

To ensure self-reliance related to the system performance within a virtual machine, you can add a virtual Watchdog Timer (VWDT) device.

If the guest operating system stops responding and cannot recover on its own due to software glitches or errors, the VWDT waits for a predefined period of time and then restarts the system.

You can enable the VWDT to start either by the guest operating system, or by the BIOS or EFI firmware. If you chose the VWDT to start by the BIOS or EFI firmware, it starts before the guest operating system boots.

The VWDT has an important role in guest-based clustering solutions where each virtual machine in the cluster can recover on its own if it fails.

Add a Virtual Watchdog Timer Device to a Virtual Machine

To prevent the virtual machine from a guest operating system failure for an extended period of time, you can add a VWDT device to the virtual machine.

- Power off the virtual machine.
 - Verify that the virtual machine hardware is of version 17 or later.
 - Verify that the guest operating system of the virtual machine supports a watchdog timer:

Note: The guest operating system may require you to explicitly enable the watchdog services. Failure to do so might cause the watchdog timer to power off or restart the virtual machine.

 - The VWDT device implements the Watchdog Resource Table (WDRT) and Watchdog Action Table (WDAT) specifications. Windows Server 2003 supports devices compatible with WDRT and Windows Server 2008 and later supports devices compatible with WDAT.
 - The Linux distributions based on 4.9 or later kernel, such as Ubuntu 18.04 and Red Hat Enterprise Linux 7.6, support devices compatible with WDAT if the `wdat_wdt.ko` driver is available.
 - The watchdog timer is not supported for other guest operating systems, such as FreeBSD and Mac OS X.
 - Required privileges:
 - **Virtual Machine > Configuration > Add or remove device**
 - **Virtual machine > Configuration > Modify device settings**
1. Right-click a virtual machine from the vSphere inventory and select **Edit Settings**.
 2. On the **Virtual Hardware** tab, click **Add New Device** and select **Watchdog Timer** from the drop-down menu. The New Watchdog timer device appears in the **Virtual Hardware** devices list.
 3. To start the virtual watchdog timer with the BIOS or EFI firmware, select **Start with BIOS/EFI boot**. The virtual watchdog timer starts before the guest operating system boots.

If the guest operating system takes too long to boot or it does not support the VWDT device, a warning message appears, and the VWDT device might constantly restart the virtual machine.
 4. Click **OK**.

You can view the status of the VWDT device in the **VM Hardware** panel on the **Summary** tab.

How do I Add a Precision Clock Device to a Virtual Machine

A Precision Clock device is a virtual clock device that provides a virtual machine with access to the system time of the primary ESXi host.

- To ensure that the Precision Clock device provides the guest operating system of a virtual machine with accurate time, synchronize the primary ESXi host to use Network Time Protocol (NTP) or Precision Time Protocol (PTP). For information how to configure the host time synchronization, see the *vSphere Single Host Management - VMware Host Client* documentation.
- Power off the virtual machine.
- Verify that the virtual machine hardware is of version 17.

- Required privileges:
 - **Virtual Machine > Configuration > Add or remove device**
 - **Virtual machine > Configuration > Modify device settings**

To synchronize the guest operating system of a virtual machine with the host in an efficient manner, add a Precision Clock device to the virtual machine. For information on how to use the Precision Clock device as a reference clock for time synchronization of the supported guest operating system, see the *vCenter Server and Host Management* documentation.

1. Right-click a virtual machine from the vSphere Client inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click **Add New Device**, and from the drop-down menu select **Precision Clock**.
The Precision Clock device appears in the **Virtual Hardware** devices list.
3. Select the time synchronization protocol and click **OK**.

Option	Description
Any	The ESXi host time synchronization type has no limitations.
NTP	You can power on the virtual machine on an ESXi host that is configured with NTP time synchronization.
PTP	You can power on the virtual machine on an ESXi host that is configured with PTP time synchronization.


Virtual Shared Graphics Acceleration

Virtual Shared Graphics Acceleration (vSGA) enables sharing a physical GPU across multiple virtual desktops.

To provide GPU hardware acceleration to VMs, you can use one of the following graphics accelerations:

- Virtual Shared Graphics Acceleration (vSGA).
- Virtual Shared Graphics Passthrough Graphics Acceleration (vGPU/DVX).
- Virtual Dedicated Graphics Acceleration (vDGA).

With vSGA, the physical GPUs in the ESXi host are virtualized and shared across multiple VMs.

Figure 66: High-level architecture of vSGA

Graphics Acceleration Types Comparison

The following table compares the features of the three types of graphics acceleration.

Table 164: Types of graphics acceleration

	vSGA	vGPU/DVX	vDGA
Consolidation	High (limited by video memory)	Up to 1:32	None (1:1)
Performance Level	Lightweight to mid-level	Mid-level or Workstation	Workstation
Application Compatibility	Limited	All supported versions	All supported versions
Maximum DirectX Level	- vSphere 7.0 Update 2 or later - DirectX 11.0 / SM5.0	All supported versions	All supported versions
Maximum OpenGL version	- vSphere 8.0 or later - OpenGL 4.3	All supported versions	All supported versions
Video encoding and decoding	Software	Hardware	Hardware
OpenCL or CUDA compute	No	Yes	Yes
vSphere vMotion Support	Yes	Limited	No
Microsoft Virtualization-based Security (VBS)	Yes	No	No

Use Case for vSGA Graphics

vSGA is designed to address the knowledge worker use case where the graphics load is usually low. The knowledge workers use various non-specialized end-user applications that do not demand complex models to render.

Advantages of vSGA

The use of vSGA results in enhanced manageability, streamlined consolidation, and efficiency.

Manageability

vSGA supports the following features which allow an Administrator to easily manage their infrastructure at scale.

- Suspend and resume a VM.
- VM migration with vMotion.
- Guest operating system and ESXi host driver compatibility across vSphere versions.
- vSphere DRS for initial placement of VMs and load-balancing.

Consolidation

vSGA allows efficient use of hardware acceleration by leveraging VMs with different performance requirements on a single GPU.

The elasticity of vSGA allows a higher density of VMs per graphics card compared to vGPU and DVX solutions.

Hardware Requirements for Hardware Accelerated Graphics

For more information about the hardware requirements for graphics acceleration solutions, see the following table.

Table 165: Hardware requirements for graphics acceleration

Component	Description
Physical Space for Graphics Cards	Many high-end graphics cards are full-height, full-length, and double-width. Most of them take two slots on the motherboard but use a single PCIe x16 slot. To accommodate the selected GPU card in the appropriate PCIe slot, verify that the ESXi host has enough disk space.
Host Power Supply Units (PSU)	Check the power requirements of the GPU and cooling requirements of the server to make sure that the PSU is powerful enough and contains the proper power cables to power the GPU.
Two Display Adapters	If the ESXi host does not have an extra graphics adapter, you can install an additional low-end display adapter to act as a primary display adapter. If the GPU is set as a primary display for the ESXi host, vSGA cannot use the GPU for rendering. If two GPUs are installed, the server BIOS might have the option to select which GPU is primary and which is secondary.

Setup Your Environment for vSGA

To use vSGA you must configure the ESXi host and setup the VM.

Configure the ESXi host

To enable vSGA feature, configure the ESXi host by performing the following high-level steps.


Prerequisites

Verify that the ESXi host is compatible with ESXi 8.0 Update 2 or later.

Procedure

1. Install the graphics card on the ESXi host.
2. Place the host in maintenance mode. For more information, see [Place a Host in Maintenance Mode](#).
3. Install the associated GPU graphics card VIB files on the host. For example, run the following ESXCLI command.

```
esxcli software vib install -v $PWD/VendorName-vsga-driver_x.y.z-1OEM.a.b.c.vib
```
4. In the vSphere Client, verify that the GPU graphics card is in **Shared** mode.
 - a. Select the host and navigate to **Configure > Hardware > Graphics**.
 - b. Click the **Graphics Device** tab and check the **Configured Type** column.
 - c. To change the **vGPU Mode**, select the graphics card and click **Edit**.
 The **Edit Graphics Device Settings** window appears.



- d. Select **Shared** and click **OK**.
5. Reboot the host and exit the maintenance mode.
6. To verify that at least one driver VM starts, from the SSH terminal, run the `crx-cli list` command.

Setup the Virtual Machine General Settings

After you configure the ESXi host, you must setup the VM general settings.

Prerequisites

- Verify that the VM is powered off.
- Verify that the VM is not suspended.

Procedure

1. Right-click a VM in the inventory and select **Edit Settings**.

Table 166: General Settings of a VM

Setting	Description
Hardware Version	Make sure that the VM is of the latest hardware version that all hosts support. • Right-click the VM and select Compatibility. • If Upgrade VM Compatibility is not active, the VM is compatible with the latest hardware version. The minimum hardware version is 11. For DirectX 11.0 and OpenGL 4.3, the hardware version must be compatible with ESXi 8.0 and later.
CPU	The required number of CPUs depend on the usage and is determined by the actual workload. To obtain optimal performance, allocate 2 virtual CPUs to the VM.
Memory	The virtual hardware memory size depends on the usage and the number of applications that are running inside the VM. Set the memory size of the VM to • 4 GB if the VM runs 3D applications. • 2 GB if the VM runs basic productivity applications.
Network Adapter	Select the VMXNET 3 adapter type.
SCSI Controller	Select the LSI Logic SAS controller type. For the highest workloads using local flash-based storage, use VMware Paravirtual.
Other Devices	Remove all devices that are not used, such as COM, LTP, DVD, or floppy.

Configure a VM for Enabling vSGA

To enable vSGA, configure the VM by performing the following high-level steps.

Procedure

1. To enable 3D graphics, select the **Enable 3D Support** check box. For more information, see [How do I Configure 3D Graphics](#).
2. Set **3D Renderer** to **Automatic** or **Hardware**.
3. Select the size of the graphics **3D Memory**.
 - 512 MB if the VM runs basic productivity applications.
 - 2048 MB if the VM runs 3D applications. For maximum consolidation, set the **3D Memory** size to the minimum required by the guest OS applications. Guest operating systems typically cannot use more graphics memory than half of the available system memory. Therefore, to run larger graphics workloads, you must increase the Guest OS memory.

Configure a Windows VM for vSGA

After you configure the general settings for a VM, download and install the corresponding VMware Tools version in the VM.

Procedure

1. Install VMware Tools. For more information, see [Installing VMware Tools](#).
2. Restart the VM and run `dxdiag.exe`.
3. Click the **Display** tab. You should see the following parameters.

Table 167: Device panel and Driver panel parameters

Device Panel	Drivers Panel
• Name: VMware SVGA 3D • Device Type: Full Display Device	• Direct3D DDI: 11.1 • Feature levels: 11_0, 10_1, 10_0, 9_3, 9_2, 9_1 • Driver Model: WDDM 1.2

Table 168: VMware Tools version support for WDDM, DirectX, and OpenGL

VMware Tools Version	Max WDDM Version	Max DirectX Version	Max OpenGL Version
VMware Tools 10.0	WDDM 1.0	DirectX 10.0	OpenGL 3.3
VMware Tools 10.3.2	WDDM 1.0	DirectX 10.1	OpenGL 3.3
VMware Tools 11.1.5	WDDM 1.1	DirectX 11.0	OpenGL 4.1
VMware Tools 12.0	WDDM 1.2	DirectX 11.0	OpenGL 4.3

Configure a Linux VM for vSGA

After you configure the general settings for a VM, install Open VM Tools (open-vm-tools) in the Linux VM.

Procedure

1. Install Open VM Tools. For more information, see the *VMware Tools* documentation [Install Open VM Tools](#).
2. Verify that the VMware graphics driver is installed.

```
test@ubuntu:~/Desktop$ lsmod | grep vmwgfx
vmwgfx          323584  7
ttm             102400  1 vmwgfx
drm_kms_helper  225280  1 vmwgfx
drm             565248  10 vmwgfx,drm_kms_helper,ttm
test@ubuntu:~/Desktop$
```

3. To make sure that OpenGL applications work in the VM, install Mesa in the VM and verify that the OpenGL version is 3.3 or later.

For example, run the following command `glxinfo | grep "OpenGL version"`.

4. Verify the vmwgfx drm (direct rendering manager) driver version by using the `drm_info` tool.

For example, run the following script.

```
drm_info | head -n 2
Node: /dev/dri/card0
```

—Driver: vmwgfx (Linux drm driver for VMware graphics devices) version 2.20.0 (20211206)

Table 169: Linux DRM, Mesa, and OpenGL versions

Linux DRM Version	Mesa Version	OpenGL Version
DRM 2.9	Mesa 11.1.0	OpenGL 3.3
DRM 2.18	Mesa 20.2.0	OpenGL 4.1
DRM 2.20	Mesa 22.0.0	OpenGL 4.3

Securing Virtual Machines with Intel Software Guard Extensions

By using the vSphere Client, you can configure Virtual Intel® Software Guard Extensions (vSGX) for virtual machines and provide additional security to your workloads.

Some modern Intel CPUs implement a security extension called Intel® Software Guard Extensions (Intel SGX). Intel SGX is a processor-specific technology that defines private regions of memory, called enclaves. Intel SGX protects the enclave contents from disclosure or modification in such a way that code running outside the enclave cannot access them.

vSGX enables virtual machines to use Intel SGX technology if available on the hardware. To use vSGX, the ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the BIOS of the ESXi host. You can use the vSphere Client to enable SGX for a virtual machine.

Enable vSGX on a Virtual Machine

You can enable vSGX on a virtual machine when you deploy a virtual machine, edit or clone an existing virtual machine.

To use vSGX, your vSphere Client environment must meet a list of requirements:

- Virtual machine requirements:
 - EFI firmware
 - Hardware version 17 or later
 - Verify that the virtual machine is powered off
 - Verify that you have the privileges to create, clone, or edit virtual machine settings. For more information, see [Create a Virtual Machine with the New Virtual Machine Wizard](#) and [Clone an Existing Virtual Machine](#)
 - To enable remote attestation, verify that the virtual machine is of hardware version 20 or later
- Component requirements:
 - vCenter Server 7.0 and later
 - ESXi 7.0 or later
 - The ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the BIOS of the ESXi host. For information about the supported CPUs, see the VMware KB article at <https://kb.vmware.com/s/article/71367>.
 - To enable the remote attestation for the host, register the host with the Intel Registration Server. This way, the virtual machine running on the host can use the remote attestation. For more information on how to register a multi-socket ESXi, see the *vCenter Server and Host Management* documentation.
- Guest OS support:
 - Linux
 - Windows Server 2016 (64-bit) and later
 - Windows 10 (64-bit) and later

Note: Some operations and features are not supported for a virtual machine when vSGX is enabled.

- Migration with vMotion

- Migration with Storage vMotion
- Suspending or resuming the virtual machine
- Taking snapshot of the virtual machine, especially if you take a snapshot of the virtual machine memory
- Fault Tolerance
- Enabling Guest Integrity (GI, platform foundation for VMware AppDefense™ 1.0).

To use remote attestation for virtual machines using SGX enclaves, hosts with a single CPU socket do not require registration with the Intel Registration Server.

With vSphere 8.0, by enabling SGX host registration, you allow remote attestation for VMs running on multi-socket hosts.

1. You can enable SGX when you deploy a virtual machine or edit an existing virtual machine.

Option	Action
Deploy a virtual machine	1. Right-click any inventory object that is a valid parent object of a virtual machine and select New Virtual Machine. 2. On the Select a creation type page, select Create a new virtual machine, and click Next. 3. Navigate through the pages of the wizard. 4. On the Customize hardware page, click the Virtual Hardware tab.
Edit a virtual machine	1. Right-click a virtual machine in the inventory and select Edit Settings. 2. Click the Virtual Hardware tab.
Clone an existing virtual machine	1. Right-click a virtual machine in the inventory and select Clone > Clone to Virtual Machine. 2. Navigate through pages of the wizard. 3. On the Select clone options page, select Customize this virtual machine's hardware and click Next. 4. Click the Virtual Hardware tab.

2. On the **Virtual Hardware** tab, expand **Security Devices**.

Security Devices
SGX (1024 MB)

SGX (Unlocked)
☒ Enabled (Maximum: 4064 MB)

Enclave page cache size (MB)
1
GB

Remote attestation
☒ Required ⓘ

Restrictions

⚠ Some operations and features are restricted ⓘ

Launch control configuration
Unlocked

3. To enable SGX, select the **Enable** check box.
4. In the **Enclave page cache size (MB)** text box, enter the size of the cache size in MB.
Note: The enclave page cache size must be multiple of 2 MB.
5. To prevent the VM from powering on hosts that do not support SGX remote attestation, such as unregistered multi-socket SGX hosts, select the **Remote attestation** check box.
6. From the **Launch control configuration** drop-down menu, select the appropriate mode.

Option	Action
Unlocked	This option enables the launch enclave configuration of the guest operating system.
Locked	This option allows you to configure the launch enclave. 1. Select the Launch enclave public key hash option. 2. To use one of the public keys configured on the host, select Use from host and from the drop-down menu, select a public key hash. 3. To enter the public key manually, select Enter manually and enter a valid SHA256 hash (64) characters key.

7. Click **OK**.

Remove vSGX from a Virtual Machine

You can remove vSGX from a virtual machine.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Security Devices**.
3. Deselect the **Enable** check box for SGX and click **OK**.

You removed vSGX from the virtual machine. The vSGX no longer appears on the virtual machine **Summary** tab of the **VM Hardware** pane.

USB Configuration from an ESXi Host to a Virtual Machine

You can add multiple USB devices to a virtual machine when the physical devices are connected to an ESXi host. USB passthrough technology supports adding USB devices to virtual machines that reside on the host to which the devices are connected, such as security dongles and mass storage devices.

How USB Device Passthrough Technology Works

When you attach a USB device to a physical host, the device is available only to virtual machines that reside on that host. The device cannot connect to virtual machines that reside on another host in the data center.

A USB device is available to only one virtual machine at a time. When you connect a device to a powered on virtual machine, the device is not available to connect to other virtual machines that run on the host. When you remove the active connection of a USB device from a virtual machine, it becomes available to the other virtual machines that run on the host.

To connect a USB passthrough device to a virtual machine that runs on the ESXi host where the device is physically attached, you require an arbitrator, a controller, and a physical USB device or device hub.

USB Arbitrator

Manages connection requests and routes the USB device traffic. The arbitrator is installed and enabled by default on ESXi hosts. It scans the host for USB devices and manages the device connection among virtual machines that reside on the host. It

USB Controller

routes the device traffic to the correct virtual machine for delivery to the guest operating system. The arbitrator monitors the USB device and prevents other virtual machines from using it until you release it from the virtual machine it is connected to.

The USB hardware chip that provides a USB function to the USB ports that it manages. The virtual USB controller is the software virtualization of the USB host controller function in the virtual machine.

USB controller hardware and modules that support USB devices, such as USB 3.2 Gen 2x2, USB 3.1 SuperSpeedPlus, USB 3.1 SuperSpeed, and USB 2.0 must exist on the host. A controller must be present before you can add a USB device to the virtual machine.

The USB arbitrator can monitor a maximum of 15 USB controllers. Devices connected to controllers numbered 16 or greater are not available to the virtual machine.

USB Devices

You can add up to 20 USB devices to a virtual machine, which is the maximum number of devices supported for a simultaneous connection to one virtual machine. The maximum number of USB devices supported on a single ESXi host for a simultaneous connection to one or more virtual machines is also 20. For more information, see [Supported USB device models for passthrough from an ESX or ESXi host to a virtual machine](#).

USB 3.1 SuperSpeed Device Requirements

Starting with vSphere 5.5 Patch 3, USB 3.1 SuperSpeed devices are available for passthrough not only from a client computer to a virtual machine, but also from an ESXi host to a virtual machine. USB 3.1 SuperSpeed devices still have the following virtual machine configuration requirement:

- The virtual machine must have an enabled xHCI controller, Windows 8 or later, Windows Server 2012 and later, or a Linux guest operating system with a 2.6.35 or later kernel.

USB 3.1 SuperSpeedPlus Device Requirements

Starting with vSphere 7.0, USB 3.1 SuperSpeedPlus devices are available for passthrough at their maximum speed (SuperSpeedPlus), not only from a client computer to a virtual machine, but also from an ESXi host to a virtual machine. To operate their maximum transfer speed, USB 3.1 SuperSpeedPlus devices have the following virtual machine configuration requirements:

- The virtual machine must have an enabled xHCI controller, Windows 10 or later, Windows Server 2016 and later, or a Linux guest operating system with a 4.6 or later kernel.
- The virtual machine must be of hardware version 17 or later.
- For requirements and steps how to enable USB 3.1 SuperSpeedPlus, see the VMware knowledge base article <https://kb.vmware.com/s/article/70748>.

USB 3.2 Gen 2x2 Device Requirements

Starting with vSphere 8.0 Update 2, USB 3.2 Gen 2x2 devices are available for passthrough at their maximum speed of 20 Gbps, not only from a client computer to a virtual machine but also from an ESXi host to a virtual machine. To operate at their maximum data transfer speed, USB 3.2 Gen 2x2 devices have the following virtual machine configuration requirements:

- The virtual machine must have an enabled xHCI controller, Windows 10 or later, Windows Server 2016 or later, or a Linux guest operating system with a 4.6 or later kernel.
- The virtual machine must be of hardware version 21 or later.

How the USB Autoconnect Feature Works

When you add a USB device connection from an ESXi host to a virtual machine, the autoconnect feature is activated for the device connection. It is not deactivated until you remove the device connection from the virtual machine.

With autoconnect enabled, the device reconnects in the following cases:

- The virtual machine is cycling through power operations, such as Power Off/Power On, Reset, Pause/Resume.
- The device is unplugged from the host, and then plugged back in to the same USB port.
- The device is power cycled but has not changed its physical connection path.
- The device is mutating identity while it is in use.
- A new virtual USB device is added.

The USB passthrough autoconnect feature identifies the device by using the USB path of the device on the host. It uses the physical topology and port location instead of the device identity.

If you plug the same device back into a different USB port on the host, it cannot reestablish the connection with the virtual machine. If you unplug the device from the host and plug in a different device to the same USB path, the new device appears. It is connected to the virtual machine by the autoconnect feature that you enabled in the previous device connection.

Autoconnect is useful when devices mutate during usage. For example, for iPhone and other similar devices, the device VID/PID changes during software or firmware upgrades. The upgrade process disconnects and reconnects the devices to the USB port.

The USB port is speed-specific. If you change a USB device with another USB device that works with different speed, the autoconnect feature might not work. For example, you might connect a USB 2.0 high-speed device to a port and connect that device to the virtual machine. If you unplug the device from the host and plug in a USB 3.2 Gen 2x2, USB 3.1 SuperSpeed, or 3.1 SuperSpeedPlus device to the same port, the device might not connect to the virtual machine.

For a list of supported USB devices for passthrough from an ESXi host to a virtual machine, see [Supported USB device models for passthrough from an ESX or ESXi host to a virtual machine](#).

vSphere Features Available with USB Passthrough

Migrations with vMotion and DRS are supported with USB device passthrough from an ESXi host to a virtual machine.

Table 170: vSphere Features Available for USB Passthrough from an ESXi Host to a Virtual Machine

Feature	Supported with USB Device Passthrough
vSphere Distributed Power Management (DPM)	No
vSphere Distributed Resource Scheduler (DRS)	Yes
vSphere Fault Tolerance	No
vSphere vMotion	Yes

For details about migration with vMotion, see [Conditions and Limitations When Configuring USB Devices for vMotion](#).

If a host with connected USB devices resides in a DRS cluster with activated DPM, you must deactivate DPM for that host. Otherwise DPM might power off the host with the device, which disconnects the device from the virtual machine.

Conditions and Limitations When Configuring USB Devices for vMotion

With USB passthrough from a host to a virtual machine, you can migrate a virtual machine to another ESXi host in the same datacenter and maintain the USB passthrough device connections to the original host.

If a virtual machine has USB devices attached that pass through to an ESXi host, you can migrate that virtual machine with the devices attached.

For a successful migration, review the following conditions:

- You must configure all USB passthrough devices connected to a virtual machine for vMotion. If one or more devices is not configured for vMotion, the migration cannot proceed. For troubleshooting details, see the [Troubleshooting USB Passthrough Devices](#) documentation.
- When you migrate a virtual machine with attached USB devices away from the host to which the devices are connected, the devices remain connected to the virtual machine. However, if you suspend or power off the virtual machine, the USB devices are disconnected and cannot reconnect when the virtual machine is resumed. The device connections can be restored only if you move the virtual machine back to the host to which the devices are attached.
- If you resume a suspended virtual machine that has a Linux guest operating system, the resume process might mount the USB devices at a different location on the file system.
- If a host with attached USB devices resides in a DRS cluster with activated distributed power management (DPM), deactivate DPM for that host. Otherwise DPM might turn off the host with the attached device. This action disconnects the device from the virtual machine because the virtual machine migrated to another host.
- Remote USB devices require that the hosts be able to communicate over the management network following migration with vMotion, so the source and destination management network IP address families must match. You cannot migrate a virtual machine from a host that is registered to vCenter Server with an IPv4 address to a host that is registered with an IPv6 address.

Avoiding Data Loss with USB Devices

When a virtual machine connects to a physical USB device on an ESXi host, virtual machine functions can affect USB device behavior and connections.

- Before you hot add memory, CPU, or PCI devices, you must remove any USB devices. Hot adding these resources disconnects USB devices, which might result in data loss.
- Before you suspend a virtual machine, make sure that a data transfer is not in progress. During the suspend or resume process, USB devices behave as if they have been disconnected, then reconnected. For information about suspend and resume behavior after migration with vMotion, see [Conditions and Limitations When Configuring USB Devices for vMotion](#).
- Before you change the state of the arbitrator, make sure that USB devices residing on the host are not attached to a virtual machine. If USB devices become unavailable to a virtual machine, a host administrator might have deactivated the arbitrator. When an administrator stops or disconnects the arbitrator for troubleshooting or other purposes, USB devices attached to that host become unavailable to the virtual machine. If a data transfer is taking place at this time, you might lose the data. To reestablish the arbitrator, you must restart the host or restart the `usbarbitrator` and `hostd` services. Restarting the services requires that you power off and then power on the virtual machine.

Connecting USB Devices to an ESXi Host

You can connect and chain multiple USB hubs and devices to an ESXi host. Careful planning and knowledge of hub behavior and limitations can help ensure that your devices work optimally.

USB physical bus topology defines how USB devices connect to the host. Support for USB device passthrough to a virtual machine is available if the physical bus topology of the device on the host does not exceed tier seven. The first tier is the USB host controller and root hub. The last tier is the target USB device. You can cascade up to five tiers of external or

internal hubs between the root hub and the target USB device. An internal USB hub attached to the root hub or built into a compound device counts as one tier.

The quality of the physical cables, hubs, devices, and power conditions can affect USB device performance. To ensure the best results, keep the host USB bus topology as simple as possible for the target USB device, and use caution when you deploy new hubs and cables into the topology. The following conditions can affect USB behavior:

- Communication delay between the host and virtual machine increases as the number of cascading hubs increases.
- Connecting or chaining multiple external USB hubs increases device enumeration and response time, which can make the power support to the connected USB devices uncertain.
- Chaining hubs together also increases the chance of port and hub error, which can cause the device to lose connection to a virtual machine.
- Certain hubs can cause USB device connections to be unreliable, so use care when you add a new hub to an existing setup. Connecting certain USB devices directly to the host rather than to a hub or extension cable might resolve their connection or performance issues.

Note: To prevent additional problems, be aware of the physical constraints of long-term deployment in a machine room environment. Small devices are easily damaged by being stepped on or knocked loose.

In some cases, you must hard reset the device and hub to restore the device to a working state.

For a list of supported USB devices for passthrough from an ESXi host to a virtual machine, see the VMware knowledge base article at <http://kb.vmware.com/kb/1021345>.

USB Compound Devices

For compound devices, the virtualization process filters out the USB hub so that it is not visible to the virtual machine. The remaining USB devices in the compound appear to the virtual machine as separate devices. You can add each device to the same virtual machine or to different virtual machines if they run on the same host.

For example, the Aladdin HASP HL Drive USB dongle package contains three devices (0529:0001 HASP dongle, 13fe:1a00 Hub, 13fe:1d00 Kingston Drive). The virtualization process filters out the USB hub. The remaining Aladdin HASP HL Drive USB dongle devices (one Aladdin HASP dongle and one Kingston Drive) appear to the virtual machine as individual devices. You must add each device separately to make it accessible to the virtual machine.

Add USB Devices to an ESXi Host

You can connect multiple USB devices to an ESXi host so that the virtual machines that run on the same host can access the devices. The number of devices that you can connect depends on the device type and how the devices and hubs chain together.

- If a host has attached USB devices and resides in a DRS cluster with activated DPM, deactivate DPM for that host. For instructions about overriding the default DPM setting for a single host, see *vSphere Resource Management*.
- To verify that the virtual machine meets the requirements for the USB device connection, see [Connecting USB Devices to an ESXi Host](#).
- To add eight xHCI controllers to the ESXi host, verify that the current version of your ESXi host is 6.0 or later.

Each ESXi host has several USB ports. The number of ports on each host depends on the physical setup of the host. When you calculate the depth of hub chaining, remember that on a typical server the front ports connect to an internal hub.

The USB arbitrator can monitor a maximum of 15 USB controllers. If your system includes more than 15 controllers and you connect USB devices to them, the devices are not available to the virtual machine.

The host treats USB CD/DVD-ROM devices as SCSI devices.

- To add a USB device to an ESXi host, connect the device to an available port or hub.

You can add the device to the virtual machine. See [Add USB Devices from an ESXi Host to a Virtual Machine](#).

Add a USB Controller to a Virtual Machine

To support USB passthrough from an ESXi host or from a client computer to a virtual machine, you can add a USB controller to the virtual machine.

- Verify that the ESXi host has USB controller hardware and modules that support USB 3.2, 2.0, and 3.1 devices.
- Verify that the client computers have USB controller hardware and modules that support USB 3.2, 2.0, and 3.1 devices present.
- To use the xHCI controller on a Linux guest operating system, verify that the Linux kernel version is 2.6.35 or later.
- Verify that the virtual machine is powered on.
- Required Privilege (ESXi host passthrough): **Virtual Machine > Configuration > Add or Remove Device**.

In the vSphere Client, you can add one xHCI controller and one EHCI+UHCI controller.

- From hardware version 11 to hardware version 16, the supported number of root hub ports per xHCI controller is eight (four local USB 3.1 SuperSpeed ports and four logical USB 2.0 ports).
- From hardware version 17 to hardware version 20, the supported number of root hub ports per xHCI controller is eight (four local USB 3.1 SuperSpeedPlus ports and four logical USB 2.0 ports).
- With hardware version 21, the supported number of root hub ports per xHCI controller is eight (four local USB 3.2 Gen 2x2 ports and four logical USB 2.0 ports).

The conditions for adding a controller vary, depending on the device version, the type of passthrough (host or client computer), and the guest operating system.

Table 171: USB Controller Support

Controller Type	Supported USB Device Version	Supported for Passthrough from ESXi host to a VM	Supported for Passthrough from Client Computer to a VM
EHCI+UHCI	2.0	Yes	Yes
xHCI	3.2, 3.1, 2.0	Yes USB 3.2, 3.1, and 2.0	Yes Windows 8 or later, Windows Server 2012 and later, or a Linux guest operating system with a 2.6.35 or later kernel.

For Mac OS X systems, the EHCI+UHCI controller is enabled by default and is required for access to a USB mouse and keyboard.

For virtual machines with Windows or Linux guest operating systems, you can add one or two controllers of different types. You cannot add two controllers of the same type.

For USB passthrough from an host to a virtual machine, the USB arbitrator can monitor a maximum of 15 USB controllers. If your system includes more than 15 controllers and you connect USB devices to them, the devices are not available to the virtual machine.

1. Right-click a virtual machine from the vSphere inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click **Add New Device** and from the drop-down menu select **USB Controller**.
The controller appears in the **Virtual Hardware** devices list.
3. To change the USB controller type, expand **New USB Controller**.
If compatibility errors appear, you must fix them before you can add the controller.

4. Click **OK**.

Add one or more USB devices to the virtual machine.

Add USB Devices from an ESXi Host to a Virtual Machine

You can add one or more USB passthrough devices from an ESXi host to a virtual machine if the physical devices are connected to the host on which the virtual machine runs.

- Verify that the virtual machine is compatible with ESX/ ESXi 4.0 and later.
- Verify that a USB controller is present. See [Add a USB Controller to a Virtual Machine](#).
- To use vMotion to migrate a virtual machine with multiple USB devices, enable all attached USB devices for vMotion. You cannot migrate individual USB devices. For vMotion limitations, see [Conditions and Limitations When Configuring USB Devices for vMotion](#).
- When you add a CD/DVD-ROM drive that is backed by a USB CD/DVD drive on the host, add the drive as a SCSI device.
- Verify that you know the virtual machine requirements for USB devices. See [USB Configuration from an ESXi Host to a Virtual Machine](#).
- Required privileges: **Virtual Machine > Configuration > HostUSBDevice**

If a USB device is connected to another virtual machine, you cannot add it until that machine releases it.

Note:

If you have the Apple Frontpanel Controller device in your environment, you can safely add it to a virtual machine. However, this device has no documented function and no known use. ESXi hosts do not use it and do not provide Xserver functionality for USB passthrough.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click the **Add New Device** button.
3. Select **Host USB Device** from the drop-down menu.
The new USB device appears at the bottom of the Virtual Hardware device list.
4. Expand **New USB Device**, and select the device to add.
You can add multiple USB devices, but only one device at a time.
5. If you do not plan to migrate a virtual machine with USB devices attached, deselect the **Support vMotion** option.
This action reduces migration complexity, which results in better performance and stability.
6. Click **OK**.

Remove USB Devices That Are Connected Through an ESXi Host

When you remove USB devices from a virtual machine, devices that use the passthrough technology from a host become available to other virtual machines that run on that host.

- Verify that the devices are not in use.

- To minimize the risk of data loss, follow the instructions to safely unmount or eject hardware for your operating system. Safely removing hardware allows accumulated data to be transmitted to a file. Windows operating systems typically include a Remove Hardware icon located in the System Tray. Linux operating systems use the `umount` command.

Note: You might need to use the `sync` command instead of or in addition to the `umount` command, for example, after you issue a `dd` command on Linux or other UNIX operating systems.

1. Unmount or eject the USB device from the guest operating system.
2. Right-click a virtual machine in the inventory and select **Edit Settings**.
3. To remove the device, move the pointer over the device and click the **Remove** icon.
4. Click **OK** to save your changes.

Remove USB Devices from an ESXi Host

You can remove USB devices from the host if you must shut down the host for maintenance or if you do not want those devices to be available to virtual machines that run on the host. When you detach a USB device from the host, the device disconnects from the virtual machine.

Verify that the USB devices are not in use.

CAUTION: If data transfer is taking place when you remove USB devices from a host, you can lose data.

- Follow the device manufacturers instructions to safely remove the device.
When you remove the device from the host, it is no longer available to the virtual machines that run on the host.

USB Configuration from a Client Computer to a Virtual Machine

You can add multiple USB devices to a virtual machine when the physical devices are connected to the client computer on which you run the vSphere Client. The vSphere Client must be logged in to a vCenter Server instance that manages the ESXi host where the virtual machine resides. USB passthrough technology supports adding multiple USB devices, such as security dongles, mass storage devices, and smartcard readers to virtual machines.

How USB Device Passthrough Technology Works

The USB controller is the USB hardware chip that provides a USB function to the USB ports that it manages. The USB controller hardware and modules that support USB 3.2 Gen 2x2, USB 3.1 SuperSpeedPlus, USB 3.1 SuperSpeed, and USB 2.0 devices must be available on the virtual machine. Two USB controllers are available for each virtual machine. The controllers support multiple USB devices, such as USB 3.2 Gen 2x2, USB 3.1 SuperSpeedPlus, USB 3.1 SuperSpeed, and USB 2.0. The controller must be present before you can add USB devices to the virtual machine.

You can add up to 20 USB devices to a virtual machine, which is the maximum number of devices supported for a simultaneous connection to one virtual machine. You can add the devices only one at a time.

The virtual machine retains its connection to the device while it is in an S1 standby state. USB device connections are preserved when you migrate virtual machines to another host in the data center.

A USB device is available to only one powered on virtual machine at a time. When a virtual machine connects to a device, that device is no longer available to other virtual machines or to the client computer. When you disconnect the device from the virtual machine or shut down the virtual machine, the device returns to the client computer and becomes available to other virtual machines that the client computer manages.

For example, when you connect a USB mass storage device to a virtual machine, it is removed from the client computer and does not appear as a drive with a removable device. When you disconnect the device from the virtual machine, it reconnects to the client computer's operating system and is listed as a removable device.

Avoiding Data Loss

Before you connect a device to a virtual machine, make sure that the device is not in use on the client computer.

If the vSphere Client disconnects from the vCenter Server system or host, or if you restart or shut down the client computer, the device connection breaks. It is best to have a dedicated client computer for USB device use or to reserve USB devices connected to a client computer for short-term use, such as updating software or adding patches to virtual machines. To maintain USB device connections to a virtual machine for an extended time, use USB passthrough from an ESXi host to the virtual machine.

USB 3.1 SuperSpeed Device Requirements

Starting with vSphere 5.5 Patch 3, USB 3.1 SuperSpeed devices are available for passthrough not only from a client computer to a virtual machine, but also from an ESXi host to a virtual machine. USB 3.1 SuperSpeed devices still have the following virtual machine configuration requirement:

- The virtual machine must have an enabled xHCI controller, Windows 8 or later, Windows Server 2012 and later, or a Linux guest operating system with a 2.6.35 or later kernel.

USB 3.1 SuperSpeedPlus Device Requirements

Starting with vSphere 7.0, USB 3.1 SuperSpeedPlus devices are available for passthrough at their maximum speed (SuperSpeedPlus), not only from a client computer to a virtual machine, but also from an ESXi host to a virtual machine. To operate their maximum transfer speed, USB 3.1 SuperSpeedPlus devices have the following virtual machine configuration requirements:

- The virtual machine must have enabled xHCI controller, Windows 10 or later, Windows Server 2016 or later, or a Linux guest operating system with a 4.6 or later kernel.
- Verify that the virtual machine hardware is of version 17 or later.
- For requirements and steps how to enable the USB 3.1 SuperSpeedPlus, see the VMware knowledge base article <https://kb.vmware.com/s/article/70748>

USB 3.2 Gen 2x2 Device Requirements

Starting with vSphere 8.0 Update 2, USB 3.2 Gen 2x2 devices are available for passthrough at their maximum speed of 20 Gbps, not only from a client computer to a virtual machine but also from an ESXi host to a virtual machine. To operate at their maximum data transfer speed, USB 3.2 Gen 2x2 devices have the following virtual machine configuration requirements:

- The virtual machine must have an enabled xHCI controller, Windows 10 or later, Windows Server 2016 or later, or a Linux guest operating system with a 4.6 or later kernel.
- The virtual machine must be of hardware version 21 or later.

How to Connect USB Devices to a Client Computer

You can connect multiple USB devices to a client computer so that virtual machines can access the devices. The number of devices that you can add depends on several factors, such as how the devices and hubs chain together and the device type.

USB physical bus topology defines how USB devices connect to the client computer. Support for USB device passthrough to a virtual machine is available if the physical bus topology of the device on the client computer does not exceed tier seven. The first tier is the USB host controller and root hub. The last tier is the target USB device. You can cascade up to five tiers of external or internal hubs between the root hub and the target USB device. An internal USB hub attached to the root hub or built into a compound device counts as one tier.

The quality of the physical cables, hubs, devices, and power conditions can affect USB device performance. To ensure the best results, keep the client computer USB bus topology as simple as possible for the target USB device, and use caution when you deploy new hubs and cables into the topology. The following conditions can affect USB behavior:

- Connecting or chaining multiple external USB hubs increases device enumeration and response time, which can make the power support to the connected USB devices uncertain.
- Chaining hubs together increases the chance of port and hub error, which can cause the device to lose connection to a virtual machine.
- Certain hubs can cause USB device connections to be unreliable, so use care when you add a new hub to an existing setup. Connecting certain USB devices directly to the client computer rather than to a hub or extension cable might resolve their connection or performance issues. In some cases, you must remove and reattach the device and hub to restore the device to a working state.

The USB arbitrator can monitor a maximum of 15 USB controllers. If your system includes controllers that exceed the 15-controller limit and you connect USB devices to them, the devices are not available to the virtual machine.

USB Compound Devices

For compound devices, the virtualization process filters out the USB hub so that it is not visible to the virtual machine. The remaining USB devices in the compound appear to the virtual machine as separate devices. You can add each device to the same virtual machine or to different virtual machines if they run on the same host.

For example, the Aladdin HASP HL Drive USB dongle package contains three devices (0529:0001 HASP dongle, 13fe:1a00 Hub, 13fe:1d00 Kingston Drive). The virtualization process filters out the USB hub. The remaining Aladdin HASP HL Drive USB dongle devices (one Aladdin HASP dongle and one Kingston Drive) appear to the virtual machine as individual devices. You must add each device separately to make it accessible to the virtual machine.

Procedure

To add a USB device to a client computer, connect the device to an available port or hub.

What to do next:

You can now add the USB device to the virtual machine.

For more information on how to add a USB controller to a virtual machine, see [Add a USB Controller to a Virtual Machine](#).

Add USB Devices from a Client Computer to a Virtual Machine

You can add one or more USB passthrough devices from a client computer to a virtual machine in the vSphere Client. The devices must be connected to a client computer that connects to the ESXi host on which the virtual machine resides.

- Verify that a USB device is connected to the client computer.
- Verify that the virtual machine is powered on.
- Verify that a USB controller is present.

Note: If you connect to a USB device on a Mac OS X client computer, you can add only one device to the virtual machine at a time.

The devices maintain their virtual machine connections in S1 standby, if the vSphere Client is running and connected. After you add the USB device to the virtual machine, a message on the client computer states that the device is disconnected. The device remains disconnected from the client computer until you disconnect it from the virtual machine.

Fault Tolerance is not supported with USB passthrough from a client computer to a virtual machine.

1. In the vSphere Client, navigate to a virtual machine.
2. Launch the VMware Remote Console application.

Note: You cannot connect a USB device to a virtual machine if you use the HTML5 console in the vSphere Client.

3. In the VMware Remote Console toolbar, click **VMRC > Removable Devices** and find the USB device.
4. Click **Connect (Disconnect from menu)**.

The USB device is connected to the virtual machine.

Remove USB Devices That Are Connected Through a Client Computer

You can remove USB devices from a virtual machine if the devices are no longer needed. When you disconnect a USB device from a virtual machine, the device is released from the virtual machine and is given back to the client computer, which starts using it.

- Verify that the virtual machine is powered on.
- To minimize the risk of data loss, follow the instructions to safely unmount or eject hardware for your operating system. Safely removing hardware allows accumulated data to be transmitted to a file. Windows operating systems typically include a Remove Hardware icon located in the System Tray. Linux operating systems use the `umount` command.

Note: You might need to use the `sync` command instead of or in addition to the `umount` command, for example after you run a `dd` command on Linux or other UNIX operating systems.

- Required Privilege: **Virtual machine > Change Configuration > Add or remove device**

1. Unmount or eject the USB device from the guest operating system.
2. On the virtual machine **Summary** tab, click the disconnect icon on the right side of the USB device entry.
3. Select a device to disconnect from the drop-down menu.

A **Disconnecting** label and a spinner appear, indicating that a disconnection is in progress. When the device is disconnected, after a slight delay, the **Summary** tab refreshes and the device is removed from the virtual machine configuration.

The device reconnects to the client computer and is available to add to another virtual machine. In some cases, Windows Explorer detects the device and opens a dialog box on the client computer. You can close this dialog box.

Remove a USB Controller from a Virtual Machine

You can remove a USB controller from the virtual machine if you do not want to connect to USB devices.

Note: The ESXi host supports hot-removal of a USB controller. You can remove a USB controller from a virtual machine when both the ESXi host and the guest operating system of the virtual machine support the hot-removal functionality.

To make sure that the guest operating system of the virtual machine supports the hot-removal functionality, refer to the vendors documentation.

- Verify that the USB controller is not in use.

- Verify that all USB devices are disconnected from the virtual machine.
 - Required Privilege: **Virtual Machine > Configuration > Add or Remove Device**
1. Navigate to a data center, folder, cluster, resource pool, host, or vApp, click the **VMs** tab and click **Virtual Machines**.
 2. Right-click a virtual machine and click **Edit Settings**.
 3. On the **Virtual Hardware** tab, move the pointer over the USB controller and click the **Remove** icon.
 4. Click **OK** to confirm the deletion and close the dialog box.

The controller is no longer connected to the virtual machine, but remains available to add at a later time.

Add a Shared Smart Card Reader to Virtual Machines

You can configure multiple virtual machines to use a virtual shared smart card reader for smart card authentication. The smart card reader must be connected to a client computer on which the vSphere Client runs. All smart card readers are treated as USB devices.

- Verify that the smart card reader is connected to the client computer.
- Verify that the virtual machine is powered on.
- Verify that a USB controller is present.
- Required Privilege: **Virtual machine > Change Configuration > Add or remove device**

A license is required for the shared smart card feature. See *vCenter Server and Host Management*.

When you log out of Windows XP guest operating systems, to log back in, you must remove the smart card from the smart card reader and re-add it. You can also disconnect the shared smart card reader and reconnect it.

If the vSphere Client disconnects from the vCenter Server or host, or if the client computer is restarted or shut down, the smart card connection breaks. For this reason, it is best to have a dedicated client computer for smart card use.

To connect a USB smart card reader that is not shared, see [USB Configuration from a Client Computer to a Virtual Machine](#).

1. Navigate to a datacenter, folder, cluster, resource pool, host, or vApp, and click the **Related Options** tab and click **Virtual Machines**.
2. Select a virtual machine, click it again, and click the **Summary** tab.
3. Click the USB icon on the right side of **USB Devices** under **VM Hardware**, and select an available shared smart card reader from the drop down menu.

Select a device that appears as **Shared** *the model name of your smart card reader* followed by a number.

A **Connecting** label and a spinner appear showing that a connection is in progress. When the device has successfully connected and the Summary tab refreshes, the device is connected and the device name appears next to **USB Devices**.

You can now use smart card authentication to log in to virtual machines in the vSphere Client inventory.

Securing Virtual Machines with Virtual Trusted Platform Module

With the Virtual Trusted Platform Module (vTPM) feature, you can add a TPM 2.0 virtual cryptoprocessor to a virtual machine.

A vTPM is a software-based representation of a physical Trusted Platform Module 2.0 chip. A vTPM acts as any other virtual device. You can add a vTPM to a virtual machine in the same way you add virtual CPUs, memory, disk controllers, or network controllers. A vTPM does not require a hardware Trusted Platform Module chip.

What Is a Virtual Trusted Platform Module

A virtual Trusted Platform Module (vTPM) is a software-based representation of a physical Trusted Platform Module 2.0 chip. A vTPM acts as any other virtual device.

vTPMs provide hardware-based, security-related functions such as random number generation, attestation, key generation, and more. When added to a virtual machine, a vTPM enables the guest operating system to create and store keys that are private. These keys are not exposed to the guest operating system itself. Therefore, the virtual machine attack surface is reduced. Usually, compromising the guest operating system compromises its secrets, but enabling a vTPM greatly reduces this risk. These keys can be used only by the guest operating system for encryption or signing. With an attached vTPM, a client can remotely attest the identity of the virtual machine, and verify the software that it is running.

Note: By default, no storage policy is associated with a virtual machine that has been enabled with a vTPM. Only the virtual machine files (VM Home) are encrypted. If you prefer, you can choose to add encryption explicitly for the virtual machine and its disks, but the virtual machine files would have already been encrypted.

How Do You Configure a vTPM for a Virtual Machine

From the perspective of the virtual machine, a vTPM is a virtual device. You can add a vTPM to either a new or an existing virtual machine. A vTPM depends on virtual machine encryption to secure vital TPM data, thus, it requires that you configure a key provider. When you configure a vTPM, the virtual machine files are encrypted but not the disks. You can choose to add encryption explicitly for the virtual machine and its disks.

When you back up a virtual machine enabled with a vTPM, the backup must include all virtual machine data, including the *.nvram file. If your backup does not include the *.nvram file, you cannot restore a virtual machine with a vTPM. Also, because the VM home files of a vTPM-enabled virtual machine are encrypted, ensure that the encryption keys are available at the time of a restore.

In vSphere 8.0 and later, when cloning a virtual machine with a vTPM, selecting the **Replace** option for a virtual machine with a vTPM starts with a new, blank vTPM, which gets its own secrets and identity. When you replace the secrets of a vTPM, all keys, including workload-related keys, are replaced. As a best practice, ensure that your workloads no longer use a vTPM before you replace the keys. Otherwise, the workloads in the cloned virtual machine might not function correctly.

vSphere Requirements for vTPMs

To use a vTPM, your vSphere environment must meet these requirements:

- Virtual machine requirements:
 - EFI firmware
 - Hardware version 14 and later
- Component requirements:
 - vCenter Server 6.7 and later for Windows virtual machines, vCenter Server 7.0 Update 2 and later for Linux virtual machines.
 - Virtual machine encryption (to encrypt the virtual machine home files).
- Guest OS support:
 - Linux
 - Windows Server 2008 and later
 - Windows 7 and later

Differences Between a Hardware TPM and a Virtual TPM

You use a hardware Trusted Platform Module (TPM) to provide secure storage of credentials or keys. A vTPM performs the same functions as a TPM, but it performs cryptographic coprocessor capabilities in software. A vTPM uses the *.nvram file, which is encrypted using virtual machine encryption, as its secure storage.

A hardware TPM includes a preloaded key called the Endorsement Key (EK). The EK has a private and public key. The EK provides the TPM with a unique identity. For a vTPM, this key is provided either by the VMware Certificate Authority (VMCA) or by a third-party Certificate Authority (CA). After the vTPM uses a key, it is typically not changed because doing so invalidates sensitive information stored in the vTPM. The vTPM does not contact the third-party CA at any time.

Create a Virtual Machine with a Virtual Trusted Platform Module

You can add a Virtual Trusted Platform Module (vTPM) when you create a virtual machine to provide enhanced security to the guest operating system. You must create a key provider before you can add a vTPM.

- The guest OS you use can be Windows Server 2008 and later, Windows 7 and later, or Linux.
- The ESXi hosts running in your environment must be ESXi 6.7 or later (Windows guest OS), or 7.0 Update 2 (Linux guest OS).
- The virtual machine must use EFI firmware.
- Verify that you have the required privileges:
 - **Cryptographic operations > Clone**
 - **Cryptographic operations > Encrypt**
 - **Cryptographic operations > Encrypt new**
 - **Cryptographic operations > Migrate**
 - **Cryptographic operations > Register VM**

Note: After creating a virtual machine with a vTPM, the **Cryptographic operations > Direct Access** privilege is required to open a console session.

The VMware virtual TPM is compatible with TPM 2.0 and creates a TPM-enabled virtual chip for use by the virtual machine and the guest OS it hosts.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, an host or a cluster.
3. Right-click the object, select **New Virtual Machine**, and follow the prompts to create a virtual machine.

Option	Action
Select a creation type	Create a new virtual machine.
Select a name and folder	Specify a name and target location.
Select a compute resource	
Select storage	Select a compatible datastore.
Select compatibility	You must select ESXi 6.7 and later for Windows guest OS, or ESXi 7.0 U2 and later for Linux guest OS.
Select a guest OS	Select Windows or Linux for use as the guest OS.
Customize hardware	Click Add New Device and select Trusted Platform Module . You can further customize the hardware, for example, by changing disk size or CPU.
Ready to complete	Review the information and click Finish .

The vTPM-enabled virtual machine appears in your inventory as specified.

Add Virtual Trusted Platform Module to an Existing Virtual Machine

You can add a Virtual Trusted Platform Module (vTPM) to an existing virtual machine to provide enhanced security to the guest operating system. You must create a key provider before you can add a vTPM.

- The guest OS you use can be Windows Server 2008 and later, Windows 7 and later, or Linux.

- Verify that the virtual machine is turned off.
- The ESXi hosts running in your environment must be ESXi 6.7 and later (Windows guest OS), or 7.0 Update 2 and later (Linux guest OS).
- The virtual machine must use EFI firmware.
- Verify that you have the required privileges:
 - **Cryptographic operations > Clone**
 - **Cryptographic operations > Encrypt**
 - **Cryptographic operations > Encrypt new**
 - **Cryptographic operations > Migrate**
 - **Cryptographic operations > Register VM**
 - **Virtual machine > Change Configuration > Add or remove device**

Note: After adding a vTPM to a virtual machine, the **Cryptographic operations > Direct Access** privilege is required to open a console session.

The VMware virtual TPM is compatible with TPM 2.0, and creates a TPM-enabled virtual chip for use by the virtual machine and the guest OS it hosts.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. In the **Edit Settings** dialog box, click **Add New Device** and select **Trusted Platform Module**.
4. Click **OK**.

The **Virtual Machine Details** pane reflects that encryption has been applied to the virtual machine.

Remove Virtual Trusted Platform Module from a Virtual Machine

You can remove Virtual Trusted Platform Module (vTPM) security from a virtual machine.

- Ensure that the virtual machine is powered off.
- Verify that you have the required privileges: **Virtual machine > Change Configuration > Add or remove device** and **Cryptographic operations > Decrypt**

Removing a vTPM device causes all encrypted information on the virtual machine to become unrecoverable. Before removing a vTPM from a virtual machine, deactivate any applications in the Guest OS that use the vTPM device, such as BitLocker. Failure to do so can cause the virtual machine not to boot. Also, you cannot remove a vTPM from a virtual machine that contains snapshots.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. In the **Virtual Hardware** tab, expand **Security Devices**.
4. Click the ellipses icon for Virtual TPM.
5. Click **Remove device**.
6. Click **Delete** to confirm you want to remove the vTPM.

The vTPM device is marked for removal.

7. Click **OK**.

Identify Virtual Trusted Platform Module Enabled Virtual Machines

You can identify which of your virtual machines are enabled to use a Virtual Trusted Platform Module (vTPM).

You can generate a list of all virtual machines in your inventory showing virtual machine name, operating system, and vTPM status. You can also export this list to a CSV file for use in compliance audits.

1. Connect to vCenter Server by using the vSphere Client.
2. Select a vCenter Server instance, a host, or a cluster.
3. Click the **VMs** tab and click **Virtual Machines**.
4. To view all virtual machines on which a TPM is enabled, click **Manage Columns** and select **TPM**.
The TPM column displays "Present" for virtual machines on which a TPM is enabled. Virtual machines without a TPM are listed as "Not present."
5. To export the contents of an inventory list view to a CSV file, click **Export**.

Securing Virtual Machines with AMD Secure Encrypted Virtualization-Encrypted State

Secure Encrypted Virtualization-Encrypted State (SEV-ES) is a hardware feature enabled in recent AMD CPUs that keeps the guest operating system's memory and register state encrypted, protecting it against access from the hypervisor.

You can add SEV-ES to your virtual machines as an extra security enhancement. SEV-ES prevents CPU registers from leaking information in registers to components like the hypervisor. SEV-ES can also detect malicious modifications to a CPU register state.

vSphere and AMD Secure Encrypted Virtualization-Encrypted State

In vSphere 7.0 Update 1 and later, you can activate Secure Encrypted Virtualization-Encrypted State (SEV-ES) on supported AMD CPUs and guest operating systems.

Currently, SEV-ES supports only AMD EPYC 7x2 CPUs (code named "Rome") and later CPUs, and only versions of Linux kernels that include specific support for SEV-ES.

SEV-ES Components and Architecture

The SEV-ES architecture consists of the following components.

- AMD CPU, specifically, the Platform Security Processor (PSP) that manages encryption keys and handles encryption.
- Enlightened operating system, that is, an operating system that uses guest-initiated calls to the hypervisor.
- Virtual Machine Monitor (VMM) and Virtual Machine Executable (VMX), to initialize an encrypted virtual machine state during virtual machine power-on, and also to handle calls from the guest operating system.
- VMkernel driver, to communicate unencrypted data between the hypervisor and the guest operating system.

Implementing and Managing SEV-ES on ESXi

You must first activate SEV-ES in a system's BIOS configuration. See the documentation for your system for more information about accessing the BIOS configuration. After you have activated SEV-ES in the BIOS for your system, you can then add SEV-ES to a virtual machine.

You use either the vSphere Client (in vSphere 7.0 Update 2 and later) or PowerCLI commands to activate and deactivate SEV-ES on virtual machines. You can create new virtual machines with SEV-ES, or activate SEV-ES on existing virtual

machines. Privileges to manage virtual machines activated with SEV-ES are the same as for managing regular virtual machines.

Unsupported VMware Features on SEV-ES

The following features are not supported when SEV-ES is activated.

- System Management Mode
- vMotion
- Powered-on snapshots (however, no-memory snapshots are supported)
- Hot add or remove of CPU or memory
- Suspend/resume
- VMware Fault Tolerance
- Clones and instant clones
- Guest Integrity
- UEFI Secure Boot

Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the vSphere Client

In vSphere 7.0 Update 2 and later, you can use the vSphere Client to add SEV-ES to a virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.
- SEV-ES must be enabled in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When enabling SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.

Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES enabled virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES enabled virtual machines per ESXi host.

- The ESXi host running in your environment must be at ESXi 7.0 Update 1 or later.
- The vCenter Server must be at vSphere 7.0 Update 2 or later.
- The guest operating system must support SEV-ES.
Currently, only Linux kernels with specific support for SEV-ES are supported.
- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option enabled, otherwise power-on fails.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, an ESXi host or a cluster.
3. Right-click the object, select **New Virtual Machine**, and follow the prompts to create a virtual machine.

Option	Action
Select a creation type	Create a virtual machine.
Select a name and folder	Specify a name and target location.
Select a compute resource	Specify an object for which you have privileges to create virtual machines.
Select storage	In the VM storage policy, select the storage policy. Select a compatible datastore.
Select compatibility	Ensure that ESXi 7.0 and later is selected.

Option	Action
Select a guest OS	Select Linux, and select a version of Linux with specific support for SEV-ES.
Customize hardware	Under VM Options > Boot Options > Firmware , ensure that EFI is selected. Under VM Options > Encryption , select the Enable check box for AMD SEV-ES.
Ready to complete	Review the information and click Finish .

The virtual machine is created with SEV-ES.

Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the vSphere Client

In vSphere 7.0 Update 2 and later, you can use the vSphere Client to add SEV-ES to an existing virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.
- SEV-ES must be activated in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When activating SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.

Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES activated virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES activated virtual machines per ESXi host.

- The ESXi host running in your environment must be at ESXi 7.0 Update 1 or later.
- The vCenter Server must be at vSphere 7.0 Update 2 or later.
- The guest operating system must support SEV-ES.

Currently, only Linux kernels with specific support for SEV-ES are supported.

- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option checked, otherwise power-on fails.
- Ensure that the virtual machine is powered off.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. Under **VM Options > Boot Options > Firmware**, ensure that EFI is selected.
4. In the **Edit Settings** dialog box, under **VM Options > Encryption**, select the **Enable** check box for AMD SEV-ES.
5. Click **OK**.

SEV-ES is added to the virtual machine.

Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the Command Line

You can use the command line to add SEV-ES to a virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.

- SEV-ES must be enabled in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When enabling SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.
Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES enabled virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES enabled virtual machines per ESXi host.
- The ESXi host running in your environment must be at ESXi 7.0 Update 1 or later.
- The guest operating system must support SEV-ES.
Currently, only Linux kernels with specific support for SEV-ES are supported.
- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option enabled, otherwise power-on fails.
- PowerCLI 12.1.0 or later must be installed on a system with access to your environment.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator to the vCenter Server that manages the ESXi host on which you want to add a virtual machine with SEV-ES.

```
Connect-VIServer -server vCenter_Server_ip_address -User admin_user -Password 'password'
```

2. Create the virtual machine with the `New-VM` cmdlet, specifying `-SEVEnabled $true`.

For example, first assign the host information to a variable, then create the virtual machine.

```
$vmhost = Get-VMHost -Name 10.193.25.83
New-VM -Name MyVM1 $vmhost -NumCPU 2 -MemoryMB 4 -DiskMB 4 -SEVEnabled $true
```

If you must specify the virtual hardware version, run the `New-VM` cmdlet with the `-HardwareVersion vmx-18` parameter. For example:

```
New-VM -Name MyVM1 $vmhost -NumCPU 2 -MemoryMB 4 -DiskMB 4 -SEVEnabled $true -HardwareVersion vmx-18
```

The virtual machine is created with SEV-ES.

Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the Command Line

You can use the command line to add SEV-ES to an existing virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.
- SEV-ES must be activated in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When activating SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.
Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES activated virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES activated virtual machines per ESXi host.
- The ESXi host running in your environment must be ESXi 7.0 Update 1 or later.
- The guest operating system must support SEV-ES.
Currently, only Linux kernels with specific support for SEV-ES are supported.
- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option checked, otherwise power-on fails.
- PowerCLI 12.1.0 or later must be installed on a system with access to your environment.

- Ensure that the virtual machine is powered off.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator to the vCenter Server that manages the ESXi host with the virtual machine to which you want to add SEV-ES.

For example:

```
Connect-VIServer -server vCenter_Server_ip_address -User admin_user -Password 'password'
```

2. Add SEV-ES to the virtual machine with the `Set-VM` cmdlet, specifying `-SEVEnabled $true`.

For example:

```
$vmhost = Get-VMHost -Name 10.193.25.83
Set-VM -Name MyVM2 $vmhost -SEVEnabled $true
```

If you must specify the virtual hardware version, run the `Set-VM` cmdlet with the `-HardwareVersion vmx-18` parameter. For example:

```
Set-VM -Name MyVM2 $vmhost -SEVEnabled $true -HardwareVersion vmx-18
```

SEV-ES is added to the virtual machine.

Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the vSphere Client

In vSphere 7.0 Update 2 and later, you can use the vSphere Client to deactivate SEV-ES on a virtual machine.

- Ensure that the virtual machine is powered off.
1. Connect to vCenter Server by using the vSphere Client.
 2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
 3. In the **Edit Settings** dialog box, under **VM Options > Encryption**, deselect the **Enable** check box for AMD SEV-ES.
 4. Click **OK**.

SEV-ES is deactivated on the virtual machine.

Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the Command Line

You can use the command line to deactivate SEV-ES on a virtual machine.

- Ensure that the virtual machine is powered off.
 - PowerCLI 12.1.0 or later must be installed on a system that has access to your environment.
1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator to the vCenter Server that manages the ESXi host with the virtual machine from which you want to remove SEV-ES.

For example:

```
Connect-VIServer -server vCenter_Server_ip_address -User admin_user -Password 'password'
```

2. Deactivate SEV-ES on the virtual machine with the `Set-VM` cmdlet, specifying `-SEVEnabled $false`.

For example, first assign the host information to a variable, then deactivate SEV-ES for the virtual machine.

```
$vmhost = Get-VMHost -Name 10.193.25.83
Set-VM -Name MyVM2 $vmhost -SEVEnabled $false
```

SEV-ES is deactivated on the virtual machine.

Configuring Virtual Machine Options

You can set or change virtual machine options to run VMware Tools scripts, control user access to the remote console, configure startup behavior, and more.

The virtual machine options define a range of virtual machine properties, such as the virtual machine name and the virtual machine behavior with the guest operating system and VMware Tools.

You can view or change virtual machine settings from the vSphere Client. Not all options are available to every virtual machine and some options rarely must change from their defaults.

The host that the virtual machine runs on and the guest operating system must support any configurations that you make.

You can view and change virtual machine settings on the **VM Options** tab of the **Edit Settings** wizard.

You can select one of the following options.

Table 172: Virtual Machine Options in the vSphere Client

Options	Description
General Options	In this section, you can view or change the following settings. - Virtual machine name - Virtual machine configuration file location - Virtual machine working location - Guest operating system and OS version Currently, you can only edit the virtual machine name. The information about the other settings is read only. To change the operating system for a VM, you have to reinstall the OS - or consider deploying a new VM with your operating system of choice.
VMware Remote Console Options	In this section, you can change the locking behavior of a virtual machine and the settings for simultaneous connections.
Encryption	In this section, you can change the encryption settings of a virtual machine.
Power Management	In this section, you can change virtual machine suspend behavior.
VMware Tools	In this section, you can change the behavior of VMware Tools scripts. You can also customize the automatic VMware Tools upgrades, automatically synchronize the guest time of the virtual machine on startup or resume with the host, and periodically synchronize the guest time with the host.
Virtualization Based Security	Activate or deactivate VBS for the virtual machine.
Boot Options	In this section, you can change the virtual machine boot options. For example, add a delay before booting, force entry into the BIOS or EFI setup screen, or set reboot options.
Advanced	In this section, you can change the following advanced virtual machine options. - Acceleration and logging settings - Debugging and statistics - Swap file location - Latency sensitivity
Fibre Channel NPIV	In this section, you can change the virtual node and port World Wide Names (WWNs).

General Virtual Machine Options

View or change general virtual machine settings, such as the name and location of the virtual machine, configuration file location, and operating system.

Change the Virtual Machine Name

A virtual machine must have a name that is unique within the folder where the virtual machine is located. If you move a virtual machine to a different datastore folder or host that has an existing virtual machine of the same name, you must change the virtual machine's name to keep it unique.

When you change the name of a virtual machine, you change the name used to identify the virtual machine in the vCenter Server inventory. This action does not change the name used as the computer name by the guest operating system.

The virtual machine name also determines the name of the virtual machine files and folder on the disk. For example, if you name the virtual machine win8, the virtual machine files are named win8.vmx, win8.vmdk, win8.nvram, and so on. If you change the virtual machine name, the names of the files on the datastore do not change.

Note: Migration with Storage vMotion changes the virtual machine filenames on the destination datastore to match the inventory name of the virtual machine. The migration renames all virtual disk, configuration, snapshot, and .nvram files. If the new names exceed the maximum filename length, the migration does not succeed.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click the **VM Options** tab and expand **General Options**.
3. Delete the existing name and enter a new name for the virtual machine in the **VM Name** text box.
4. Click **OK**.

View the Virtual Machine Configuration and Working File Location

You can view the location of the virtual machine configuration and working files. You can use this information when you configure backup systems.

Verify that the virtual machine is powered off.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click **VM Options** tab and expand **General Options**.
The path to the location of the virtual machine configuration file appears in the **VM Config File** text box. The path to the virtual machine working location appears in the **VM Working Location** text box.

Change the Configured Guest Operating System

When you create a virtual machine, the system uses a template to specify all the parameters necessary to run the guest OS you selected so that the virtual machine can run correctly. For example, does the guest OS support hot adding devices, memory, and CPUs. After you create a virtual machine, you can add or modify various parameters. For example, you can add host bus adapters, NICs, memory sizing, storage devices, and so on.

Power off the virtual machine.

The directory where the virtual machine is stored includes a file with a .VMX extension (VMX file). The VMX file contains parameters that define and control the virtual machine.

One of the parameters in the VMX file is the guestOS parameter. This parameter does not specify the guest OS that the virtual machine uses; it determines how the low-level (implicit) parts of the virtual machine are initialized so that it can run the guest OS.

The guestOS parameter and the other parameters established through the creation template are defined and tested to ensure that the guest OS will operate adequately.

Changing the guestOS parameter is not recommended, as it can lead to undesirable results. This change only affects the guestOS parameter; no other parameters will be changed, added, or removed. The virtual machine may appear to run

correctly initially, but in many cases, it will fail to power on due to incompatibilities between the guestOS parameter and other settings in the VMX file.

Note:

Do not change the guestOS parameter when the virtual machine is powered on.

You can create two virtual machines – one for the current guestOS type and one for the destination guestOS. Review the virtual machine creation templates to identify the parameters you must add, delete, or change for the virtual machine. Changing the guestOS parameter of the newly created virtual machine is the customer's responsibility. VMware does not validate the modified virtual machine's suitability.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click the **VM Options** tab and expand **General Options**.
3. From the **Guest OS** drop-down menu, select the guest operating system family.
4. From the **Guest OS Version** drop-down menu, select the guest operating system version.
5. Click **OK**.

VMware Remote Console Options

Change the VMware Remote Console options to control the access to the virtual machine.

Change the Virtual Machine Console Options for Remote Users

You can limit the number of simultaneous connections to a virtual machine and lock the guest operating system when the last remote user disconnects from the virtual machine console.

- Verify that VMware Tools is installed and running.
- To use the **Guest OS lock** option, verify that you have a Windows XP or later guest operating system.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click the **VM Options** tab, and expand **VMware Remote Console Options**.
3. Optional: Select the **Guest OS lock** check box to lock the guest operating system when the last remote user disconnects.
4. Optional: In the **Maximum number of sessions** text box, specify the number of simultaneous connections to the virtual machine.
5. Click **OK**.

Configure Virtual Machine Encryption Options

Starting with vSphere 7.0, you can take advantage of virtual machine encryption. Encryption protects not only your virtual machine but also virtual machine disks and other files.

You set up a trusted connection between vCenter Server and a key management server (KMS). vCenter Server can then retrieve keys from the KMS as needed.

For detailed information about virtual machine encryption, see the *vSphere Security* documentation.

Encrypt an Existing Virtual Machine or Virtual Disk

You can encrypt an existing virtual machine or virtual disk by changing its storage policy. You can encrypt virtual disks only for encrypted virtual machines.

- Establish a trusted connection with the KMS and select a default KMS.

- Create an encryption storage policy, or use the bundled sample, VM Encryption Policy.
 - Ensure that the virtual machine is powered off.
 - Verify that you have the required privileges:
 - **Cryptographic operations > Encrypt new**
 - If the host encryption mode is not Enabled, you also need **Cryptographic operations > Register host**.
1. Connect to vCenter Server by using the vSphere Client.
 2. Right-click the virtual machine that you want to change and select **VM Policies > Edit VM Storage Policies**.
You can set the storage policy for the virtual machine files, represented by VM home, and the storage policy for virtual disks.
 3. Select the storage policy.
 - To encrypt the VM and its hard disks, select an encryption storage policy and click **OK**.
 - To encrypt the VM but not the virtual disks, toggle on **Configure per disk**, select the encryption storage policy for VM Home and other storage policies for the virtual disks, and click **OK**.

You cannot encrypt the virtual disk of an unencrypted virtual machine. However, if you use the vSphere Client to encrypt the VM Home files, you can then reconfigure the unencrypted virtual machine with the encrypted disk.
 4. If you prefer, you can encrypt the virtual machine, or both virtual machine and disks, from the **Edit Settings** menu in the vSphere Client.
 - a) Right-click the virtual machine and select **Edit Settings**.
 - b) Select the **VM Options** tab, and open **Encryption**. Choose an encryption policy. If you deselect all disks, only the VM home is encrypted.
 - c) Click **OK**.

Decrypt an Encrypted Virtual Machine or Virtual Disk

You can decrypt a virtual machine, its disks, or both, by changing the storage policy.

- The virtual machine must be encrypted.
- The virtual machine must be powered off or in maintenance mode.
- Required privileges: **Cryptographic operations > Decrypt**

This task describes how to decrypt an encrypted virtual machine using the vSphere Client.

All encrypted virtual machines require encrypted vMotion. During virtual machine decryption, the Encrypted vMotion setting remains. To change this setting so that Encrypted vMotion is no longer used, change the setting explicitly.

This task explains how to perform decryption using storage policies. For virtual disks, you can also perform decryption using the **Edit Settings** menu.

Note: In the Virtual Machine Details pane, a vTPM-enabled virtual machine displays both a lock icon and an "Encrypted with *key_provider*" message. To remove a vTPM from a virtual machine, see [Remove Virtual Trusted Platform Module from a Virtual Machine](#).

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine that you want to change and select **VM Policies > Edit VM Storage Policies**.
You can set the storage policy for the virtual machine files, represented by VM home, and the storage policy for virtual disks.

3. Select a storage policy.

- To decrypt the VM and its hard disks, toggle off **Configure per disk**, select a storage policy from the drop-down menu, and click **OK**.
- To decrypt a virtual disk but not the virtual machine, toggle on **Configure per disk**, select the encryption storage policy for VM Home and other storage policies for the virtual disks, and click **OK**.

You cannot decrypt the virtual machine and leave the disk encrypted.

4. If you prefer, you can use the vSphere Client to decrypt the virtual machine and disks from the **Edit Settings** menu.

- Right-click the virtual machine and select **Edit Settings**.
- Select the **VM Options** tab and expand **Encryption**.
- To decrypt the VM and its hard disks, choose **None** from the **Encrypt VM** drop-down menu.
- To decrypt a virtual disk but not the virtual machine, deselect the disk.
- Click **OK**.

5. Optional: You can change the Encrypted vMotion setting.

- Right-click the virtual machine and click **Edit Settings**.
- Click **VM Options**, and open **Encryption**.
- Set the **Encrypted vMotion** value.

Clone an Encrypted Virtual Machine

When you clone an encrypted virtual machine, the clone is encrypted with the same keys unless you change them. To change keys, you can use the vSphere Client, the PowerCLI, or the API.

- A key provider must be configured and enabled.
- Create an encryption storage policy, or use the bundled sample, VM Encryption Policy.
- Required privileges:
 - **Cryptographic operations > Clone**
 - **Cryptographic operations > Encrypt**
 - **Cryptographic operations > Decrypt**
 - **Cryptographic operations > Recrypt**
 - If the host encryption mode is not **Enabled**, you also must have **Cryptographic operations > Register host** privilege.

If you use the PowerCLI or the API, you can clone the encrypted virtual machine and change keys in one step. For more information, see *vSphere Web Services SDK Programming Guide*.

You can perform the following operations during clone.

- Create an encrypted virtual machine from an unencrypted virtual machine or virtual machine template.
- Create a decrypted virtual machine from an encrypted virtual machine or virtual machine template.
- Recrypt the destination virtual machine with different keys from that of source virtual machine.
- Starting in vSphere 8.0, selecting the **Replace** option for a virtual machine with a vTPM device starts with a new, blank vTPM, which gets its own secrets and identity.

Note: vSphere 8.0 includes the `vpdx.clone.tpmProvisionPolicy` advanced setting to make the default clone behavior for vTPMs to be "replace."

You can create an instant clone virtual machine from an encrypted virtual machine with the caveat that the instant clone shares the same key with the source virtual machine. You cannot recrypt keys on either the source or the instant clone virtual machine. See *vSphere Web Services SDK Programming Guide*.

1. Navigate to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and select **Clone > Clone to Virtual Machine > .**
3. Navigate through pages of the wizard.
 - a) On the **Select a name and folder** page, enter a name and select a data center or folder in which to deploy it.
 - b) On the **Select a compute resource**, select an object for which you have privileges to create encrypted virtual machines. For information about prerequisites and required privileges for encryption tasks, see the *vSphere Security* documentation.
 - c) Change the keys for the cloned vTPM.

Cloning a virtual machine duplicates the entire virtual machine, including the vTPM and its secrets, which can be used to determine a system's identity. To change secrets on a vTPM, select **Replace** for **TPM Provision Policy**.

Note:

When you replace the secrets of a vTPM, all keys, including workload-related keys, are replaced. As a best practice, ensure that your workloads no longer use a vTPM before you replace the keys. Otherwise, the workloads in the cloned virtual machine might not function correctly.

- d) On the **Select storage** page, select the datastore or datastore cluster in which to store the template configuration files and all virtual disks. You can change the storage policy as part of the clone operation. For example, changing from using an encryption to non-encryption policy decrypts the disks.
 - e) On the **Select clone options**, select additional customization options.
 - f) On the **Ready to complete** page, review the information and click **Finish**.
4. Optional: Change the keys for the cloned virtual machine.

By default, the cloned virtual machine is created with the same keys as its parent. Best practice is to change the cloned virtual machine keys to ensure that multiple virtual machines do not have the same keys.

 - a) Decide upon a shallow or deep recrypt.

To use a different DEK and KEK, perform a deep recrypt of the cloned virtual machine. To use a different KEK, perform a shallow recrypt of the cloned virtual machine. For a deep recrypt, you must power off the virtual machine. You can perform a shallow recrypt operation while the virtual machine is powered on, and if the virtual machine has snapshots present. Shallow recrypt of an encrypted virtual machine with snapshots is permitted only on a single snapshot branch (disk chain). Multiple snapshot branches are not supported. If the shallow recrypt fails before updating all links in the chain with the new KEK, you can still access the encrypted virtual machine if you have the old and new KEKs.
 - b) Perform a recrypt of the clone using the API. For more information, see *vSphere Web Services SDK Programming Guide*.

Virtual Machine Power Management Options

Configure virtual machine Power management options to define how the virtual machine responds when the guest OS is placed on standby.

Manage Power Management Settings for a Virtual Machine

If the guest operating system is placed on standby, the VM can either remain powered on or be suspended. You can use the Power Management settings to control this behavior. Some desktop-based guests, such as Windows 7, have standby enabled by default, so that the guest goes into standby after a predetermined time.

Note: To avoid having the guest operating system go into standby mode unintentionally, verify the settings before you deploy the virtual machine.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click **VM Options** and expand **Power Management**.
3. In the **Standby response** section, select the standby response of the virtual machine.
 - The **Suspend the virtual machine** option stops all processes, which saves resources, and copies the contents of the virtual machine's memory to the virtual machine's `.vmss` file. Writing the memory to the `.vmss` file is useful if you need to copy the file to help with a troubleshooting scenario.
 - The **Put the guest operating system in standby mode and leave the virtual machine powered on** option stops all processes, but leaves the virtual devices connected to the virtual machine.
4. In the **Wake on LAN** section, select one or multiple virtual network adapters to which to apply the Wake on LAN option.
5. To save your changes, click **OK**.

Configuring VMware Tools Options

Configure the VMware Tools options to define the power operations for the virtual machine and decide when to run VMware Tools Scripts. Through VMware Tools configuration, you can automatically synchronize the virtual machine guest operating system time with the host.

Configure the Virtual Machine Power States

Changing virtual machine power states is useful when you do maintenance on the host. You can use the system default settings for the virtual machine power controls, or you can configure the controls to interact with the guest operating system. For example, you can configure the **Power off** control to power off the virtual machine or shut down the guest operating system.

- Verify that you have privileges to perform the intended power operation on the virtual machine.
- To set optional power functions, install VMware Tools in the virtual machine.
- Power off the virtual machine before editing the VMware Tools options.

You can modify many virtual machine configurations while the virtual machine is running, but you might need to change the virtual machine power state for some configurations.

You cannot configure a **Power on** () action. This action powers on a virtual machine when a virtual machine is stopped, or resumes the virtual machine and runs a script when it is suspended and VMware Tools is installed and available. If VMware Tools is not installed, it resumes the virtual machine and does not run a script.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click the **VM Options** tab and expand **VMware Tools**.
3. Select an option for the virtual machine **Power Off** () control from the drop-down menu.

Option	Description
Power Off	Immediately stops the virtual machine. The power off action shuts down the guest operating system or powers off the virtual machine. A message indicates that the guest

Option	Description
	operating system might not shut down properly. Use this power off option only when necessary.
Shut Down Guest (Default)	Follows system settings. The current value of the system settings appears in parentheses. Uses VMware Tools to initiate an orderly system shut down of the virtual machine. Soft power operations are possible only if the tools are installed in the guest operating system.

4. Select an option for the **Suspend** () control from the drop-down menu.

Option	Description
Suspend (Default)	Suspends the virtual machine and leaves it connected to the network.
Suspend Guest	Follows system settings. The current value of the system setting appears in parentheses. Pauses all virtual machine activity. When VMware Tools is installed and available, the suspend action runs a script in the guest operating system and suspends the virtual machine. If VMware tools is not installed, the suspend action suspends the virtual machine without running a script.

5. Select an option for the **Reset** () control from the drop-down menu.

Option	Description
Reset	Shuts down and restarts the guest operating system without powering off the virtual machine. If VMware Tools is not installed, the reset action resets the virtual machine.
Default (Restart Guest)	Follows system settings. The current value of the system setting appears in parentheses. Uses VMware Tools to initiate an orderly restart. Soft power operations are possible only if the tools are installed in the guest operating system.

6. To save your changes, click **OK**.

Synchronize the Time of a Virtual Machine Guest Operating System with the Host

You can configure VMware Tools options to set the time of a virtual machine guest operating system to be the same as the time of the host. Starting with vSphere 7.0 Update 1, VMware Tools provides accurate and synchronized time between guest and host operating systems.

To synchronize the virtual machine guest operating system time with that on the host, install VMware Tools on the virtual machine.

You can synchronize the time between guest and host operating systems while the virtual machine is running. This operation does not depend on the hardware version of the virtual machine.

For information about the timekeeping best practices for Windows guest operating systems, see <https://kb.vmware.com/s/article/1318>.

For information about the timekeeping best practices for Linux guest operating systems, see <https://kb.vmware.com/s/article/1006427>.

1. To synchronize the time of the virtual machine guest operating system with the host time, open VMware Tools options for your virtual machine while creating, editing, cloning or deploying a virtual machine.

Option	Action
Create a virtual machine	1. Right-click any inventory object that is a valid parent object of a virtual machine and select New Virtual Machine. 2. On the Select a creation type page, select Create a new virtual machine and click Next. 3. Navigate through the pages of the wizard.

Option	Action
	4. On the Customize hardware page, click the VM Options tab and expand VMware Tools .
Edit a virtual machine	1. Right-click a virtual machine in the inventory and select Edit Settings. 2. Click the VM Options tab and expand VMware Tools.
Clone an existing virtual machine	1. Right-click a virtual machine in the inventory and select Clone > Clone to Virtual Machine. 2. Navigate through the pages of the wizard. 3. On the Select clone options page, select Customize this virtual machine's hardware and click Next. 4. On the Customize hardware page, click the VM Options tab and expand VMware Tools.
Deploy a virtual machine from a template	1. Right-click a template in the inventory and select New VM from This Template. 2. Navigate through the pages of the wizard. 3. On the Select clone options page, select Customize this virtual machine's hardware and click Next. 4. On the Customize hardware page, click the VM Options tab and expand VMware Tools.

2. Select the time synchronization options.

Option	Action
Synchronize at startup and resume (Default)	This option is selected by default to ensure the best time synchronization between guest and host operating systems after performing certain operations, for example: • When you resume a virtual machine from a suspended operation. • When you migrate a virtual machine with vMotion. • When you revert a snapshot.
Synchronize time periodically	Periodically synchronize the time of the guest operating system of a virtual machine with the host. Note: Select this option if the guest operating system of a virtual machine does not have a native time synchronization software.

3. To save your changes, click **OK**.

Virtualization Based Security

Microsoft VBS, a feature of Windows 10, Windows Server 2016, and Windows Server 2019 operating systems, uses hardware and software virtualization to enhance system security by creating an isolated, hypervisor-restricted, specialized subsystem. Starting with vSphere 6.7, you can enable Microsoft virtualization-based security (VBS) on supported Windows guest operating systems.

For more detailed information about VBS, see the *vSphere Security* documentation.

Activate Virtualization-based Security on an Existing Virtual Machine

You can activate Microsoft virtualization-based security (VBS) on existing virtual machines for supported Windows guest operating systems.

See the *vSphere Security* documentation for information about acceptable CPUs and VBS best practices.

Using Intel CPUs for VBS requires vSphere 6.7 or later. The virtual machine must have been created using hardware version 14 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit) or later releases
- Windows Server 2016 (64 bit) or later releases

Using AMD CPUs for VBS requires vSphere 7.0 Update 2 or later. The virtual machine must have been created using hardware version 19 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit), version 1809 or later releases
- Windows Server 2019 (64 bit) or later releases

Ensure that you install the latest patches for Windows 10, version 1809, and Windows Server 2019, before activating VBS.

For more information about activating VBS on virtual machines on AMD platforms, see the VMware KB article at <https://kb.vmware.com/s/article/89880>.

Configuring VBS is a process that involves first activating VBS in the virtual machine then activating VBS in the guest OS.

Note: New virtual machines configured for Windows 10, Windows Server 2016, and Windows Server 2019 on hardware versions less than version 14 are created using Legacy BIOS by default. If you change the virtual machine's firmware type from Legacy BIOS to UEFI, you must reinstall the guest operating system.

1. In the vSphere Client, browse to the virtual machine.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click the **VM Options** tab.
4. Select the **Enable** check box for Virtualization Based Security.
5. Click **OK**.

The Virtual Machine Details tile under the **Summary** tab displays "Virtualization Based Security - Enable".

See [Activate Virtualization-based Security on the Guest Operating System](#).

Activate Virtualization-based Security on the Guest Operating System

You can activate Microsoft virtualization-based security (VBS) for supported Windows guest operating systems.

- Ensure that virtualization-based security has been activated on the virtual machine.

You activate VBS from within the Windows Guest OS. Windows configures and enforces VBS through a Group Policy Object (GPO). The GPO gives you the ability to turn off and on the various services, such as Secure Boot, Device Guard, and Credential Guard, that VBS offers. Certain Windows versions also require you to perform the additional step of enabling the Hyper-V platform.

See Microsoft's documentation about deploying Device Guard to activate virtualization-based security for details.

1. In Microsoft Windows, edit the group policy to turn on VBS and choose other VBS-related security options.
2. Optional: For Microsoft Windows versions less than Redstone 4, in the Windows Features control panel, enable the Hyper-V platform.
3. Reboot the guest operating system.

Deactivate Virtualization-based Security

If you no longer use virtualization-based security (VBS) with a virtual machine, you can deactivate VBS. When you deactivate VBS for the virtual machine, the Windows VBS options remain unchanged but might induce performance issues. Before deactivating VBS on the virtual machine, deactivate VBS options within Windows.

Ensure that the virtual machine is powered off.

1. In the vSphere Client, browse to the virtual machine that is using VBS.
See [Identify VBS-Enabled Virtual Machines](#) for help in locating virtual machines that use VBS.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click **VM Options**.
4. Deselect the **Enable** check box for Virtualization Based Security.
A message reminds you to deactivate VBS in the guest OS.
5. Click **OK**.
6. Verify that the virtual machine's **Summary** tab no longer displays "Virtualization Based Security - Enable" in the Guest OS description.

Identify VBS-Enabled Virtual Machines

You can identify which of your virtual machines have VBS enabled, for reporting and compliance purposes.

1. Connect to vCenter Server by using the vSphere Client.
2. Select a vCenter Server instance, a data center, or a host in the inventory.
3. Click the **VMs** tab and click **Virtual Machines**.
4. To show the **VBS** column, click **Manage Columns** and select the **VBS** check box.
5. Scan for "Present" in the **VBS** column.

Configure Virtual Machine Boot Options

Edit Boot Options to activate or deactivate UEFI Secure Boot and configure the boot behavior of the virtual machine.

Activate or Deactivate UEFI Secure Boot for a Virtual Machine

UEFI Secure Boot is a security standard that helps ensure that your PC boots using only software that is trusted by the PC manufacturer. For certain virtual machine hardware versions and operating systems, you can activate secure boot just as you can for a physical machine.

You can activate secure boot only if all prerequisites are met. If prerequisites are not met, the check box is not visible in the vSphere Client.

- Verify that the virtual machine operating system and firmware support UEFI boot.
 - EFI firmware
 - Virtual hardware version 13 or later.
 - Operating system that supports UEFI secure boot.

Note: Some guest operating systems do not support changing from BIOS boot to UEFI boot without guest OS modifications. Consult your guest OS documentation before changing to UEFI boot. If you upgrade a virtual machine that already uses UEFI boot to an operating system that supports UEFI secure boot, you can activate Secure Boot for that virtual machine.

- Turn off the virtual machine. If the virtual machine is running, the check box is dimmed.

In an operating system that supports UEFI secure boot, each piece of boot software is signed, including the bootloader, the operating system kernel, and operating system drivers. The virtual machine's default configuration includes several code signing certificates.

- A Microsoft certificate that is used only for booting Windows.
- A Microsoft certificate that is used for third-party code that is signed by Microsoft, such as Linux bootloaders.
- A VMware certificate that is used only for booting ESXi inside a virtual machine.

The virtual machine's default configuration includes one certificate for authenticating requests to modify the secure boot configuration, including the secure boot revocation list, from inside the virtual machine, which is a Microsoft KEK (Key Exchange Key) certificate.

VMware Tools version 10.1 or later is required for virtual machines that use UEFI secure boot. You can upgrade those virtual machines to a later version of VMware Tools when it becomes available.

For Linux virtual machines, VMware Host-Guest Filesystem is not supported in secure boot mode. Remove VMware Host-Guest Filesystem from VMware Tools before you activate secure boot.

Note: If you turn on secure boot for a virtual machine, you can load only signed drivers into that virtual machine.

This task describes how to use the vSphere Client to activate and deactivate secure boot for a virtual machine. You can also write scripts to manage virtual machine settings. For example, you can automate changing the firmware from BIOS to EFI for virtual machines with the following PowerCLI code:

```
$vm = Get-VM TestVM

$spec = New-Object VMware.Vim.VirtualMachineConfigSpec
$spec.Firmware = [VMware.Vim.GuestOsDescriptorFirmwareType]::efi
$vm.ExtensionData.ReconfigVM($spec)
```

See *VMware PowerCLI User's Guide* for more information.

1. Browse to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click the **VM Options** tab, and expand **Boot Options**.
4. Under **Boot Options**, ensure that firmware is set to **EFI**.
5. Select your task.
 - Select the **Secure Boot** check box to activate secure boot.
 - Deselect the **Secure Boot** check box to deactivate secure boot.
6. Click **OK**.

When the virtual machine boots, only components with valid signatures are allowed. The boot process stops with an error if it encounters a component with a missing or invalid signature.

Delay the Boot Sequence

Delaying the boot operation is useful when you change BIOS or EFI settings such as the boot order. For example, you can change the BIOS or EFI settings to force a virtual machine to boot from a CD-ROM.

- Verify that vSphere Client is logged in to a vCenter Server.
 - Verify that you have access to at least one virtual machine in the inventory.
 - Verify that you have privileges to edit boot options for the virtual machine.
1. Right-click a virtual machine in the inventory and select **Edit Settings**.
 2. Click **VM Options** tab and expand **Boot Options**.
 3. In the **Boot Delay** section, select the time in milliseconds to delay the boot operation.
 4. Optional: In the **Force setup** section, select whether to force entry into the BIOS or EFI setup screen the next time the virtual machine boots.
 5. Optional: In the **Failed Boot Recovery** section, select whether the virtual machine should reboot after a boot failure and enter the time in seconds.
 6. Click **OK**.

Configuring Virtual Machine Advanced Options

You can edit the Advanced virtual machines settings when you need to solve issues caused by an application or when you need log files and debugging information for troubleshooting purposes.

You can also add or change configuration parameters and change the latency sensitivity of a virtual machine.

Enable Virtual Machine Logging

You can enable logging to collect log files to help troubleshoot problems with your virtual machine.

Required privilege: **Virtual machine > Configuration > Settings**

ESXi hosts store virtual machine log files in the same directory as the virtual machine's configuration files. By default, the log file name is `vmware.log`. Archived log files are stored as `vmware-n.log`, where *n* is a number in sequential order beginning with 1.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click the **VM Options** tab and expand **Advanced**.
3. In the Settings row, select **Enable logging** and click **OK**.

You can view and compare log files in the same storage location as the virtual machine configuration files.

Configure Virtual Machine Debugging and Statistics

You can run a virtual machine so that it collects additional debugging information that is helpful to VMware technical support in resolving issues.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click the **VM Options** tab and expand **Advanced**.
3. Select a debugging and statistics option from the drop-down menu.
 - **Run normally**
 - **Record Debugging Information**
 - **Record Statistics**
 - **Record Statistics and Debugging Information**

The number of debugging and statistics options available depends on the host software type and version. On some hosts, some options are not available.

4. Click **OK**.

Change the Swap File Location

When a virtual machine is powered on, the system creates a VMkernel swap file to serve as a backing store for the virtual machine RAM contents.

Power off the virtual machine.

You can accept the default swap file location or save the file to a different location. By default, the swap file is stored in the same location as the virtual machine configuration file.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. Click the **VM Options** tab and expand **Advanced**.
3. Select a swap file location option.

Option	Description
Default	Stores the virtual machine swap file at the default location defined by the host or cluster swap file settings.
Virtual machine directory	Stores the virtual machine swap file in the same folder as the virtual machine configuration file.
Datastore specified by host	If the host or cluster settings define a location for the swap file, this location is used. Otherwise, the swap file is stored with the virtual machine.

4. Click **OK**.

Configure the Virtual Machine Latency Sensitivity

To improve the latency sensitive workloads, you can configure the latency sensitivity of a virtual machine.

Verify that the virtual machine is of hardware version 20 or later.

With vSphere 8.0, you can use the Virtual Hyperthreading (vHT) feature to configure low latency VMs and expose vHT to the guest operating systems. When you enable vHT, each guest vCPU is treated as a single hyperthread of a virtual core.

For more information about the vHT feature, see the *vSphere Resource Management* documentation.

1. Navigate to a virtual machine in the inventory and click **Actions**.
2. Click **Edit Settings**.
3. On the **VM Options** tab, click **Advanced**.
4. From the **Latency Sensitivity** drop-down menu, select a setting.

Option	Description
Normal	The default setting. The CPU scheduler for a VM with normal latency sensitivity can move the virtual CPUs of the VM across any physical CPU of the host, depending on the system load.
High	High Latency Sensitivity requires you to set CPU and 100% memory reservation for the VM. Scheduling of the VM is adjusted for low-latency workload requirements. Each virtual CPU is granted exclusive access to a physical core.
High with Hyperthreading	High Latency Sensitivity requires you to set CPU and 100% memory reservation for the VM. Each virtual CPU is granted exclusive access to a hyperthread of a physical core. Each consecutive pair of virtual machines vCPUs is assigned to the pair of hyperthreads on a physical core.

For more information on how to calculate the full CPU reservation for vHT, see [vHT Full CPU Reservation](#) section in the *vSphere Resource Management* documentation.

5. Click **OK**.

How to Set a Stun Time Limit for Your vGPU Virtual Machines

Learn how to set a stun time limit per virtual machine when you migrate NVIDIA vGPU-powered virtual machines with vSphere vMotion. Setting a stun time limit may prevent vCenter Server from powering on the virtual machine or migrating it to a host and network whose estimated maximum stun time exceeds that limit.

- Verify that the vCenter Server instance is version 8.0 Update 2.
- Verify that the source and destination ESXi hosts are version 8.0 Update 2.
- Verify that vMotion is enabled for vGPU virtual machines. The `vgpu.hotmigrate.enabled` advanced setting must be set to `true`. For more information about how to configure the vCenter Server advanced settings, see *Configure Advanced Settings* in the *vCenter Server Configuration* documentation.

Virtual machine stun time is the time when the virtual machine is inaccessible to users during vMotion. Starting with vSphere 8.0 Update 2, you can set a vMotion stun time limit for vGPU virtual machines.

The limit that you set must be higher than the estimated maximum stun time for the current device configuration. In case of multiple PCI devices, the stun time limit that you set must be higher than the sum of the contributions of all PCI devices. Setting a stun time limit that is lower than the estimated maximum stun time may prevent the virtual machine from powering on.

Note: Maximum stun time estimates are only available for C-Series and Q-Series vGPU types.

The maximum stun time is calculated on the basis of the bandwidth of the host on which the VM runs currently. The calculations may change at the time of migration if the destination host has lower bandwidth. For example, if a virtual

machine is running on a host with a 25 Gbps vMotion NIC, but the destination host has a 10 Gbps vMotion NIC. At the time of migration the maximum stun time calculation is based on the 10 Gbps vMotion NIC.

1. In the vSphere Client, navigate to a vGPU virtual machine.
2. Right-click the virtual machine and select **Edit Settings**.

The **Edit Settings** dialog box appears.

3. Click the **VM Options** tab and expand **Advanced**.

Edit Settings | vgpu-multivgpu

1

Virtual HardwareVM OptionsAdvanced Parameters

> General Options

VM Name: vgpu-multivgpu

> VMware Remote Console Options

Expand for VMware Remote Console settings

> Encryption

Expand for encryption settings

> VMware Tools

Expand for VMware Tools settings

> Boot Options

Expand for boot options

> Power management

Expand for power management settings

< Advanced2

Settings

☒ Enable logging

Debugging and statistics

Run normally

Swap file location

☒ Default

CANCEL

4. In the **vMotion Stun Time Limit** input field, enter the stun time limit in seconds.

CANCEL

5. Click **OK**.

N-port ID virtualization (NPIV) provides the ability to share a single physical Fibre Channel HBA port among multiple virtual ports, each with unique identifiers. This capability lets you control virtual machine access to LUNs on a per-virtual machine basis.

- Each virtual port is identified by a pair of world wide names (WWNs) that vCenter Server assigns. The pair consists of a world wide port name (WWPN) and a world wide node name (WWNN).

NPIV support is subject to the following limitations:

- NPIV must be enabled on the SAN switch. Contact the switch vendor for information about enabling NPIV on their devices.
 - NPIV is supported only for virtual machines with RDM disks. Virtual machines with regular virtual disks continue to use the WWNs of the host's physical HBAs.
 - The physical HBAs on the ESXi host must have access to a LUN through their WWNs in order for the virtual machines on that host to have access to that LUN using their NPIV WWNs. Ensure that access is provided to both the host and the virtual machines.
 - The physical HBAs on the ESXi host must support NPIV. If the physical HBAs do not support NPIV, the virtual machines on that host start using the WWNs of the physical HBAs for LUN access.
 - Each virtual machine can have up to 4 virtual ports. NPIV-enabled virtual machines are assigned exactly 4 NPIV-related WWNs, which are used to communicate with physical HBAs through virtual ports. Therefore, virtual machines can use up to 4 physical HBAs for NPIV purposes.
1. Right-click a virtual machine in the inventory and select **Edit Settings**.
 2. Click the **VM Options** tab and expand **Fibre Channel NPIV**.
 3. Optional: Select the **Temporarily Disable NPIV for this virtual machine** check box.
 4. Select an option for assigning WWNs.
 - To leave WWNs unchanged, select **Leave unchanged**.
 - To have vCenter Server or the ESXi host generate new WWNs, select **Generate New WWNs**.
 - To remove the current WWN assignments, select **Remove WWN assignment**.
 5. Click **OK**.

Configure the CPU Topology of a Virtual Machine

Different virtual CPU topologies impact scheduling and placement from within the guest operating system of a virtual machine.

- Virtual machine requirements:
 - Hardware version 20 or later.
 - Verify that the virtual machine is powered off.

By default, the ESXi host automatically selects a virtual topology that aligns with the underlying hardware. With vSphere 8.0, you can manually configure the virtual CPU topology. For example, you can change the number virtual CPUs per NUMA node, assign a virtual device to a virtual NUMA node or leave it unassigned.

The configuration is preserved when you edit a virtual machine or migrate a virtual machine with vMotion.

Option	Description
	For more information, see the VMware KB article at https://kb.vmware.com/s/article/87651 .
NUMA Nodes	Configure the number of NUMA nodes automatically at power on or manually.
Device Assignment	If you select to configure the NUMA nodes manually, to assign devices to NUMA nodes, click the vertical ellipsis icon () that appears in front of each device and select a NUMA node.

- Click **OK**.

From the **Summary** tab of the virtual machine, you can monitor the CPU topology exposed by the ESXi host - either automatic or manual.

Configure Virtual Machine Advanced File Parameters

You can change or add virtual machine configuration parameters when instructed by a VMware technical support representative, or if you see VMware documentation that instructs you to add or change a parameter to fix a problem with your system.

Important: Changing or adding parameters when a system does not have problems might lead to decreased system performance and instability.

The following conditions apply:

- To change a parameter, you change the existing value for the keyword/value pair. For example, if you start with the keyword/value pair, keyword/value, and change it to keyword/value2, the result is keyword=value2.
- You cannot delete a configuration parameter entry.

CAUTION: You must assign a value to configuration parameter keywords. If you do not assign a value, the keyword can return a value of 0 or false, which can result in a virtual machine that cannot power on.

- Right-click a virtual machine in the inventory and select **Edit Settings**.
- Click the **Advanced Parameters** tab.
- Optional: To add a parameter, enter an attribute name and value for the parameter and click **Add**.
- Optional: To change a parameter, enter a new value in the **Value** text box for that parameter and click  to confirm changes.
- Click **OK**.

Managing Multi-Tiered Applications with vSphere vApp

With vSphere vApp, you can package multiple interoperating virtual machines and software applications into a single unit that you can manage and distribute in OVF format.

A vApp can contain one or more virtual machines. Any operation carried out with the vApp, such as clone or power off, affects all virtual machines in the vApp container.

In the vSphere Client, you can navigate to the vApp **Summary** tab, where you can view the current status of the vApp, and you can manage the vApp.

Note: Because the vApp metadata resides in the vCenter Server database, a vApp can be distributed across multiple ESXi hosts. The metadata information might be lost if the vCenter Server database is cleared or if the standalone ESXi

host that contains the vApp is removed from the vCenter Server. Back up your vApps to an OVF package to avoid losing metadata.

vApp metadata for the virtual machines within the vApp does not follow the snapshot semantics for virtual machine configuration. The vApp properties that you delete, modify, or define after you take a snapshot of a virtual machine remain respectively deleted, modified, or defined, if the virtual machine reverts to that snapshot or any prior snapshots.

Create a vApp

A vApp allows you to perform resource management and certain other management activities such as power operations for multiple virtual machines at the same time.

Verify that one of those objects is available in your data center.

- A standalone host that is running ESX 4.0 or later
- A DRS cluster

You can think of the vApp as the container for the virtual machines, and you can perform the operations on the container.

When you create a vApp, you can add it to a folder, standalone host, resource pool, DRS cluster, or another vApp.

1. In the vSphere Client, right-click an object that supports vApp creation and click **New vApp**.

The **New vApp** wizard starts.

2. On the **Select creation type** page, select **Create a new vApp** and click **Next**.

3. On the **Select a name and location** page, type a name and select a location for the vApp, and click **Next**.

- If you start the creation process from a folder or vApp, you are prompted for a host, cluster, or resource pool.
- If you start the creation process from a resource pool, host, or cluster, you are prompted for a folder or a data center.

4. On the Resource allocation page, allocate CPU and memory resources to the vApp.

Option	Description
Shares	Defines the CPU or memory shares for this vApp with respect to the parent's total. Sibling vApps share resources according to their relative share values bounded by the reservation and limit. Select Low , Normal , or High , which specify share values respectively in a 1:2:4 ratio. Select Custom to give each vApp a specific number of shares that expresses a proportional weight.
Reservation	Defines the guaranteed CPU or memory allocation for this vApp.
Reservation Type	Defines whether the reservation is expandable. Select the Expandable check box to make the reservation expandable. When the vApp is powered on, if the combined reservations of its virtual machines are larger than the reservation of the vApp, the vApp can use resources from its parent or ancestors.
Limit	Defines the upper limit for this vApp's CPU or memory allocation. Select Unlimited to specify no upper limit.

5. On the **Review and finish** page, review the vApp settings and click **Finish**.

Perform vApp Power Operations

One of the advantages of a vApp is that you can perform power operations on all virtual machines it contains at the same time.

Prerequisites depend on the task that you want to perform.

Task	Required Privileges
Power on a vApp	vApp > Power On on the vApp.

Task	Required Privileges
Power off a vApp	vApp > Power Off on the vApp.
Suspend a vApp	vApp > Suspend

1. Navigate to a vApp in the inventory.
2. Select one of the power operation options.

Task	Action
Power on	Right-click the vApp and select Power > Power On. You can power on a vApp to power on all its virtual machines and child vApps. Virtual machines are powered on according to the startup order configuration. If a delay is set in the startup settings of a virtual machine in the vApp, the vApp waits for the set length of time before powering on that virtual machine.
Power off	Right-click the vApp and select Power > Power Off You can power off a vApp to power off all its virtual machines and child vApps. Virtual machines are powered off in reverse startup order. If a delay is set in the shutdown settings of a virtual machine in the vApp, the vApp waits for the set length of time before powering off that virtual machine.
Suspend	Right-click the vApp and select Power > Suspend. You can suspend a vApp to suspend all its virtual machines and child vApps. Virtual machines are suspended in the reverse order of the specified startup order. All virtual machines are suspended regardless of the Suspend behavior you specified in the Power Management VM Option for the virtual machine.
Resume	Right-click a vApp that is powered off or suspended and select Power On. Virtual machines are resumed according to their startup order configuration.

On the **Summary** tab, the **Status** indicates the vApp status.

Create or Add an Object to a vApp

You can populate a vApp with objects by creating a virtual machine, resource pool, or child vApp within the vApp. You can also add an existing object from the inventory, such as a virtual machine or another vApp, to the vApp.

Select your task.

- Create an object inside a vApp.
 - Right-click a vApp in the inventory, and select **New Virtual Machine** to create a new virtual machine within the vApp.
 - Right-click a vApp in the inventory, and select **New Child vApp** to create a child vApp inside the vApp.
 - Right-click a vApp in the inventory, and select **New Resource Pool** to create a new resource pool within the vApp.
 - Right-click a vApp in the inventory, and select **Deploy OVF Template** to deploy an OVF template and add the corresponding virtual machine to the selected vApp.
- Add an existing object to a vApp.
 - a. Navigate to an object in the inventory.
 - b. Drag the object to the destination vApp.
 - c. Release the mouse button.

If the move is not permitted, the object is not added to the vApp.

The new object is now a part of the vApp in the vApp inventory.

Clone a vApp

Cloning a vApp is similar to cloning a virtual machine. When you clone a vApp, you clone all virtual machines and vApps within the vApp.

When you clone a vApp, you can add the clone to a folder, standalone host, resource pool, cluster enabled for DRS, or another vApp.

Verify that one of those objects is available in your datacenter.

- A standalone host that runs ESXi 3.0 or later
 - A DRS cluster
1. Start the cloning wizard.
 - Start the cloning wizard from a DRS cluster by right-clicking the cluster and selecting **New vApp > New vApp** (🧩).
 - Start the cloning wizard from an existing vApp by right-clicking the vApp and selecting **Clone > Clone**.
 2. On the Select creation type page, select **Clone an existing vApp** and click **Next**.
 3. On the Select source vApp page, select an existing vApp to clone, and click **Next**.
The Select source vApp page appears only if you start the wizard from a DRS cluster.
 4. On the Select destination page, select a valid host, vApp, or resource pool in which to run the vApp, and click **Next**.
 5. On the Select a name and location page, type a name for the vApp, select location, and click **Next**.
 6. On the Select storage page, select the virtual disk format and the target datastore and click **Next**.
 7. On the Map networks page, configure the network mappings for the networks that the virtual machines in the vApp use.
 8. On the Resource allocation page, allocate CPU and memory resources for the vApp and click **Next**.
 9. On the Review and finish page, review the vApp settings and click **Finish**.

Edit vApp Notes

You can add or edit notes for a particular vApp.

1. Right-click a vApp in the inventory and select **Edit Notes**.
2. Type your comments in the **Edit Notes** window.
3. Click **OK**.

Your comments appear on the **Summary** tab for the vApp.

Configure vApp Properties

Define and use custom properties to present custom information to all the virtual machines in the vApp. You can later assign values and edit those properties. If you deploy the vApp from an OVF file with properties predefined in that OVF, you can edit those properties as well.

- Power off the vApp.

- Required privilege: **vApp > vApp application configuration** on the vApp.

The **Properties** pane contains a list of all the properties that are defined for a vApp. You can use the filters to browse through the list more easily.

1. Navigate to a vApp in the inventory.
2. On the **Configure** tab, select **Settings > vApp properties**.
The **Properties** pane shows the list of properties and the permissible actions.
3. Select your task by clicking the respective button.

Option	Description
Add	Creates a new property. Important: If the virtual machine is connected to a distributed switch and has the vApp option enabled, you cannot select any of the following Dynamic property options: IP address, Subnet, Netmask, Gateway, Domain name, HTTP Proxy, Host prefix, DNS servers, DNS search path, Network name.
Edit	Edits the property. You can change the general information for the property, such as the property label, category and description. You can also edit the type parameters. Important: If the virtual machine is connected to a distributed switch and has the vApp option enabled, you cannot select any of the following Dynamic property options: IP address, Subnet, Netmask, Gateway, Domain name, HTTP Proxy, Host prefix, DNS servers, DNS search path, Network name.
Set Value	Sets a value for the property. That value is different from the default value that you define when you create a new property.
Delete	Removes a property from the list.

Edit vApp Settings

You can edit and configure several vApp settings, including startup order, resources, and custom properties.

1. Right-click a vApp in the inventory and click **Edit Settings**.
2. On the **Resources** tab, edit CPU and memory settings, such as shares, reservation, and limit.
3. On the **Start Order** tab, set and edit the start order of the virtual machines.
4. On the **IP Allocation** tab, specify IP protocol and choose an IP allocation scheme.
5. On the **Details** tab, view product information, such as name, vendor, product URL, and vendor URL.

Configure vApp CPU and Memory Resources

You can configure the CPU and memory resource allocation for the vApp.

Reservations on vApps and all their child resource pools, child vApps, and child virtual machines count against the parent resources only if those objects are powered on.

1. Right-click a vApp in the inventory and click **Edit Settings**.
2. To allocate CPU resources to the vApp, click the **Resources** tab and expand **CPU**.

Option	Description
Shares	CPU shares for this vApp with respect to the parent's total. Sibling vApps share resources according to their relative share values bounded by the reservation and limit. Select Low , Normal , or High , which specify share values respectively in a 1:2:4 ratio. Select Custom to give each vApp a specific number of shares, which express a proportional weight.
Reservation	Guaranteed CPU allocation for this vApp.

Option	Description
Reservation Type	To make the reservation expandable, select the Expandable check box. When the vApp is powered on, if the combined reservations of its virtual machines are larger than the reservation of the vApp, the vApp can use resources from its parent or ancestors.
Limit	Upper limit for this vApp's CPU allocation. Select Unlimited to specify no upper limit.

3. To allocate memory resources to the vApp, click the **Resources** tab and expand **Memory**.

Option	Description
Shares	Memory shares for this vApp with respect to the parent's total. Sibling vApps share resources according to their relative share values bounded by the reservation and limit. Select Low , Normal , or High , which specify share values respectively in a 1:2:4 ratio. Select Custom to give each vApp a specific number of shares, which express a proportional weight.
Reservation	Guaranteed memory allocation for this vApp.
Reservation Type	To make the reservation expandable, select the Expandable check box to make the reservation expandable. When the vApp is powered on, if the combined reservations of its virtual machines are larger than the reservation of the vApp, the vApp can use resources from its parent or ancestors.
Limit	Upper limit for this vApp's memory allocation. Select Unlimited to specify no upper limit.

4. Click **OK**.

Configure vApp IP Allocation Policy

If your vApp is set up to allow it, and if you have the required privileges, you can edit how IP addresses are allocated for the vApp.

Required privilege: **vApp > vApp instance configuration**

You cannot configure the IP allocation policy during the vApp creation process.

Before you configure the IP allocation policy, you must specify the IP protocol and the IP allocation scheme that the vApp supports.

If you deploy a vApp from an OVF template, the IP allocation policy might still be editable.

- Right-click a vApp in the inventory and click **Edit Settings**.
- In the **Edit vApp** dialogue box, click the **IP Allocation** tab.
This tab is only available in the vSphere Client.
- In the Authoring section, define the IP protocol and the IP allocation scheme that the vApp supports.

The IP protocol and the IP allocation scheme determine what options for IP allocation are available.

A vApp can obtain its network configuration through the OVF environment or through a DHCP server. If you do not select any of these options, the IP addresses are manually allocated.

The IP protocols that a vApp can support are IPv4, IPv6, or both.

- In the Deployment section, select an IP allocation policy from the **IP allocation** drop-down menu.

Option	Description
Static - Manual	IP addresses are manually configured. No automatic allocation is performed.
Transient - IP Pool	IP addresses are automatically allocated using IP pools from a specified range when the vApp is powered on. The IP addresses are released when the appliance is powered off.

Option	Description
DHCP	A DHCP server is used to allocate the IP addresses. The addresses assigned by the DHCP server are visible in the OVF environments of virtual machines started in the vApp.
Static - IP Pool	IP addresses are automatically allocated from the managed IP network range of vCenter Server at power-on, and remain allocated at power-off.

Static - IP Pool and Transient - IP Pool have in common that IP allocation is done through the range managed by the vSphere platform as specified by the IP pool range in a network protocol profile. The difference is that for a static IP Pool, the IP addresses are allocated at first power-on and remain allocated, while for a transient IP Pool, the IP addresses are allocated when needed, typically at power-on, but released during power-off.

- Click **OK**.

Configure vApp Startup and Shutdown Options

You can change the order in which virtual machines and nested vApps within a vApp start up and shut down. You can also specify the delays and actions performed at startup and shutdown.

- Right-click a vApp in the inventory and click **Edit Settings**.
- Select a virtual machine and select its order group.

Virtual machines and vApps in the same group are started before the objects in the next group. The first group of virtual machines to power on is Group 1, followed by Group 2, Group 3, and so forth. The reverse order is used for shutdown.

- Click the **Start Order** tab and select a virtual machine from the list.
- From the **Group** drop-down menu, select a group for the virtual machine.
- Optional: Select the startup action for the virtual machine.
The default is **Power On**. Select **None** to power on the virtual machine manually.
- Optional: Specify when the startup action is to happen.
 - Enter a time delay in seconds for the startup action.
 - To perform the startup action when VMware Tools has started, select **Continue when VMware Tools are ready**.
- Optional: Select the shutdown action for the virtual machine.
The default shutdown action is **Power Off**. You can also select **Guest Shutdown** to shut down the guest OS and leave the virtual machine running, **Suspend**, or **None**.
- Optional: Enter a time delay in seconds for the shutdown action.
- Click **OK**.

Configure vApp Product Properties

You can configure product and vendor information for a vApp.

Required privilege: **vApp > vApp application configuration** on the vApp.

- Right-click a vApp in the inventory and click **Edit Settings**.
- To enter product and vendor information, click the **Details** tab.

vApp Setting	Description
Name	Product name.
Product URL	If you enter a product URL, a user can click the product name on the virtual machine summary page and go to the product's web page.

vApp Setting	Description
Vendor	Vendor name.
Vendor URL	If you enter a vendor URL, a user can click the vendor name from the Summary page of the virtual machine and go to the vendor's web page.

3. Click **OK**.

View vApp License Agreements

You can view the license agreement for the vApp that you edit.

1. Navigate to a vApp in the inventory.
2. On the **Configure** tab, expand **Settings** and click **License agreements**.

Add a Network Protocol Profile

A network protocol profile contains a pool of IPv4 and IPv6 addresses that vCenter Server assigns to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

You can configure network protocol profile ranges for IPv4, IPv6, or both. vCenter Server uses these ranges to dynamically allocate IP addresses to the virtual machines within a vApp, when the vApp uses transient IP allocation policy.

Network protocol profiles also contain settings for the IP subnet, the DNS, and HTTP proxy servers.

Note: If you move a vApp or a virtual machine that retrieves its network settings from a protocol profile to another data center, to power on the vApp or virtual machine you must assign a protocol profile to the connected port group on the destination data center.

1. Navigate to a data center that is associated with a vApp.
2. On the **Configure** tab, select **More > Network Protocol Profiles**.
Existing network protocol profiles are listed.
3. Click the **Add** button.
The **Add Network Protocol Profile** wizard opens.
4. On the **Name and network** page, enter the name of the network protocol profile and select the networks that use this profile. Click **Next**.
A network can be associated with one network protocol profile at a time.

5. On the **IPv4** page, configure the relevant IPv4 settings.
 - a) In the **Subnet** and the **Gateway** text boxes, enter the IP subnet and gateway.
 - b) To indicate that the DHCP server is available on the network, select the **DHCP Present** radio button.
 - c) In the **DNS server addresses** text box, enter the DNS server information.
 - d) To specify an IP pool range, enable the **IP Pool** option.
 - e) If you enable IP pools, enter a comma-separated list of host address ranges in the **IP pool range** text box.
 A range consists of an IP address, a pound sign (#), and a number indicating the length of the range.
 For example, `10.20.60.4#10, 10.20.61.0#2` indicates that the IPv4 addresses can range from 10.20.60.4 to 10.20.60.13 and 10.20.61.0 to 10.20.61.1.
 The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** text box cannot include the gateway address.
 - f) Click **Next**.
6. On the **IPv6** page, configure the relevant IPv6 settings.
 - a) In the **Subnet** and the **Gateway** text boxes, enter the IP subnet and gateway.
 - b) Select the **DHCP Present** radio button to indicate that the DHCP server is available on this network.
 - c) In the **DNS server addresses**, enter the DNS server information.
 - d) Enable the **IP Pool** option to specify an IP pool range.
 - e) If you enable IP pools, enter a comma-separated list of host address ranges in the **IP pool range** text box.
 A range consists of an IP address, a pound sign (#), and a number indicating the length of the range.
 For example, assume that you specify the following IP pool range: `fe80:0:0:0:2bff:fe59:5a:2b#10, fe80:0:0:0:2bff:fe59:5f:b1#2`. Then the addresses are in this range:
`fe80:0:0:0:2bff:fe59:5a:2b - fe80:0:0:0:2bff:fe59:5a:34`
 and
`fe80:0:0:0:2bff:fe59:5f:b1 - fe80:0:0:0:2bff:fe59:5f:b2`.
 The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** text box cannot include the gateway address.
 - f) Click **Next**.
7. On the **Other network configurations** page, specify additional network configurations.
 - a) Enter the DNS domain.
 - b) Enter the host prefix.
 - c) Enter the DNS search path.
 The search paths are specified as a list of DNS domains separated by commas, semi-colons, or spaces.
 - d) Enter the server name and port number for the proxy server.
 The server name must include a colon and a port number. For example, `web-proxy:3912` is a valid proxy server.
 - e) Click **Next**.
8. On the **Name and Network Assignment** page, review the settings and click **Finish**.

Assign a Port Group or Network to a Network Protocol Profile

To apply the range of IP addresses from a network protocol profile to a virtual machine that is a part of a vApp or has vApp functionality enabled, assign the network or distributed port group that controls the networking of the virtual machine to the network protocol profile.

1. Navigate to a data center that is associated with a vApp.
2. On the **Configure** tab, select **More > Network Protocol Profiles**.
Existing network protocol profiles are listed.
3. Select a network protocol profile from the list and click the **Assign** button.
The **Assign Networks** dialog box opens.
4. Select a port group or a network to assign to the network protocol profile.
 - On the **Distributed portgroups** tab, you see a list of the distributed port groups.
 - On the **Networks** tab, you see a list of the port groups of standard switches.
 You can select multiply port groups before you close the dialog box.
5. Click **Save**.

The port groups that you selected are now associated with the network protocol profile.

To apply to a virtual machine or vApp the range of IP addresses that the network protocol profile contains, configure the virtual machine or vApp to use the network protocol profile. For more information, see [Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine or vApp](#).

Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine or vApp

After you associate a network protocol profile with a port group of a standard switch or a distributed switch, you can use the profile to dynamically allocate IP addresses to a virtual machine that is within a vApp.

Verify that the virtual machine is connected to a port group that is associated with the network protocol profile.

Select your task.

Option	Description
Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine	1. Navigate to a virtual machine in the vCenter Server inventory. 2. On the Configure tab, expand Settings and select vApp Options. 3. Click the Edit button. The Edit vApp options dialog box opens. 4. If vApp options are not enabled, select the Enable vApp options check box. 5. Click the IP Allocation tab. 6. In the Authoring section, select OVF environment as an IP allocation scheme. 7. In the Deployment section, set the IP allocation to Transient - IP Pool or Static - IP Pool. 8. Click OK.
Use a Network Protocol Profile to Allocate IP Addresses to a vApp	1. Navigate to a vApp in the vCenter Server inventory 2. Right-click the vApp and select Edit Settings. The Edit vApp dialog box opens. 3. Click the IP Allocation tab. 4. In the Authoring section, select OVF environment as an IP allocation scheme. 5. In the Deployment section, set the IP allocation to Transient - IP Pool or Static - IP Pool. 6. Click OK.

Both the **Static - IP Pool** and **Transient - IP Pool** options allocate an IP address from the range defined in the network protocol profile that is associated with the port group. If you select **Static - IP Pool**, the IP address is assigned the first time the virtual machine or vApp is powered on. The assigned IP address persists across restarts. If you select **Transient - IP Pool**, an IP address is assigned every the virtual machine or vApp is powered on.

When the virtual machine is powered on, the adapters connected to the port group receive IP addresses from the range in the protocol profile. When the virtual machine is powered off, the IP addresses are released.

Virtual Machine vApp Options

You can enable and configure vApp functionality for virtual machines that are not within a vApp. When a virtual machine vApp options are enabled, you can configure OVF properties, use the OVF environment, and specify IP allocation and product information for the virtual machine.

You can configure and modify the authoring vApp options of a virtual machine to specify custom information that is preserved and used when you export the virtual machine as an OVF template. If you later deploy that OVF template, the information that you specified is available for editing in the deployment vApp options of the virtual machine.

For a virtual machine with vApp options enabled, the authoring vApp options are preserved when you export the virtual machine as an OVF template and those options are used when you deploy a new virtual machine from that template. Deployment vApp options are available for virtual machines that are deployed from OVF templates.

Activate vApp Options for a Virtual Machine

To activate the vApp functionality for a virtual machine that is not part of a vApp, you must enable vApp options at the virtual machine level. Virtual machine vApp options are saved when you export the virtual machine as an OVF template. Those vApp options are later used when you deploy the OVF template.

If you activate vApp options and export a virtual machine to OVF, the virtual machine receives an OVF Environment XML descriptor at boot time. The OVF descriptor might include values for custom properties, including network configuration and IP addresses.

The OVF environment can be transported to the guest in two ways:

- As a CD-ROM that contains the XML document. The CD-ROM is mounted on the CD-ROM drive.
- Through VMware Tools. The guest OS environment variable `guestinfo.ovfEnv` contains the XML document.

For information on how to create a vApp, see [Create a vApp](#).

1. Navigate to a virtual machine from the inventory.
2. On the **Configure** tab, expand **Settings** and select **vApp options**.
3. Click the **Edit** button.

The **Edit vApp Options** dialog box opens.

4. If vApp options are deactivated, select the **Enable vApp options** check box and click **OK**.

Edit Application Properties and OVF Deployment Options for a Virtual Machine

If a virtual machine is a deployed OVF template, you can view the application properties and the OVF deployment options that are defined in the OVF. The deployment options include information about unrecognized OVF sections and the IP allocation policy.

Verify that the virtual machine has the vApp options enabled. See [Activate vApp Options for a Virtual Machine](#).

1. Navigate to a virtual machine from the inventory.
2. On the **Configure** tab, expand **Settings** and select **vApp options**.
3. Click the **Edit** button.

The **Edit vApp Options** dialog box opens.

4. If the OVF template included editable IP allocation options, click the **IP Allocation** tab and modify the IP allocation options in the **Deployment** section.

Option	Description
Static - Manual	IP addresses are manually configured. No automatic allocation is performed.
Transient - IP Pool	IP addresses are automatically allocated using IP pools from a specified range when the vApp is powered on. The IP addresses are released when the appliance is powered off.
DHCP	A DHCP server is used to allocate the IP addresses. The addresses assigned by the DHCP server are visible in the OVF environments of virtual machines started in the vApp.
Static - IP Pool	IP addresses are automatically allocated from the managed IP network range of vCenter Server at power-on, and remain allocated at power-off.

With the **Static - IP Pool** and **Transient - IP Pool** options, the allocation of IP addresses is done through an IP pool range defined in a network protocol profile. The difference between the two options is that with a static IP pool, the IP addresses are allocated at the first power on of the virtual machine and remain allocated, while with a transient IP pool, the IP addresses are allocated when needed, typically at power-on, and released during power-off.

OVF Authoring Options for a Virtual Machine

You can use the OVF authoring options that are included in a virtual machine vApp options to specify custom information that is preserved when you export the virtual machine as an OVF template.

vApp properties are a central concept to vApp deployment and self-configuration. They can turn a general OVF package into a running vApp instance with a custom configuration.

The set of properties assigned to a running vApp is determined by the OVF package from which the vApp is deployed.

- When an OVF package is created the author adds the set of properties necessary for the vApp to function in an unknown environment. For example, properties that contain network configuration, a property that contains the email address of the system administrator, or a property that contains the number of expected vApp users.
- Some property values are entered by the user when the vApp is deployed, while other property values are configured by vCenter Server when the vApp is powered on. How properties are handled depends on the property type and the vCenter Server configuration.

When vCenter Server powers on a vApp, it creates an XML document that contains all properties and their values. This document is made available to each virtual machine in the vApp, and allows the virtual machines to apply the properties to their own environment.

Edit vApp Product Information for a Virtual Machine

If you want to export a virtual machine as an OVF template, you can specify product information that becomes available when you deploy a new virtual machine from the OVF template.

1. Navigate to a virtual machine from the **inventory**.
2. On the **Configure** tab, expand **Settings** and select **vApp options**.
3. Click the **Edit** button.
The **Edit vApp Options** dialog box opens.
4. To enter product and vendor information, click the **Details** tab.

vApp Setting	Description
Name	Product name.
Product URL	If you enter a product URL, a user can click the product name on the virtual machine summary page and go to the product's web page.
Vendor	Vendor name.
Vendor URL	If you enter a vendor URL, a user can click the vendor name from the Summary page of the virtual machine and go to the vendor's web page.

5. Click **OK**.

Manage vApp Custom Properties for a Virtual Machine

You can define and manage custom properties that are stored in the OVF template when you export a virtual machine or vApp and are used by vCenter Server when you deploy the OVF template. OVF templates support static properties, which are often configured by the user, and dynamic properties, which are always set by the vCenter Server.

To customize your virtual machine or vApp with properties, perform the following steps.

1. Define the OVF properties, for example a DNS address or gateway, in the virtual machine or vApp.
2. If you plan to export the virtual machine or the vApp to an OVF template:
 - a. Set up the OVF environment transport to carry the settings into the virtual machine. See [Edit OVF Details for a Virtual Machine](#).
 - b. Write some glue-code to access and apply the information to the virtual machine.

See the VMware vApp Developer blog topic *Self-Configuration and the OVF Environment* for a discussion, sample code, and a video.

1. Navigate to a virtual machine from the **inventory**.
2. On the **Configure** tab, expand **Settings** and select **vApp options**.
3. You can manage custom properties from the **Properties** panel.

Option	Description
Add	To create a property, click Add .
Edit	To edit an existing property, select the property and click Edit .
Set Value	To set a new value to the existing property, select the property and click Set Value .
Delete	To delete an existing property, select the property and click Delete .

Edit vApp IP Allocation Policy for a Virtual Machine

You can set or edit the IP allocation policy that the virtual machine uses when you export it to an OVF template and deploy the OVF template.

1. Navigate to a virtual machine from the inventory.
2. On the **Configure** tab, expand **Settings** and select **vApp options**.

3. Click the **Edit** button.

The **Edit vApp Options** dialog box opens.

4. If vApp options are not enabled, select the **Enable vApp options** check box.
5. Select the **IP protocol** and an **IP allocation scheme**.

The supported protocols are IPv4, IPv6, or both.

Consult the following table to learn more about the IP allocation schemes.

Option	Description
OVF environment	IP allocation is determined by the environment in which you deploy the OVF template.
DHCP	The IP addresses are allocated through a DHCP server when the virtual machine is powered on.

The information you specify in the Authoring section will be used if you export the virtual machine to OVF and deploy the OVF at a later time.

6. Click **OK**.

Edit OVF Details for a Virtual Machine

A virtual machine's OVF settings allow you to customize the OVF environment, OVF transport, and boot behavior after OVF deployment. You can edit and configure settings that affect the OVF environment in the Virtual Machine Properties dialog box.

vApp options must be enabled in order to access these options.

1. Navigate to a virtual machine from the inventory.
2. On the **Configure** tab, expand **Settings** and select **vApp options**.

3. Click the **Edit** button.

The **Edit vApp Options** dialog box opens.

4. If vApp options are not enabled, select the **Enable vApp options** check box.
5. To customize the OVF settings for the virtual machine, click the **OVF Details** tab.

Option	Description
OVF environment transport	• If you select ISO image, an ISO image that contains the OVF template information is mounted in the CD-ROM drive. • If you select VMware Tools, the VMware Tools <code>guestInfo.ovfEnv</code> variable is initialized with the OVF environment document.
Installation boot	If you click Enable , the virtual machine reboots after the OVF deployment completes. You can specify the delay time in seconds before the virtual machine starts the reboot operation.

6. Click **OK** to save your changes.

In the **vApp Options are enabled** pane, you can view the **OVF Settings** panel with information about the OVF settings.

7. Optional: To view information about the OVF environment settings, power on the virtual machine and click the **View OVF Environment** button in the **vApp Options are enabled** pane.

The information is displayed in XML format.

Monitoring Solutions with the vCenter Solutions Manager

A solution is an extension of the vCenter Server that adds new functions to the vCenter Server instance. You can monitor the detailed information of an inventory and the health status of the solutions.

VMware products that integrate with vCenter Server are also considered solutions. For example, vSphere ESX Agent Manager is a VMware solution that lets you manage host agents that add new capabilities to ESX and ESXi hosts.

You can also install a solution to add functionality from third-party technologies to the standard functions of vCenter Server. Solutions are typically delivered as OVF packages. You can install and deploy solutions from the vSphere Client. You can integrate the solutions into the vCenter Solutions Manager, which provides a list of all installed solutions.

If a virtual machine or vApp runs a solution, a custom icon represents it in the inventory of the vSphere Client. Each solution registers a unique icon to show that the solution manages the virtual machine or vApp. The icons show the power states (powered on, paused, or powered off). Solutions display more than one type of icon if they manage more than one type of virtual machine or vApp.

When you power on or power off a virtual machine or vApp, you are notified that you are operating with an object that is managed by the vCenter Solutions Manager. When you attempt another operation on a virtual machine or a vApp that is managed by a solution, an informational warning message appears.

For more information, see the *Developing and Deploying vSphere Solutions, vServices, and ESX Agents* documentation.

View Solutions

vCenter Solutions Manager helps to deploy, monitor, and interact with solutions that are installed in a vCenter Server instance.

The vCenter Solutions Manager displays information about the solution, such as the solution name, the vendor name, and the version of the product. The vCenter Solutions manager also displays information about the health of a solution.

1. Navigate to the vCenter Solutions Manager.
 - a) In the vSphere Client, select **Menu > Administration**.
 - b) Expand **Solutions** and click **vCenter Server Extensions**.

2. Click a solution from the list.

For instance, vService Manager or vSphere ESX Agent Manager.

3. Navigate through the tabs to view information about the solution.

- **Summary:** You can view details about the solution, such as the product name, a short description, and links to the product and vendor websites. You can also view the solution configuration and the solution UI. Select the vCenter Server link to view the **Summary** page of the Virtual Machines or vApp.
- **Monitor:** You can view the tasks and events related to the solution.
- **VMs:** You can view a list of all the Virtual Machines and vApps that belong to the solution.

Managing Virtual Machines

You can manage individual virtual machines or a group of virtual machines that belongs to a host or cluster.

From the virtual machine's console, you can change the guest operating system settings, use applications, browse the file system, monitor system performance, and so on. Use snapshots to capture the state of the virtual machine at the time you take the snapshot.

To migrate virtual machines using cold or hot migration, including vMotion, vMotion in environments without shared storage, and Storage vMotion, see the *vCenter Server and Host Management* documentation.

Install a Guest Operating System

A virtual machine is not complete until you install the guest operating system and VMware Tools. Installing a guest operating system in your virtual machine is essentially the same as installing it in a physical computer.

The basic steps for a typical operating system are described in this section. See the *Guest Operating System Installation Guide* at <http://partnerweb.vmware.com/GOSIG/home.html>.

Using PXE with Virtual Machines

You can start a virtual machine from a network device and remotely install a guest operating system using a Preboot Execution Environment (PXE). You do not need the operating system installation media. When you turn on the virtual machine, the virtual machine detects the PXE server.

PXE booting is supported for Guest Operating Systems that are listed in the VMware Guest Operating System Compatibility list and whose operating system vendor supports PXE booting of the operating system.

The virtual machine must meet the following requirements:

- Have a virtual disk without operating system software and with enough free disk space to store the intended system software.
- Have a network adapter connected to the network where the PXE server resides.

For details about guest operating system installation, see the *Guest Operating System Installation Guide* at <http://partnerweb.vmware.com/GOSIG/home.html>.

Install a Guest Operating System from Media

You can install a guest operating system from a CD-ROM or from an ISO image. Installing from an ISO image is typically faster and more convenient than a CD-ROM installation.

- Verify that the installation ISO image is present on a VMFS datastore or network file system (NFS) volume accessible to the ESXi host.
Alternatively, verify that an ISO image is present in a content library.
- Verify that you have the installation instructions that the operating system vendor provides.

1. Log in to the vCenter Server system or host on which the virtual machine resides.
2. Select an installation method.

Option	Action
CD-ROM	Insert the installation CD-ROM for your guest operating system into the CD-ROM drive of your ESXi host.
ISO image	1. Right-click the virtual machine and select Edit Settings. The virtual machine Edit Settings dialog box opens. If the Virtual Hardware tab is not preselected, select it. 2. Select Datastore ISO File from the CD/DVD drop-down menu, and browse for the ISO image for your guest operating system.
ISO image from a Content Library	1. Right-click the virtual machine and select Edit Settings. The virtual machine Edit Settings dialog box opens. If the Virtual Hardware tab is not preselected, select it.

Option	Action
	2. Select Content Library ISO File from the CD/DVD drop-down menu, and select an ISO image from the content library items.

3. Right-click the virtual machine and select **Power On**.

A green right arrow appears next to the virtual machine icon in the inventory list.

4. Follow the installation instructions that the operating system vendor provides.

Install VMware Tools. VMware highly recommends running the latest version of VMware Tools on your guest operating systems. Although the guest operating system can run without VMware Tools, you lose important functionality and convenience without them.

Upload ISO Image Installation Media for a Guest Operating System

You can upload an ISO image file to a datastore from your local computer. You can do this when a virtual machine, host, or cluster does not have access to a datastore or to a shared datastore that has the guest operating system installation media that you require.

Required privileges:

- **Datastore > Browse datastore** on the datastore.
- **Datastore > Low level file operations** on the datastore.
- **Host > Configuration > System Management**

1. In the vSphere Client, select **Menu > Storage**.
2. Select the datastore from the inventory to which you will upload the file.
3. Optional: On the **Files** tab, click the **New Folder** icon to create a new folder.
4. Select an existing folder or the folder that you created, and click the **Upload Files** icon.
5. On the local computer, find the file and upload it.
ISO upload times vary, depending on file size and network upload speed.
6. Refresh the datastore file browser to see the uploaded file in the list.

After you upload the ISO image installation media, you can configure the virtual machine CD-ROM drive to access the file.

Customizing Guest Operating Systems

When you clone a virtual machine or deploy a virtual machine from a template, you can customize the guest operating system of the virtual machine. You can change the computer name, network settings, and license settings.

Customizing guest operating systems helps prevent conflicts that occur if virtual machines with identical settings are deployed, for example conflicts due to duplicate computer names. You can apply customization as part of virtual machine deployment or later.

- During the cloning or deployment process, you can specify customization settings or you can select an existing customization specification.
- You can create a customization specification explicitly from the **Policies and Profiles** and apply it to a virtual machine.

Guest Operating System Customization Requirements

To customize the guest operating system, you must configure the virtual machine and guest operating system to meet VMware Tools and virtual disk requirements.

Other requirements apply, depending on the guest operating system type.

VMware Tools Requirements

The latest version of VMware Tools must be installed on the virtual machine or template to customize the guest operating system during cloning or deployment. For information about VMware Tools support matrix, see the *VMware Product Interoperability Matrixes* at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php.

Virtual Disk Requirements

The guest operating system being customized must be installed on a disk attached as SCSI node 0:0 in the virtual machine configuration.

Windows Requirements

Customization of Windows guest operating systems requires the virtual machine to be running on an ESXi host running version 3.5 or later.

Linux Requirements

Customization of Linux guest operating systems requires that Perl is installed in the Linux guest operating system.

Note: Certain Linux OS versions do not include the D-Bus package in the minimal Linux installation option. Therefore, running commands such as `dbus-uuidgen` might fail during customization.

Verifying Customization Support for a Guest Operating System

To verify customization support for Windows operating systems or Linux distributions and compatible ESXi hosts, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>. You can use this online tool to search for the guest operating system and ESXi version. After the tool generates your list, click the guest operating system to see whether guest customization is supported.

Create a vCenter Server Application to Generate Computer Names and IP Addresses

Instead of entering computer names and IP addresses for virtual NICs when you customize guest operating systems, you can create a custom application and configure it in such a way that vCenter Server generates the names and addresses.

Verify that Perl is installed on vCenter Server.

The application can be an arbitrary executable binary or script file appropriate for the corresponding operating system in which vCenter Server is running. After you configure an application and make it available to vCenter Server, every time you initiate a guest operating system customization for a virtual machine, vCenter Server runs the application.

The application must comply with the reference XML file in the VMware knowledge base article at <http://kb.vmware.com/kb/2007557>.

1. Create the application and save it on the vCenter Server system's local disk.
2. Select a vCenter Server instance in the inventory.
3. Click the **Configure** tab, click **Settings**, and click **Advanced Settings**.
4. Click **Edit Settings** and enter the configuration parameters for the script.
 - a) In the **Name** text box, enter `config.guestcust.name-ip-generator.arg1`.
 - b) In the **Value** text box, enter `c:\sample-generate-name-ip.pl` and click **Add**.
 - c) In the **Name** text box, enter `config.guestcust.name-ip-generator.arg2`.
 - d) In the **Value** text box, enter the path to the script file on the vCenter Server system and click **Add**. For example, enter `c:\sample-generate-name-ip.pl`.
 - e) In the **Name** text box, enter `config.guestcust.name-ip-generator.program`.
 - f) In the **Value** text box, enter `c:\perl\bin\perl.exe` and click **Add**.
5. Click **Save**.

You can select the option to use an application to generate computer names or IP addresses during guest operating system customization.

Customize Windows During Cloning or Deployment

You can customize the Windows guest operating systems of the virtual machines when you deploy a new virtual machine from a template or clone an existing virtual machine. Customizing the guest operating system helps prevent conflicts that might result if you deploy virtual machines with identical settings, such as duplicate computer names.

- Verify that all requirements for customization are met. See [Guest Operating System Customization Requirements](#).
- Verify that there are customization specifications available for use. For information about creating a guest customization specification, see [Create a Customization Specification for Windows](#).

You can prevent Windows from assigning new virtual machines or templates with the same Security IDs (SIDs) as the original virtual machine. Duplicate SIDs do not cause problems when the computers are part of a domain and only domain user accounts are used. However, if the computers are part of a Workgroup or local user accounts are used, duplicate SIDs can compromise file access controls. For more information, see the documentation for your Microsoft Windows operating system.

Important: After customization, the default administrator password is not preserved for Windows Server 2008. During customization, the Windows Sysprep utility deletes and recreates the administrator account on Windows Server 2008. You must reset the administrator password when the virtual machine starts the first time after customization.

1. Right-click any vSphere inventory object that is a valid parent object of a virtual machine, such as data center, cluster, vApp, resource pool, or host, and select **New Virtual Machine**.
2. On the **Select a creation page**, select **Clone an existing virtual machine** or **Deploy from template**.
3. Click **Next**.
4. Follow the prompts until you reach the **Select clone options** page.
5. On the **Select clone options** page, select the **Customize the operating system** check box and click **Next**.
6. On the **Customize guest OS** page, apply a customization specification to the virtual machine, and click **Next**.

Option	Action
Select an existing specification	Select a customization specification from the list.

Option	Action
Override	To change the guest customization specification for this deployment only, click Override , complete the steps in the Override VM Customization Specification wizard, and click OK .

- On the **User settings** page, specify the required settings for the virtual machine.

This page of the wizard appears only if the selected specification requires additional customization.

- On the **Ready to complete** page, review the details and click **Finish**.

When the new virtual machine starts for the first time, the guest operating system runs finalization scripts to complete the customization process. The virtual machine might restart several times during this process.

If the guest operating system pauses when the new virtual machine starts, it might be waiting for you to correct errors, such as an incorrect product key or an invalid user name. To determine whether the system is waiting for information, open the virtual machine console.

After you deploy certain Windows operating systems that are not volume licensed, you might need to reactivate your operating system on the new virtual machine.

If the new virtual machine encounters customization errors while it is starting, the errors are logged to %WINDIR%\temp\vmware-imc. To view the error log file, from the Windows **Start** menu navigate to **Programs > Administrative Tools > Event Viewer**.

Customize Linux During Cloning or Deployment

When you deploy a virtual machine from a template or clone an existing virtual machine, you can customize the Linux guest operating systems of the virtual machine.

- Verify that all requirements for customization are met. See [Guest Operating System Customization Requirements](#).
 - Verify that there are customization specifications available for use. For information about creating a guest customization specification, see [Create a Customization Specification for Linux](#).
- Right-click any vSphere Client inventory object that is a valid parent object of a virtual machine, such as data center, cluster, vApp, resource pool, or host, and select **New Virtual Machine**.
 - On the **Select a creation page**, select **Clone an existing virtual machine** or **Deploy from template**.
 - Click **Next**.
 - Follow the prompts until you reach the **Select clone options** page.
 - On the **Select clone options** page, select the **Customize the operating system** check box and click **Next**.
 - On the **Customize guest OS** page, apply a customization specification to the virtual machine, and click **Next**.

Option	Action
Select an existing specification	Select a customization specification from the list.
Override	To change the guest customization specification for this deployment only, click Override , complete the steps in the Override VM Customization Specification wizard, and click OK .

- On the **User settings** page, specify the required settings for the virtual machine.

This page of the wizard appears only if the selected specification requires additional customization.

- On the **Ready to complete** page, review the details and click **Finish**.

When the new virtual machine starts for the first time, the guest operating system runs finalization scripts to complete the customization process. The virtual machine might restart several times during this process.

If the guest operating system pauses when the new virtual machine starts, it might be waiting for you to correct errors, such as an incorrect product key or an invalid user name. To determine whether the system is waiting for information, open the virtual machine console.

If the new virtual machine encounters customization errors while it is starting, the errors are reported using the guest's system logging mechanism. View the errors in the `/var/log/vmware-imc/toolsDeployPkg.log` file.

Apply a Customization Specification to an Existing Virtual Machine

To prevent conflicts that can occur if you deploy virtual machines with identical settings, such as duplicate computer names, you can apply a customization specification to an existing virtual machine.

- Verify the guest operating system is installed.
- Verify that VMware Tools is installed and running.
- Power off the virtual machine.

When you clone an existing virtual machine, or deploy a virtual machine from a VM template in a folder, you can customize the guest operating system of the resulting virtual machine during the clone or the deployment tasks.

When you deploy a virtual machine from a template in a content library, you can customize the guest operating system only after the deployment task is complete.

1. Right-click a virtual machine in the vSphere inventory, and select **Guest OS > Customize Guest OS**.

The **Customize Guest OS** dialog box opens.

2. Select a customization specification from the list and click **OK**.

If the specification requires you to configure additional settings, a new dialog box opens and you are prompted to enter information about the required settings.

Create and Manage Customization Specifications for Windows and Linux

You can create and manage customization specifications for Windows and Linux guest operating systems. Customization specifications are XML files that contain guest operating system settings for virtual machines.

When you apply a specification to the guest operating system during virtual machine cloning or deployment, you prevent conflicts that might occur if you deploy virtual machines with identical settings, such as duplicate computer names.

vCenter Server saves the customized configuration parameters in the vCenter Server database. If the customization settings are saved, the administrator and domain administrator passwords are stored in encrypted format in the database. Because the certificate used to encrypt the passwords is unique to each vCenter Server system, if you reinstall vCenter Server or attach a new instance of the server to the database, the encrypted passwords become invalid. You must reenter the passwords before you can use them.

To learn how you can create and manage customization specifications in the vSphere Client, watch the following video.

Create a Customization Specification for Linux

Save the system settings for a Linux guest operating system in a customization specification, which you can apply when cloning virtual machines or deploying virtual machines from templates. You can use the cloud-init utility as a standardized way to initialize a virtual machine when you deploy it in a cloud environment.

- Verify that all requirements for customization are met. See [Guest Operating System Customization Requirements](#).

- To run the customization script:
 - Verify that VMware Tools version 10.1.0 or later is installed. The customization fails if VMware Tools is version earlier than 10.1.0 and you attempt to run the customization script.
 - In the VMware Tools configuration, the `enable-custom-scripts` option is deactivated by default for security reasons. When you attempt to run the customization script with a deactivated `enable-custom-scripts` option, the customization fails with a customization error.

For example, to enable the `enable-custom-scripts` option, you must run `vmware-toolbox-cmd` as a root user with the `config` command:

```
vmware-toolbox-cmd config set deployPkg enable-custom-scripts true
cat /etc/vmware-tools/tools.conf
[deployPkg]
enable-custom-scripts = true
```

To verify that you set the option correctly, you can run the following command:

```
vmware-toolbox-cmd config get deployPkg enable-custom-scripts
[deployPkg] enable-custom-scripts = true
```

For more information, see the *VMware Tools User Guide*.

- To customize the guest operating system of a virtual machine with cloud-init metadata and user data:
 - Verify that VMware Tools version 11.3.0 or later is installed.
 - Verify that cloud-init version 21.1 or later is installed.

- Click the menu icon () and click **Policies and Profiles**.
- Click **VM Customization Specifications**.
- To create a specification, from the **VM Customization Specifications** pane, click **New**.
The **New VM Guest Customization Specification** wizard opens.
- On the **Name and target OS** page, enter a name and a description for the customization specification, and select **Linux** as a target guest OS.
- Optional: To initiate customization with raw cloud-init data:
 - Select the **Use cloud-init config data** check box and click **Next**.
 - On the **Cloud-init metadata** page, apply a cloud-init metadata to the guest operating system of the VM.
For more information about the metadata and user data format, see the VMware KB article at: <https://kb.vmware.com/s/article/82250>.

Option	Action
Upload a cloud-init metadata file	The cloud-init metadata file is a plain-text in YAML or JSON file format. The maximum file size is 512 KB. Click Upload and navigate to the file on your local machine. The contents of the file appear in the text box.

Option	Action
Enter a cloud-init metadata script	Enter the metadata script directly into the text box.

c) Click **Next**.

d) Optional: On the **Cloud-init user data**, apply a cloud-init user configuration to the guest operating system of the VM. For example, you can add a new user to the guest operating system.

For more information about the cloud-init user data, see the Cloud-init documentation at: <https://cloudinit.readthedocs.io/en/latest/topics/format.html>.

Upload a cloud-init user data file	The cloud-init user data file is a plain-text in raw cloud-init format. The maximum file size is 512 KB. Click Upload and navigate to the file on your local machine. The contents of the file appear in the text box.
Enter a cloud-init user data	Enter the cloud-init user data script directly into the text box.

e) Click **Next**.

f) On the **Ready to complete** page, review the details and click **Finish** to save your changes.

6. Optional: To initialize traditional customization:

a) On the **Computer name** page, enter a computer name for the guest operating system and a domain name.

The guest operating system uses the computer name to identify itself on the network. On Linux system, it is called the host name.

Option	Action
Use the virtual machine name	Select this option to use the virtual machine name. The computer name that vCenter Server creates is identical to the name of the virtual machine on which the guest operating system is running. If the name exceeds 63 characters, it is truncated.
Enter a name in the Clone/Deploy wizard	Select this option to be prompted to enter a name during cloning or deployment.
Enter a name	- Enter a name. The name can contain alphanumeric characters and a hyphen (-). It cannot contain a period (.), blank spaces, or special characters, and cannot contain digits only. Names are not case-sensitive. - To ensure that the name is unique, select the Append a numeric value check box. This action appends a hyphen followed by a numeric value to the virtual machine name. The name is truncated if it exceeds 63 characters when combined with the numeric value.

Option	Action
Generate a name using the custom application configured with vCenter Server	Enter a parameter that can be passed to the custom application.

- b) Enter the **Domain Name** for the computer and click **Next**.
- c) On the **Time zone** page, select the time zone for the virtual machine and click **Next**.
- d) On the **Customization script** page, apply a customization script to the guest operating system of the VM and click **Next**.

Option	Action
Upload a file containing the customization script	Click Browse and navigate to the file on your local machine. The contents of the script appear in the Script text box.
Enter the customization script	Enter the customization script directly into the Script text box.

The customization script cannot exceed 1500 characters.

Note:

The default timeout period for the guest customization to complete is set to 100 seconds and includes the time for the script to run when you use a "precustomization" command-line parameter. If you run scripts that take a time exceeding the timeout, the guest customization fails.

When you add a customization script with the "precustomization" command-line parameter, it is called before the guest customization begins. As a result, the virtual NIC is disconnected and you cannot access the network.

When you add a customization script with the "postcustomization" command-line parameter, it is called after the guest customization finishes. As a result, the script is scheduled in the initialization process after the virtual machine powers on, the NIC is connected, and you can access the network. The time for the script to run is not included in the default timeout period and you avoid a guest customization failure.

Customization Script Example

```
#!/bin/sh
if [ x$1 == x"precustomization" ]; then
echo Do Precustomization tasks
elif [ x$1 == x"postcustomization" ]; then
echo Do Postcustomization tasks
fi
```

- e) On the **Network** page, select the type of network settings to apply to the guest operating system and click **Next**.

Option	Action
Use standard network settings	Select Use standard network settings so that vCenter Server configures all network interfaces from a DHCP server by using the default settings.
Manually select custom settings	- Select a network adapter from the list or add a new one. - For the selected NIC, click Edit. The Edit Network dialog box opens.

Option	Action
	- To configure the virtual machine to use an IPv4 network, click the IPv4 tab. If you select the Prompt the user for an IPv4 address when the specification is used option, vCenter Server prompts for an IP address when you select to apply the customization specification during cloning or deployment. You are also prompted to configure the gateways during cloning and deployment. - To configure the virtual machine to use an IPv6 network, click the IPv6 tab. If you select the Prompt the user for an address when the specification is used option, vCenter Server prompts for an IP address when you select to apply the customization specification during cloning or deployment. You are also prompted to configure the gateways during cloning and deployment. - Click OK

f) On the **DNS settings** page, enter DNS server and domain settings.

The **Primary DNS**, **Secondary DNS**, and **Tertiary DNS** text boxes accept both IPv4 and IPv6 addresses.

g) On the **Ready to complete** page, review the details and click **Finish** to save your changes.

The customization specification that you created is listed in the Customization Specification Manager. You can use the specification to customize virtual machine guest operating systems.

Create a Customization Specification for Windows

Save specific Windows guest operating system settings in a customization specification, which you can apply when cloning virtual machines or deploying from templates.

Ensure that all requirements for customization are met. See [Guest Operating System Customization Requirements](#).

Note: The default administrator password is not preserved for Windows Server 2008 after customization. During customization, the Windows Sysprep utility deletes and recreates the administrator account on Windows Server 2008. You must reset the administrator password when the virtual machine starts the first time after customization.

Starting with vSphere 8.0 Update 2, you can specify an organizational unit (OU) for your remote desktops. An OU is a subdivision in Active Directory that contains users, groups, computers, or other OUs.

1. Select **Menu > Policies and Profiles**, and under Policies and Profiles, click **VM Customization Specifications**.

2. Click the **Create a new specification** icon.

The **New VM Guest Customization Specification** wizard opens.

3. On the **Name and target OS** page, enter a name and description for the customization specification and select **Windows** as a target guest OS.

4. Optional: Select the **Generate a new security identity (SID)** option and click **Next**.

A Windows Security ID (SID) is used in some Windows operating systems to uniquely identify systems and users. If you do not select this option, the new virtual machine has the same SID as the virtual machine or template from which it was cloned or deployed.

Duplicate SIDs do not cause problems when the computers are part of a domain and only domain user accounts are used. However, if the computers are part of a Workgroup or local user accounts are used, duplicate SIDs can compromise file access controls. For more information, see the documentation for your Microsoft Windows operating system.

5. On the **Set Registration Information** page, enter the virtual machine owner's name and organization and click **Next**.
6. On the **Computer name** page, enter a computer name for the guest operating system and a domain name.
The operating system uses the computer name to identify itself on the network. On Linux systems, it is called the host name.

Option	Action
Use the virtual machine name	Select this option to use the virtual machine name. The computer name that creates is identical to the name of the virtual machine on which the guest operating system is running. If the name exceeds 63 characters, it is truncated.
Enter a name in the Clone/Deploy wizard	Select this option to be prompted to enter a name during cloning or deployment.
Enter a name	1. Enter a name. The name can contain alphanumeric characters and a hyphen (-). It cannot contain a period (.), blank spaces, or special characters, and cannot contain digits only. Names are not case-sensitive. 2. To ensure that the name is unique, select the Append a numeric value check box. This action appends a hyphen followed by a numeric value to the virtual machine name. The name is truncated if it exceeds 63 characters when combined with the numeric value.
Generate a name using the custom application configured with	Enter a parameter that can be passed to the custom application.

7. On the **Windows license** page, provide licensing information for the Windows operating system and click **Next**.

Option	Action
For nonserver operating systems	Type the Windows product key for the new guest operating system.
For server operating systems	1. Type the Windows product key for the new guest operating system. 2. Select Include Server License Information. 3. Select either Per seat or Per server. 4. If you select Per server, enter the maximum number of simultaneous connections for the server to accept.

8. On the **Set Administrator Password** page, configure the administrator password for the virtual machine and click **Next**.
 - a) Enter a password for the administrator account and confirm the password by typing it again.
 - b) Optional: Select the **Automatically logon as Administrator** check box to log users in to the guest operating system as Administrator, and select the number of times to log in automatically.
9. On the **Time zone** page, select the time zone for the virtual machine and click **Next**.
10. Optional: On the **Run Once** page, specify commands to run the first time a user logs in to the guest operating system and click **Next**.
See the Microsoft Sysprep documentation for information about `RunOnce` commands.

11. On the **Network** page, select the type of network settings to apply to the guest operating system and click **Next**.

- Select **Use standard network settings** so that configures all network interfaces from a DHCP server by using the default settings.
- Select **Manually select custom settings** and configure each network interface yourself.
 - a. Select a network adapter from the list or add a new one.
 - b. For the selected NIC, click the vertical ellipsis icon and select **Edit**.
The **Edit Network** dialog box opens.
 - c. Click the **IPv4** tab to configure the virtual machine to use IPv4 network.
You can configure all the settings at that stage or you can select the **Prompt the user for an IPv4 address when the specification is used** option. In that case, prompts for an IP address when you select to apply that customization specification during cloning or deployment. With that option, you can also configure the gateways during cloning or deployment.
 - d. Click the **IPv6** tab to configure the virtual machine to use IPv6 network.
You can configure all the settings at that stage or you can select the **Prompt the user for an address when the specification is used** option. In that case, prompts for an IP address when you select to apply that customization specification during cloning or deployment. With that option, you can also configure the gateways during cloning or deployment.
 - e. Click the **DNS** tab to specify DNS server details.
 - f. Click **WINS** to specify primary and secondary WINS server information.
 - g. Click **OK** to close the **Edit Network** dialog box.

12. On the **Workgroup or Domain** page, select how the virtual machine participates in the network and click **Next**.

Option	Action
Workgroup	Enter a workgroup name. For example, MSHOME.
Windows Server Domain	1. Enter the domain name. 2. To add a computer to the specified domain, enter the user name and password for a user account that has permission. 3. Enter the OU path. For example, OU=MyOU, DC=MyDom, DC=MyCompany, DC=com

13. On the **Ready to complete** page, review the details and click **Finish** to save your changes.

The customization specification that you created is listed in the Customization Specification Manager. You can use the specification to customize virtual machine guest operating systems.

Create a Customization Specification for Windows Using a Custom Sysprep Answer File

A custom sysprep answer file is a file that stores various customization settings such as computer name, licensing information, and workgroup or domain settings. You can supply a custom sysprep answer file as an alternative to specifying many of the settings in the Guest Customization wizard.

Ensure that all requirements for customization are met. See [Guest Operating System Customization Requirements](#).

Windows Server 2003 and Windows XP use a text file called `sysprep.inf`. Windows Server 2008, Windows Vista, and Windows 7 use an XML file called `sysprep.xml`. You can create these files using a text editor, or use the Microsoft Setup Manager utility to generate them. For more information about how to create a custom sysprep answer file, see the documentation for the relevant operating system.

Important: If you use a custom sysprep answer file to deploy a virtual machine with an operating system Windows Vista or later, you must specify the network customization specifications in the sysprep file. The custom network settings that you configure in the **New VM Guest Customization Specification** wizard are not applied. For more information, see VMware KB article 1029174 at <https://kb.vmware.com/s/article/1029174>.

You can prevent Windows from assigning new virtual machines or templates with the same Security IDs (SIDs) as the original virtual machine. Duplicate SIDs do not cause problems when the computers are part of a domain and only domain user accounts are used. However, if the computers are part of a Workgroup or local user accounts are used, duplicate SIDs can compromise file access controls. For more information, see the documentation for your Microsoft Windows operating system.

1. Select **Menu > Policies and Profiles**, and under Policies and Profiles, click **VM Customization Specifications**.
2. Click the **Create a new specification** icon.

The **New VM Guest Customization Specification** wizard opens.

3. On the Name and target OS page, enter a name and description for the customization specification and select **Windows** as a target guest OS.

4. Optional: Select the **Generate a new security identity (SID)** option.

A Windows Security ID (SID) is used in some Windows operating systems to uniquely identify systems and users. If you do not select this option, the new virtual machine has the same SID as the virtual machine or template from which it was cloned or deployed.

Duplicate SIDs do not cause problems when the computers are part of a domain and only domain user accounts are used. However, if the computers are part of a Workgroup or local user accounts are used, duplicate SIDs can compromise file access controls. For more information, see the documentation for your Microsoft Windows operating system.

5. Select **Use Custom Sysprep Answer File** and click **Next**.

6. On the Custom sysprep file, select the option to import or create a sysprep answer file and click **Next**.

Option	Description
Import a Sysprep answer file	Click Browse and browse to the file.
Create a Sysprep answer file	Type the contents of the file in the text box.

7. On the Network page, select the type of network settings to apply to the guest operating system and click **Next**.

- Select **Use standard network settings** so that configures all network interfaces from a DHCP server by using the default settings.
- Select **Manually select custom settings** and configure each network interface yourself.

- a. Select a network adapter from the list or add a new one.

- b. For the selected NIC, click the vertical ellipsis icon and select **Edit**.

The **Edit Network** dialog box opens.

- c. Click the **IPv4** tab to configure the virtual machine to use IPv4 network.

You can configure all the settings at that stage or you can select the **Prompt the user for an IPv4 address when the specification is used** option. In that case, prompts for an IP address when you select to apply that customization specification during cloning or deployment. With that option, you can also configure the gateways during cloning or deployment.

- d. Click the **IPv6** tab to configure the virtual machine to use IPv6 network.

You can configure all the settings at that stage or you can select the **Prompt the user for an address when the specification is used** option. In that case, prompts for an IP address when you select to apply that customization specification during cloning or deployment. With that option, you can also configure the gateways during cloning or deployment.

- e. Click the **DNS** tab to specify DNS server details.

- f. Click **WINS** to specify primary and secondary WINS server information.

- g. Click **OK** to close the **Edit Network** dialog box.

8. On the Ready to complete page, review the details and click **Finish** to save your changes.

The customization specification that you created is listed in the Customization Specification Manager. You can use the specification to customize virtual machine guest operating systems.

Manage Customization Specifications

You can edit, duplicate, export, or delete existing specifications.

1. In the , select **Menu > Policies and Profiles** and click **VM Customization Specifications**.
2. Select a customization specification and select your task.

Option	Description
Edit customization spec	You can make changes to the customization spec, such as changing the networking configuration. Click Edit and make the necessary changes.
Duplicate customization spec	If you need a customization specification that is only slightly different from an existing specification, you can use the Customization Specification Manager to create a copy of the existing specification and modify it. For example, you might need to change the IP address or the administrator password.
Export customization spec	You can export customization specifications and save them as <code>.xml</code> files. To apply an exported specification to a virtual machine, import the <code>.xml</code> file using the Import button.
Delete specification spec	You can remove customization specifications to free up storage.

Import a Customization Specification

You can import an existing specification and use the specification to customize the guest operating system of a virtual machine.

Before you begin, you must have at least one customization specification saved as an xml file located on a file system accessible from the vSphere Client.

1. In the , select **Menu > Policies and Profiles** and click **VM Customization Specifications**.
2. Click the **Import** icon.
3. Browse to the `.xml` file to import, specify a name and optional description, and click **OK**.

The imported specification is added to the list of customization specifications.

Configure User Mappings on Guest Operating Systems

As a vSphere administrator, you can enable guest OS access on certain Single Sign-On (SSO) accounts.

Enabling SSO accounts to log in to a guest OS provides users with additional capabilities to perform administrative tasks on guest virtual machines, such as installing or upgrading the VMware Tools or configuring apps.

Functionality to allow vSphere administrators to configure a guest operating system to use VGAAuth authentication. The vSphere administrator must know the guest administrator password for the enrollment process.

To enroll SSO users to a guest user account, you must enroll SSO users to accounts in guest operating systems. The enrollment process maps a vSphere user to a particular account in the guest by using SSO certificates. Subsequent guest management requests use an SSO SAML token to log in to the guest.

You must configure VMs to accept X.509 certificates. X.509 certificates allow the vSphere administrators in your data center to use SAML tokens issued by single sign-on service to access guest OSs.

View Existing SSO User Mappings

You can view the existing guest user mappings for guest operating systems on the selected virtual machine. You need to authenticate your credentials to view guest mappings.

1. Navigate to the virtual machine and click the **Configure** tab.
2. Click the **Guest User Mappings** tab.
3. To log in to your guest OS account, enter your user name and password, and click **Log In**.
The existing in-guest user mappings are displayed.

Add SSO users to Guest Operating Systems

You can map a new SSO user to a guest user account by creating a user map. Mapping can be established for any type of SSO users, for example solution and users.

Power on the virtual machine.

1. Navigate to the virtual machine and click **Configure** tab.
2. Click the **Guest User Mappings** tab.
3. Enter your user name and password and click **Log In**.
4. In the **Guest User Mappings** pane, click the **Add** button.
The **Add New User Mapping** dialog box opens.
5. From the list of SSO users, select the SSO user that you want to map to a guest account.
6. Specify a guest OS user name and click **OK**.
The SSO user is mapped to a guest user account. A new guest user account is added to the list of **Guest User Mappings**.

Remove SSO Users from Guest Operating Systems

You can remove an existing SSO account from guest user mappings.

Power on your virtual machine.

1. Navigate to a virtual machine and click the **Configure** tab.
2. Click **Guest User Mappings**, enter you user name and password, and click **Log In**.
3. In the **Guest User Mappings** pane, select the SSO user from the list that you want to remove.
4. Click the **Remove** button.
5. Click **Yes** to confirm.
The mapping between the selected SSO user account and guest OS account is removed.

Managing Power States of a Virtual Machine

The basic power operations for a virtual machine include powering on, powering off, suspending, reset, and hard stop. These power options are analogous to the power operations on a physical computer.

Verify that you have the following privileges:

- **VirtualMachine > Interaction > PowerOn**
- **VirtualMachine > Interaction > PowerOff**
- **VirtualMachine > Interaction > Suspend**

- **VirtualMachine > Interaction > Reset**

For information about how to change the virtual machine power states, see [Configure the Virtual Machine Power States](#).

1. Navigate to a virtual machine in the inventory.
2. Right-click the virtual machine or click **Actions**, and select **Power**.
3. Select a power operation.

Option	Description
 Power On	Powers on a virtual machine when the virtual machine is stopped.
 Power Off	Powers off a virtual machine. Powering off a virtual machine might cause loss of data within the guest operating system.
 Suspend	Suspends a running virtual machine and leaves it connected to the network. When you resume a suspended virtual machine, the virtual machine continues operating at the same point the virtual machine was at when it was suspended.
 Reset	Restarts the guest operating system. By using this operation, the virtual machine powers off and then powers on. By resetting a virtual machine, you might lose all unsaved information within the guest operating system.
 Hard stop	You can use this operation only when you cannot power off a virtual machine or if a virtual machine becomes unresponsive. When you hard stop a virtual machine, all processes are ended and the virtual machine is powered off. By using this operation, you might lose all unsaved information. If you use the Hard stop operation, vSphere HA restarts the virtual machine.

Edit Virtual Machine Startup and Shutdown Settings

You can configure virtual machines running on an ESXi host to start up and shut down with the host or after a delay. You can also set the default timing and startup order for virtual machines. This way, the operating system has enough time to save data when the host enters maintenance mode or is being powered off for another reason.

The Virtual Machine Startup and Shutdown (automatic startup) setting is deactivated for all virtual machines residing on hosts that are in a vSphere HA cluster. Automatic startup is not supported with vSphere HA.

Note: You can also create a scheduled task to change the power settings for a virtual machine. See *vCenter Server and Host Management*.

1. In the vSphere Client, navigate to and select the host where the virtual machine is located.
2. Click the **Configure** tab.
3. Under **Virtual Machines**, select **VM Startup/Shutdown** and click **Edit**.
The **Edit VM Startup/Shutdown Configuration** dialog box opens.
4. Select **Automatically start and stop the virtual machines with the system**.
5. Optional: In the **Default VM Settings** pane, configure the default startup and shutdown behavior for all virtual machines on the host.

Setting	Description
Startup Delay	After you start the ESXi host, it starts powering on the virtual machines that are configured for automatic startup. After the ESXi host powers on the first virtual machine, the host waits for the specified delay time and then powers on the next virtual machine. The virtual

Setting	Description
	machines are powered on in the startup order specified under the Default VM Settings pane.
Continue if VMware Tools is started	Shortens the startup delay of the virtual machine. If VMware Tools starts before the specified delay time passes, the ESXi host powers on the next virtual machine without waiting for the delay time to pass.
Shutdown delay	Shutdown delay is the maximum time the ESXi host waits for a shutdown command to complete. When you power off the ESXi host, the autostart manager initiates the automatic shutdown on the first virtual machine and waits within the specific delay time for the virtual machine to complete the power action. The power action can be Power Off, Guest Shutdown, or Suspended. The order in which virtual machines are shut down is the reverse of their startup order. After the ESXi host shuts down the first virtual machine within the time that you specify, the host shuts down the next virtual machine. If a virtual machine does not shut down within the specified delay time, the host runs a power off command and then starts shutting down the next virtual machine. The ESXi host shuts down only after all virtual machines are shut down.
Shutdown action	Select a shutdown action that is applicable to the virtual machines on the host when the host shuts down. • Guest Shutdown • Power Off • Suspend • None

6. Optional: You can also configure the startup order and behavior for individual virtual machines. Use this option when you need the delay of the virtual machine to be different from the default delay for all machines. The settings that you configure for individual virtual machines override their default settings.
 - a) To change the startup order of virtual machines, select a virtual machine from the **Manual Startup** category and use the up arrow to move it up to the **Automatic** or **Automatic Ordered** categories.

Use the up and down arrows to change the startup order for virtual machines in the **Automatic** and **Manual Startup** categories. During shutdown, the virtual machines shut down in the reverse order.
 - b) To edit the startup and shutdown behavior of a virtual machine, select a virtual machine, use the up arrow to move it, and click the **Edit** icon.

The **Virtual Machine Startup/Shutdown setting** dialog box opens.
 - c) In the **Startup Settings** pane, configure the startup behavior of the virtual machine.

You can decide to use the default startup delay or you can specify a new one. If you select **Continue immediately if VMware Tools starts**, the ESXi host powers on the next virtual machine without waiting for the delay to pass.
 - d) In the **Shutdown Settings** pane, configure the shutdown behavior of the virtual machine.

You can use the default shutdown delay or specify a new one and select the shutdown action.
 - e) Click **OK**.
7. Click **OK**.

Using a Virtual Machine Console

You can access a virtual machine console by displaying it in a separate Web browser, or from the VMware Remote Console (VMRC).

From the virtual machine remote console, you can perform tasks in the virtual machine such as installing an operating system, configuring the operating system settings, running applications, monitoring performance, and so on. The vSphere Client offers these choices:

- Launch the Web console to display the VM console in a separate browser tab.
- Download the VMware Remote Console (VMRC) standalone application, which opens in a separate window. The VMware Remote Console standalone application enables you to connect to client devices and launch virtual machine consoles on remote hosts.

Install the VMware Remote Console Application

The VMware Remote Console (VMRC) is a standalone console application. VMRC enables you to connect to client devices and open virtual machine consoles on remote hosts.

Download the VMRC if you do not already have the VMRC installed.

VM-001 | ACTION

Summary Monitor Configure Permissions

Guest OS

Powered Off

LAUNCH REMOTE CONSOLE ⓘ

LAUNCH WE

VM Hardware

CPU

Memory

Hard disk 1

Network adapter 1

CD/DVD drive 1

Compatibility

Virtual Ma

VMware Remote Console

Use VMware Remote Console (VMRC) to connect to client devices and access advanced keyboard features.

If you already have VMware Remote Console installed, it will start automatically when you click the Launch Remote Console button. If not, you can use the link below to download it.

[Download Remote Console](#)

1. In the vSphere Client, navigate to a virtual machine in the inventory.
2. Click the **Summary** tab.
- 3.


From the **Launch Remote Console** link, click the information icon () and click **Download Remote Console**.

4. Download the VMRC installer from the Broadcom website at <https://support.broadcom.com>.

Note:

You must have a profile at <https://support.broadcom.com> to download the VMRC installer.

- a) Log in to the Broadcom Support Portal with your credentials.
- b) From the drop-down menu, select **VMware Cloud Foundation**.



- c) From the side navigation menu, click **My Downloads** and enter `VMware vSphere` in the **Search Product Name** bar.
- d) Click **Show Results** and click **VMware vSphere**.
- e) Click the **Products** tab and expand **VMware vSphere - Enterprise** from the list.
- f) Click the release version and click the **Drivers & Tools** tab.
- g) In the **Search** bar, enter `VMware Remote Console`.
The list with VMRC versions appears.
- h)
Expand the version you want to download and click the download () icon.

Start the VMware Remote Console Application

You can use the standalone VMRC application to connect to client devices.

- Verify that VMRC is installed on your local system. You can download the VMRC installer from the Broadcom website at <https://support.broadcom.com>.
- Verify that you have the **Virtual machine > Interaction > Console interaction** privilege.

With VMRC, you can access the mouse and keyboard connected to remote virtual machines. To perform administrative tasks, you must log in to the VMRC as an administrator.

1. In the vSphere Client, navigate to a virtual machine in the inventory.
2. On the **Summary** tab, click **Launch Remote Console**.
A dialog box opens that requires you to confirm that you want to open the remote console.

3. In the **Open VMware Remote Console** dialog box, confirm that you want to open VMRC.
The VMRC opens as a standalone application for the selected virtual machine. You can also run more than one console to access remotely several virtual machines at the same time.

Open the Web Console

You can access a virtual machine desktop from the vSphere Client by opening the web console and performing various tasks in the virtual machine. For example, you can install an operating system, configure the operating system settings, run applications, monitor performance, and so on.

- Verify that the virtual machine has a guest operating system and that VMware Tools is installed.
 - Verify that the virtual machine is powered on.
 - Verify that you have the **Virtual machine > Interaction > Console interaction** privilege.
1. In the vSphere Client, navigate to a virtual machine in the inventory.
 2. In the **Summary** tab, select **Launch Web Console**.
The console opens in a new browser tab.
 3. Click anywhere inside the console window to start using your mouse, keyboard, and other input devices in the console.

Note: For information about supported international keyboards, refer to the VMware HTML Console SDK Release Notes at <https://www.vmware.com/support/developer/html-console/html-console-21-releasenotes.html#knownissues>.

4. Optional: Click **Send Ctrl-Alt-Delete** to send the Ctrl+Alt+Delete keystroke combination to the guest operating system.
5. Optional: Press Ctrl+Alt to release the pointer from the console window and work outside the console window.
6. Optional: Click **Full Screen** to view the console in full screen mode.
7. Optional: Press Ctrl+Alt+Enter to enter or exit full screen mode.

Manage the VMware Remote Console Proxy Configuration

VMware Remote Console proxy for vSphere (VMRC proxy) is a service in the vCenter Server system which transmits the network traffic between VMRC and the ESXi hosts. When you use the VMRC proxy, VMRC does not require a direct network connection to the ESXi host.

You can activate or deactivate the VMRC proxy. You also can change the VMRC proxy settings to reduce the workload on the vCenter Server system.

Activate the VMware Remote Console Proxy

The VMRC proxy is deactivated by default. You can activate the VMRC proxy service through **Advanced Settings** of your vCenter Server system.

Verify that you have the **Global.Settings** privilege.

1. In the vSphere Client, navigate to and select the vCenter Server instance.
 2. On the **Configure** tab, select **Advanced Settings**.
 3. Click **Edit Settings**.
- The **Edit Advanced vCenter Server Settings** dialog box opens.
4. In the **Name** text box, enter the name of the service - `config.mksdevproxy.enable`.
 5. In the **Value** text box, enter `true` and click **Add**.
 6. Click **Save**.

The proxy setting appears in the list containing all configuration parameters.

Deactivate the VMware Remote Console Proxy

To simplify your vCenter Server configuration, you can deactivate the VMRC proxy.

Verify that you have the **Global.Settings** privilege.

1. In the vSphere Client, navigate to and select the vCenter Server instance.
2. On the **Configure** tab, select **Advanced Settings**.
3. Click **Edit Settings**.

The **Edit Advanced vCenter Server Settings** dialog box opens.

4. Click the **Filter** icon in the **Name** column.
5. To view the VMRC proxy parameter, in the text box, enter `config.mksdevproxy.enable` and close the dialog box.
6. In the **Value** text box, enter `false`, and click **Save**.

Manage the VMware Remote Console Proxy Settings

To ensure that the vCenter Server system performs efficiently, you can change the VMRC proxy settings. You can configure the number of concurrent VMRC proxy connections and the maximum bandwidth for each concurrent VMRC proxy connection.

Configure a Concurrent VMRC Proxy Connection

The VMRC network traffic might impact the work of the vCenter Server instance and you can limit the number of the concurrent VMRC proxy connections.

Verify that you have the **Global.Settings** privilege.

To ensure that the vCenter Server performance is optimum, configure the number of the concurrent VMRC proxy connections to be between 1 and 1024. To meet your configuration needs, you can change the maximum concurrent VMRC proxy connections. By default, the maximum number of concurrent VMRC proxy connections is 32.

1. In the vSphere Client, navigate to and select the vCenter Server instance.
2. On the **Configure** tab, select **Advanced Settings**.
3. Click **Edit Settings**.
4. Configure the number of the concurrent VMRC proxy connections.

Option	Action
Configure a concurrent VMRC proxy connection	1. In the Name text box, enter <code>config.mksdevproxy.connLimit</code>. 2. In the Value text box, enter a value for the maximum number of concurrent connections allowed. 3. Click Add.
Limit the number of the VMRC proxy connections	1. Click the Filter icon in the Name column. 2. To view the VMRC proxy parameter, in the text box enter <code>config.mksdevproxy.connLimit</code> and close the dialog box. 3. In the Value text box, change the number of the concurrent connections.

5. Click **Save**.

Configure the Maximum Bandwidth for a Concurrent VMRC Proxy Connection

You can limit the number of the available network bandwidth.

Verify that you have the **Global.Settings** privilege.

Each concurrent connection has a maximum of 300 KBps network bandwidth. VMRC requires minimum 50 KBps for the basic mouse, keyboard, and screen transfer.

1. In the vSphere Client, navigate to and select the vCenter Server instance.
2. On the **Configure** tab, select **Advanced Settings**.
3. Click **Edit Settings**.
4. Configure the maximum bandwidth for each concurrent VMRC proxy connection.

Option	Action
Configure the VMRC inbound traffic to the vCenter Server system	1. In the Name text box, enter <code>config.mksdevproxy.readthrottler</code>. 2. In the Value text box, enter the value of the inbound traffic in KBps. 3. Click Add.
Configure the vCenter Server outbound traffic to VMRC	1. In the Name text box, enter <code>config.mksdevproxy.writethrottler</code>. 2. In the Value text box, enter the value of the outbound traffic in KBps. 3. Click Add.
Limit the VMRC inbound traffic to the vCenter Server system	1. Click the Filter icon in the Name column. 2. To display the VMRC proxy parameter, in the text box, enter <code>config.mksdevproxy.readthrottler</code> and close the dialog box. 3. In the Value text box, enter the inbound bandwidth limit in KBps.
Limit the vCenter Server outbound traffic to VMRC	1. Click the Filter icon in the Name column. 2. To display the VMRC proxy parameter, in the text box, enter <code>config.mksdevproxy.writethrottler</code> and close the dialog box. 3. In the Value text box, enter the outbound bandwidth limit in KBps.

5. Click **Save**.

Answer Virtual Machine Questions

Virtual machine questions are messages that are generated by vCenter Server. Virtual machine questions appear whenever the virtual machine needs user intervention to continue its operation. In most cases, virtual machine questions appear when you power on a virtual machine .

Verify that the virtual machine hardware version is 11 or later.

To save time and ensure the consistency of your virtual environment, you can apply the same answer to multiple or to all virtual machines that have the same pending question.

1. Navigate to a virtual machine with a question.
2. Right-click the virtual machine and select **Guest OS > Answer Question**.
The **Answer Question** wizard opens.

3. In the **Answer Question** dialog box, select your answer.
4. Optional: Apply the selected answer to other virtual machines that have the same pending question.
 - a) Click the **Select other virtual machines** hyperlink.
A list of all virtual machines with the same pending question appears.
 - b) Select the virtual machines to which to apply the answer.
5. Click **OK**.

Remove and Reregister VMs and VM Templates

You can remove VMs and VM templates from the vCenter Server inventory or delete them from a disk. If you only removed the VM from the inventory, you can add it back from the datastore.

Adding Existing Virtual Machines to vCenter Server

When you add a host to vCenter Server, it discovers all the virtual machines on that managed host and adds them to the vCenter Server inventory.

If a managed host is disconnected, the already discovered virtual machines continue to be listed in the inventory.

If a managed host is disconnected and reconnected, any changes to the virtual machines on that managed host are identified, and the vSphere Client updates the list of virtual machines. For example, if node3 is removed and node4 is added, the new list of virtual machines adds node4 and shows node3 as orphaned.

Remove VMs or VM Templates from vCenter Server or from the Datastore

You can temporarily remove a virtual machine or a VM template from vCenter Server or you can permanently delete it from the datastore.

Power off the virtual machine.

The process is the same for a virtual machine or a VM template:

- When you remove a virtual machine from the inventory, you unregister it from the host and vCenter Server, you do not delete it from the datastore. Virtual machine files remain at the same storage location and you can later re-register the virtual machine by using the datastore browser. This helps if you want to edit the virtual machine configuration file. It is also useful to temporarily remove a virtual machine when you have reached the maximum number of virtual machines that your license or hardware allows.
- If you no longer need a virtual machine and want to free up space on the datastore, you can remove the virtual machine from vCenter Server and delete all virtual machine files from the datastore, including the configuration file and virtual disk files.

Log in to the vSphere Client and perform the task:

Option	Description
Temporarily remove the virtual machine or VM template	1. Right-click the virtual machine. 2. Select Remove From Inventory and click Yes.
Permanently delete the virtual machine or VM template	1. Right-click the virtual machine. 2. Select Delete from Disk and click Yes.

Register a VM or VM Template with vCenter Server

If you removed a VM or VM template from vCenter Server but did not delete it from disk, you can return it to the vCenter Server inventory by registering it with the vCenter Server.

1. In the vSphere Client inventory, right-click the datastore on which the virtual machine configuration file is stored and select **Register VM**.
2. Browse to, select the virtual machine configuration (`.vmx`) file or the VM template configuration file (`.vmtx` file) and click **OK**.

The **Register Virtual Machine** wizard opens.

3. On the Select a name and folder page, use the existing name or type a new name, select a datacenter or folder location and click **Next**.
4. Select a resource pool in which to run the virtual machine and click **Next**.
5. On the Ready to complete page, review your selections and click **Finish**.

Manage Virtual Machine Templates

After you clone a virtual machine to a template, you can perform different operations with the template. You can rename the template, delete it from the vCenter Server inventory, or delete it from the disk. You can also register the template with vCenter Server.

Change the Template Name

If you move a template to another host or a data center folder, you can change the template name to make it unique in that folder.

1. Right-click the template and select **Rename**.
2. Enter a new name and click **OK**.

Delete and Remove Templates

You can delete a template by removing it from the inventory or deleting the template from the disk. If you remove the template from the inventory, it remains on the disk and can be reregistered with vCenter Server to restore it to the inventory.

By removing a template, you unregister it from the vCenter Server inventory, but it remains on the disk at the same storage location and you can use the datastore browser to reregister the template at a later time. For example, if you decide to update the template rather than create one.

Note: Templates that you delete are permanently removed from the system and you cannot recover such templates.

Action	Description
Remove a template from the inventory	1. Right-click the template and select Remove from Inventory. 2. To confirm removing the template from the vCenter Server inventory, click Yes. The template is unregistered from the vCenter Server inventory.
Delete a template from the disk	1. Right-click the template and select Delete from Disk. 2. To confirm removing the template from the datastore, click Yes.

Reregister Templates

You unregister templates when you remove them from the vCenter Server inventory or when the hosts with which they are associated are removed from the vCenter Server inventory.

You can reregister templates to the vCenter Server inventory.

1. In the vCenter Server inventory, navigate to the datastore that contains the template.
2. Select the datastore and click the **Files** tab.
3. Locate the template folder and click it to display the template files.
4. Select the `.vmtx` file and click the **Register VM** icon. **Result:** The **Register VM Template** wizard opens.
5. On the **Select a name and folder** page, specify a name and location for the template and click **Next**.
6. On the **Select a compute resource** page, select a host or cluster on which to store the template and click **Next**.
7. On the **Ready to complete** page, review your selections and click **Finish**.
8. (Optional) To verify that the template is registered, check the host or cluster inventory.

Inventory Object	Action
Host	1. Navigate to the host. 2. Click the VMs tab and select VM Templates.
Cluster	1. Navigate to the cluster. 2. Click the VMs tab and select VM Templates.

Result: The template is registered to the host and you can view the template when you navigate to the host and click **VM Templates**.

Manage Virtual Machines With Snapshots

Snapshots preserve the state and data of a virtual machine at the time you take the snapshot. When you take a snapshot of a virtual machine, an image of the virtual machine in a given state is copied and stored.

Snapshots are useful when you want to revert repeatedly to a virtual machine state, but you do not want to create multiple virtual machines.

You can take multiple snapshots of a virtual machine to create restoration positions in a linear process. With multiple snapshots, you can save many positions to accommodate many kinds of work processes. Snapshots operate on individual virtual machines. Taking snapshots of multiple virtual machines, for example, taking a snapshot of a VM for each member of a team, requires that you take a separate snapshot of each team member's virtual machine.

Snapshots are useful as a short term solution for testing software with unknown or potentially harmful effects. For example, you can use a snapshot as a restoration point during a linear or iterative process, such as installing update packages, or during a branching process, such as installing different versions of a program. Using snapshots ensures that each installation begins from an identical baseline.

With snapshots, you can preserve a baseline before you change a virtual machine.

Several operations for creating and managing virtual machine snapshots and snapshot trees are available in the vSphere Client. These operations enable you to create snapshots, revert any snapshot in the snapshot hierarchy, delete snapshots, and more. You can create snapshot trees where you save the virtual machine state at any specific time so that you can revert that virtual machine state later. Each branch in a snapshot tree can have up to 32 snapshots.

A snapshot preserves the following information:

- Virtual machine settings. The virtual machine directory, which includes the disks added or changed after you take the snapshot.
- Power state. The virtual machine can be powered on, powered off, or suspended.
- Disk state. State of all the virtual machine's virtual disks.

- (Optional) Memory state. The contents of the virtual machine's memory.

The Snapshot Hierarchy

The vSphere Client presents the snapshot hierarchy as a tree with one or more branches. Snapshots in the hierarchy have parent to child relationships. In linear processes, each snapshot has one parent snapshot and one child snapshot, except for the last snapshot, which has no child snapshot. Each parent snapshot can have more than one child. You can revert to the current parent snapshot or to any parent or child snapshot in the snapshot tree and create more snapshots from that snapshot. Each time you revert a snapshot and take another snapshot, a branch (child snapshot) is created.

Parent Snapshots

The first virtual machine snapshot that you create is the base parent snapshot. The parent snapshot is the most recently saved version of the current state of the virtual machine. Taking a snapshot creates a delta disk file for each disk attached to the virtual machine and optionally, a memory file. The delta disk files and memory file are stored with the base `.vmdk` file. The parent snapshot is always the snapshot that appears immediately above the **You are here** icon in the Snapshot Manager. If you revert a snapshot, that snapshot becomes the parent of the **You are here** current state.

Note: The parent snapshot is not always the snapshot that you took most recently.

Child Snapshots

A snapshot of a virtual machine taken after the parent snapshot. Each child snapshot contains delta files for each attached virtual disk, and optionally a memory file that points from the present state of the virtual disk (You are here). Each child snapshot's delta files merge with each previous child snapshot until reaching the parent disks. A child disk can later be a parent disk for future child disks.

The relationship of parent and child snapshots can change if you have multiple branches in the snapshot tree. A parent snapshot can have more than one child. Many snapshots have no children.

CAUTION: Do not manually manipulate individual child disks or any of the snapshot configuration files because doing so can compromise the snapshot tree and result in data loss. This restriction includes disk resizing and making modifications to the base parent disk by using the `vmkfstools` command.

Snapshot Behavior

Taking a snapshot preserves the disk state at a specific time by creating a series of delta disks for each attached virtual disk or virtual RDM and optionally preserves the memory and power state by creating a memory file. Taking a snapshot creates a snapshot object in the Snapshot Manager that represents the virtual machine state and settings.

Each snapshot creates an additional delta `.vmdk` disk file. When you take a snapshot, the snapshot mechanism prevents the guest operating system from writing to the base `.vmdk` file and instead directs all writes to the delta disk file. The delta disk represents the difference between the current state of the virtual disk and the state that existed at the time that you took the previous snapshot. If more than one snapshot exists, delta disks can represent the difference between each snapshot. Delta disk files can expand quickly and become as large as the entire virtual disk if the guest operating system writes to every block of the virtual disk.

Snapshot Files and Limitations

When you take a snapshot, you capture the state of the virtual machine settings and the virtual disk. If you are taking a memory snapshot, you also capture the memory state of the virtual machine. These states are saved to files that reside with the base files of the virtual machines.

Snapshot Files

A snapshot consists of files that are stored on a supported storage device. A Take Snapshot operation creates `.vmdk`, `-delta.vmdk` or `-sesparse.vmdk`, `.vmsd`, and `.vmsn` files. By default, the first and all delta disks are stored with the base `.vmdk` file. The `.vmsd` and `.vmsn` files are stored in the virtual machine directory.

SEsparse is a default format for all delta disks on the VMFS6 datastores.

Delta disk files

A `.vmdk` file to which the guest operating system can write. The delta disk represents the difference between the current state of the virtual disk and the state that existed at the time that the previous snapshot was taken. When you take a snapshot, the state of the virtual disk is preserved, the guest operating system stops writing to it, and a delta or child disk is created.

A delta disk has two files. One is a small descriptor file that contains information about the virtual disk, such as geometry and child-parent relationship information. The other one is a corresponding file that contains the raw data.

The files that make up the delta disk are called child disks or redo logs.

Flat file

A `-flat.vmdk` file that is one of two files that comprises the base disk. The flat disk contains the raw data for the base disk. This file does not appear as a separate file in the Datastore Browser.

Database file

A `.vmsd` file that contains the virtual machine's snapshot information and is the primary source of information for the Snapshot Manager. This file contains line entries, which define the relationships between snapshots and between child disks for each snapshot.

Memory file

A `.vmsn` file that includes the active state of the virtual machine. Capturing the memory state of the virtual machine lets you revert to a powered on virtual machine state. With nonmemory snapshots, you can only revert to a powered off virtual machine state. Memory snapshots take longer to create than nonmemory snapshots. The time the ESXi host takes to write the memory onto the disk depends on the amount of memory the virtual machine is configured to use.

A Take Snapshot operation creates `.vmdk`, `-delta.vmdk` or `-sesparse.vmdk`, and `.vmsn` files and updates `.vmsd` file.

File	Description
The base disk names in a VM are: <code>vmname_number.vmdk</code> and <code>vmname_number-flat.vmdk</code> The newly created disk names in a VM are: <code>vmname_number-number.vmdk</code> , <code>vmname_number-number-sesparse.vmdk</code> or <code>vmname_number-number-delta.vmdk</code>	Snapshot file that represents the difference between the current state of the virtual disk and the state that existed at the time the previous snapshot was taken. The base disk filename uses the following syntax, <code>vm_000001.vmdk</code> where <code>vm</code> is the name of the virtual machine and the six-digit number, <code>000001</code>, is based on the files that already exist in the directory. The number does not consider the number of disks that are attached to the virtual machine. The newly created disk filenames are, for example: <code>test-512-000001-sesparse.vmdk</code> <code>test-512-000001.vmdk</code>

File	Description
	- test-512_1-000001-sesparse.vmdk - test-512_1-000001.vmdk
<i>vmname.vmsd</i>	Database of the virtual machine's snapshot information and the primary source of information for the Snapshot Manager.
<i>vmname-Snapshotnumber.vmsn</i>	Memory state of the virtual machine at the time you take the snapshot. The filename uses the following syntax, <i>vm1-Snapshot1.vmsn</i>, where <i>vm1</i> is the virtual machine name, and <i>Snapshot1</i> is the first snapshot. Note: A <i>.vmsn</i> file is created each time you take a snapshot, regardless of the memory selection. A <i>.vmsn</i> file without memory is much smaller than one with memory.

Snapshot Limitations

Snapshots can affect the virtual machine performance and do not support some disk types or virtual machines configured with bus sharing. Snapshots are useful as short-term solutions for capturing point-in-time virtual machine states and are not appropriate for long-term virtual machine backups.

- VMware does not support snapshots of raw disks, RDM physical mode disks, or guest operating systems that use an iSCSI initiator in the guest.
- Virtual machines with independent disks must be powered off before you take a snapshot.
- You can take a memory snapshot of a virtual machine with an independent disk only to analyze the guest operating system behavior of a virtual machine. You cannot use such snapshots for virtual machine backups because restoring this type of snapshots is unsupported.
- Quiesced snapshots require VMware Tools installation and guest operating system support.
- Snapshots are not supported with PCI vSphere DirectPath I/O devices.
- VMware does not support snapshots of virtual machines configured for bus sharing. If you require bus sharing, consider running backup software in your guest operating system as an alternative solution. If your virtual machine currently has snapshots that prevent you from configuring bus sharing, delete (consolidate) the snapshots.
- Snapshots provide a point-in-time image of the disk that backup solutions can use, but Snapshots are not intended for a robust method of backup and recovery. If the files containing a virtual machine are lost, its snapshot files are also lost. Also, large numbers of snapshots are difficult to manage, consume large amounts of disk space, and are not protected if there is hardware failure.
- Snapshots can negatively affect the performance of a virtual machine. Performance degradation is based on how long the snapshot or snapshot tree is in place, the depth of the tree, and how much the virtual machine and its guest operating system have changed from the time you took the snapshot. Also, you might see a delay in the amount of time it takes the virtual machine to power on. Do not run production virtual machines from snapshots on a permanent basis.

Manage Snapshots

You can view and manage all snapshots for an active virtual machine. You can review the snapshots information, revert to the latest snapshot, change the name and description, or delete a snapshot.

You can manage the snapshots when you select a virtual machine in the vSphere Client inventory and click the **Snapshots** tab.

The snapshot tree displays all snapshots of the virtual machine and the power state of the virtual machine when a snapshot was taken. The detailed information region contains the snapshot name and description, time of creation, and the disk space. You can also see whether you took a snapshot of the virtual machine memory and if you quiesced the guest file system.

The **You are here** pin represents the current and active state of the virtual machine and it is always visible.

Take Snapshots of a Virtual Machine

You can take one or more snapshots of a virtual machine to capture the settings state, disk state, and memory state at specific times. When you take a snapshot, you can also quiesce the virtual machine files and exclude the virtual machine disks from snapshots.

When you take a snapshot, other activity that is occurring in the virtual machine might affect the snapshot process when you revert to that snapshot. The best time to take a snapshot from a storage perspective, is when you are not incurring a large I/O load. The best time to take a snapshot from a service perspective is when no applications in the virtual machine are communicating with other computers. The potential for problems is greatest if the virtual machine is communicating with another computer, especially in a production environment. For example, if you take a snapshot while the virtual machine is downloading a file from a server on the network, the virtual machine continues downloading the file and communicating its progress to the server. If you revert to the snapshot, communications between the virtual machine and the server are confused and the file transfer fails. Depending on the task that you are performing, you can create a memory snapshot or you can quiesce the file system in the virtual machine.

Memory Snapshots

The default selection for taking snapshots. When you capture the virtual machine's memory state, the snapshot retains the live state of the virtual machine. Memory snapshots create a snapshot at a precise time, for example, to upgrade software that is still working. If you take a memory snapshot and the upgrade does not complete as expected, or the software does not meet your expectations, you can revert the virtual machine to its previous state.

When you capture the memory state, the virtual machine's files do not require quiescing. If you do not capture the memory state, the snapshot does not save the live state of the virtual machine and the disks are crash consistent unless you quiesce them.

Quiesced Snapshots

When you quiesce a virtual machine, VMware Tools quiesces the file system of the virtual machine. A quiesce operation ensures that a snapshot disk represents a consistent state of the guest file systems. Quiesced snapshots are appropriate for automated or periodic backups. For example, if you are unaware of the virtual machine's activity, but want several recent backups to revert to, you can quiesce the files.

If the virtual machine is powered off or VMware Tools is not available, the `Quiesce` parameter is not available. You cannot quiesce virtual machines that have large capacity disks.

Important: Do not use snapshots as your only backup solution or as a long-term backup solution.

Change Disk Mode to Exclude Virtual Disks from Snapshots

You can set a virtual disk to independent mode to exclude the disk from any snapshots taken of its virtual machine.

Prerequisites

Power off the virtual machine and delete any existing snapshots before you change the disk mode. Deleting a snapshot involves committing the existing data on the snapshot disk to the parent disk.

Required privileges:

- **Virtual machine > Snapshot management > Remove Snapshot**

- **Virtual machine > Configuration > Modify device settings**

Note: An independent disk does not participate in virtual machine snapshots. That is, the disk state is independent of the snapshot state and creating, consolidating, or reverting to snapshots does not have effect on the disk.

You can take a memory snapshot of a virtual machine with an independent disk, but the snapshot cannot be restored. The main purpose of these snapshots is to capture the state of a virtual machine for further analysis through the Vmss2core or similar tools.

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Hard disk**, and select an independent disk mode option.

Option	Description
Dependent	Dependent mode is the default disk mode. When you take a snapshot of a virtual machine, dependent disks are included in the snapshot. When you revert to the previous snapshot, all data are reverted to the point of taking a snapshot.
Independent - Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk even if you revert a snapshot. When you power off or reset a virtual machine, the disk and all its snapshots are preserved.
Independent - Nonpersistent	Disks in nonpersistent mode behave like read-only disks. Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset the virtual machine, or when you delete a snapshot.

3. Click **OK**.

Take a Snapshot of a Virtual Machine

You can take a snapshot when a virtual machine is powered on, powered off, or suspended. If you are suspending a virtual machine, wait until the suspend operation finishes before you take a snapshot.

When you create a memory snapshot, the snapshot captures the state of the virtual machine memory and the virtual machine power settings. When you capture the virtual machine memory state, the snapshot operation takes longer to complete. You might also see a momentary lapse in response over the network.

When you quiesce a virtual machine, VMware Tools quiesces the file system in the virtual machine. The quiesce operation pauses or alters the state of running processes on the virtual machine, especially processes that might modify information stored on the disk during a revert operation.

Application-consistent quiescing is not supported for virtual machines with IDE or SATA disks.

Note: If you take a snapshot of a Dynamic Disk (a Microsoft-specific disk type), the snapshot technology preserves the quiesce state of the file system, but does not preserve the quiesce state of the application.

Prerequisites

- If you are taking a memory snapshot of a virtual machine that has multiple disks in different disk modes, verify that the virtual machine is powered off. For example, if you have a special purpose configuration that requires you to use an independent disk, you must power off the virtual machine before taking a snapshot.

- To capture the memory state of the virtual machine, verify that the virtual machine is powered on.
- To quiesce the virtual machine files, verify that the virtual machine is powered on and that VMware Tools is installed.
- Verify that you have the **Virtual machine > Snapshot management > Create snapshot** privilege on the virtual machine.

Procedure

1. In the vSphere Client, navigate to a virtual machine and click the **Snapshots** tab.
2. Click **Take Snapshot**.
Result: The **Take snapshot** dialog box opens.
3. Enter a name for the snapshot.
4. (Optional) Enter a description for the snapshot.
5. (Optional) To capture the memory of the virtual machine, select the **Snapshot the virtual machine's memory** check box.
6. (Optional) To pause running processes on the guest operating system so that file system contents are in a known consistent state when you take a snapshot, select the **Quiesce guest file system (requires VMware Tools)** check box.

Note: You can quiesce the virtual machine files only when the virtual machine is powered on and the **Snapshot the virtual machine's memory** check box is deselected.

For more information about the virtual machine behavior when you quiesce a virtual machine, see the KB article <https://kb.vmware.com/s/article/5962168>.

7. Click **Create**.

Revert a Virtual Machine Snapshot

To return a virtual machine to its original state, or to return to another snapshot in the snapshot hierarchy, you can use the revert options.

Verify that you have the **Virtual machine > Snapshot management > Revert to snapshot** privilege on the virtual machine.

When you revert a snapshot, you return the virtual machine's memory, settings, and the state of the virtual machine disks to the state they were in when you took the snapshot. You can revert any snapshot in the snapshot tree and make that snapshot the parent snapshot of the current state of the virtual machine. Subsequent snapshots from this point create a new branch of the snapshot tree.

Restoring snapshots has the following effects:

- The current disk and memory states are discarded, and the virtual machine reverts to the disk and memory states of the parent snapshot.
- Existing snapshots are not removed. You can revert those snapshots at any time.
- If the snapshot includes the memory state, the virtual machine will be in the same power state as when you created the snapshot.

Table 173: Virtual Machine Power State After Restoring a Snapshot

Virtual Machine State When Parent Snapshot Is Taken	Virtual Machine State After Restoration
Powered on (includes memory)	Reverts to the parent snapshot, and the virtual machine is powered on and running.
Powered on (does not include memory)	Reverts to the parent snapshot and the virtual machine is powered off.

Virtual Machine State When Parent Snapshot Is Taken	Virtual Machine State After Restoration
Powered off (does not include memory)	Reverts to the parent snapshot and the virtual machine is powered off.

Virtual machines running certain kinds of workloads can take several minutes to resume responsiveness after reverting from a snapshot.

Note: vApp metadata for virtual machines in vApps does not follow the snapshot semantics for the virtual machine configuration. vApp properties that are deleted, modified, or defined after a snapshot is taken remain intact (deleted, modified, or defined) after the virtual machine reverts to that snapshot or any previous snapshots.

When you revert to a snapshot, disks that you added or changed after the snapshot was taken are reverted to the snapshot point. For example, when you take a snapshot of a virtual machine, add a disk, and revert the snapshot, the added disk is removed.

Independent disks are also removed when you revert to a snapshot that was taken before the disk was added. If the latest snapshot includes an independent disk, its contents do not change when you revert to that snapshot.

1. To revert a snapshot, navigate to a virtual machine in the vSphere Client inventory and click the **Snapshots** tab.
2. Navigate to a snapshot in the snapshot tree, click **Revert**, and click the **Revert** button.

Delete a Snapshot

Deleting a snapshot permanently removes the snapshot from the snapshot tree. The snapshot files are consolidated and written to the parent snapshot disk and merge with the virtual machine base disk. You can delete a single snapshot or all snapshots in a snapshot tree.

- Familiarize yourself with the delete and delete all actions and how they affect virtual machine performance.
- Required Privilege: **Virtual machine > Snapshot management > Remove Snapshot** on the virtual machine.

Deleting a snapshot does not change the virtual machine or other snapshots. Deleting a snapshot consolidates the changes between snapshots and previous disk states. Then it writes all the data from the delta disk that contains the information about the deleted snapshot to the parent disk. When you delete the base parent snapshot, all changes merge with the base virtual machine disk.

To delete a snapshot, a large amount of information must be read and written to a disk. This process can reduce the virtual machine performance until the consolidation is complete. Consolidating snapshots removes redundant disks, which improves the virtual machine performance and saves storage space. The time to delete snapshots and consolidate the snapshot files depends on the amount of data that the guest operating system writes to the virtual disks after you take the last snapshot. If the virtual machine is powered on, the required time is proportional to the amount of data the virtual machine is writing during consolidation.

Failure of disk consolidation can reduce the performance of virtual machines. You can check whether any virtual machines require separate consolidation operations by viewing a list. For information about locating and viewing the consolidation state of multiple virtual machines and running a separate consolidation operation, see [Consolidate Snapshots](#).

Delete

Use the **Delete** option to remove a single parent or child snapshot from the snapshot tree. This option writes disk changes that occur between the state of the snapshot and the previous disk state to the parent snapshot.

Note: Deleting a single snapshot preserves the current state of the virtual machine and does not affect any other snapshot.

You can also use the **Delete** option to remove a corrupt snapshot and its files from an abandoned branch of the snapshot tree without merging them with the parent snapshot.

Delete All

Use the **Delete All** option to delete all snapshots from the snapshot tree. The **Delete all** option consolidates and writes the changes that occur between snapshots and the previous delta disk states to the base parent disk. It then merges them with the base virtual machine disk.

To prevent snapshot files from merging with the parent snapshot if, for example, an update or installation fails, first use the **Revert** button to revert to a previous snapshot. This action invalidates the snapshot delta disks and deletes the memory file. You can then use the **Delete** option to remove the snapshot and any associated files.

CAUTION: Use care when you delete snapshots. You cannot revert a deleted snapshot. For example, you might want to install several browsers, a, b, and c, and capture the virtual machine state after you install each browser. The first, or base snapshot, captures the virtual machine with browser a and the second snapshot captures browser b. If you revert the base snapshot that includes browser a and take a third snapshot to capture browser c, and delete the snapshot that contains browser b, you cannot return to the virtual machine state that includes browser b.

To delete snapshots from a snapshot tree, navigate to a virtual machine in the vSphere Client inventory and click the **Snapshots** tab.

Option	Action
Delete a single snapshot	1. Navigate to and select a snapshot in the snapshots tree. 2. Click Delete and click the Delete button. The snapshot data is consolidated to the parent snapshot and the selected snapshot is removed from the snapshot tree.
Delete all snapshots	1. Click Delete All and click the Delete all button. All immediate snapshots before the You are here current state are consolidated to the base parent disk. All existing snapshots are removed from the snapshot tree and the virtual machine.

Consolidate Snapshots


The presence of redundant delta disks can adversely affect the virtual machine performance. You can combine such disks without violating a data dependency. After consolidation, redundant disks are removed, which improves the virtual machine performance and saves storage space.

Required privilege: **Virtual machine > Snapshot management > Remove Snapshot**

Snapshot consolidation is useful when snapshot disks fail to compress after a **Revert**, **Delete**, or **Delete all** operation. This might happen, for example, if you delete a snapshot but its associated disk does not commit back to the base disk.

1. Navigate to a virtual machine in the vSphere Client inventory and click the **Snapshots** tab.
2. Perform the necessary snapshot operations.
If the virtual machine snapshot files must be consolidated, the **Consolidation is required** message appears.
3. Click the **Consolidate** button.
The **Consolidate** dialog box appears.

4. Click **OK**.
5. To verify that the consolidation is successful, check the **Needs Consolidation** column.
 - a) Navigate to an inventory object that contains a list of virtual machines, for example a vCenter Server instance, a host, or a cluster.
 - b) Click the **VMs** tab and click **Virtual Machines**.
 - c) Click  at the bottom-left corner of the VMs list panel.
The **Show Columns** window appears.
 - d) Select **Needs Consolidation**.



A **Yes** status indicates that the snapshot files for the virtual machine must be consolidated. A **Not Required** status indicates that the files are consolidated.

How do I Schedule Tasks for Snapshot Deletion

By scheduling tasks for snapshot deletion, you can significantly reduce resource consumption in your environment. This is achieved by decreasing the number of VM snapshots and configuring the period in which they are preserved before deletion.

Required privilege: **Virtual machine > Snapshot management > Remove Snapshot** on the virtual machine.

You can set these tasks to run once in the future or at a recurring interval.

The deletion process runs for 3 hours by default. You can set the exact time when it should stop. Ending the deletion process will not halt it for VM snapshots that have already started but will only prevent initiating a deletion task for new snapshots.

1. Navigate to a virtual machine in the inventory and click the **Configure** tab.
2. Select **Scheduled Tasks** and from the **New Scheduled Task** drop-down menu select **Delete Snapshots**.
The **Schedule New Task** window opens.

3. On the **Scheduling Options** page set the scheduling options.

Schedule New Tasks (Delete Snapshots)

1 Scheduling Options

2 Deletion Settings

Scheduling Options

Fields marked with * are required

Task name *

Snapshot Deletion

Description

Delete snapshots of VM-001 each Tuesday.

Target

VM-001

Run *

Weekly

Repeat every1week(s)

On *

MONTUEWEDTHUFRI

Start on05/23/2024

Hour1:30PM

EndNever

Email notification upon completion

example1@email.com,example2@email.com

Option	Description
Task name	Enter a task name.
Description	Enter a brief summary of the criteria and schedule for the deletion process.
Run	Set when the deletion process should occur, the specific day and time, and set the specific day and time when the process should stop. • Once • After vCenter startup • Hourly • Daily • Weekly • Monthly
Email notification upon completion	Enter the email addresses to receive a notification when the deletion process is complete.

4. Click **Next**.
5. On the **Deletion Settings** page, specify the number of days the snapshots older than this will be removed.
6. Click **Schedule The Task**.

How do I Schedule Tasks for Snapshot Deletion for a Group of Virtual Machines

With vSphere 8.0 Update 3, you have the flexibility to schedule a task for snapshot deletion for multiple virtual machines so that you can optimize your work.

Required privileges:

- **Virtual machine > Snapshot management > Remove Snapshot** on the virtual machine.
- **Scheduled task > Create tasks**

To schedule a task for snapshot deletion of multiple virtual machines, perform the following steps.

1. Navigate to and select an ESXi host in the inventory.
2. Click the **VMs** tab and select the virtual machines you want to schedule a snapshot deletion.
3. Right-click the selection and select **Snapshots > Schedule Snapshots Deletion**.
4. To schedule the task, complete the steps in the wizard.

After scheduling snapshot deletion tasks, you can view detailed information about each task, such as the planned period and the exact time the following deletion process will run.

Monitor the Scheduled Tasks for Snapshot Deletion

After you schedule tasks for snapshot deletion on a virtual machine, you can monitor all planned tasks for better visibility of the process.

To observe the scheduled tasks for snapshot deletion, perform the following steps.

1. Navigate to and select a virtual machine in the inventory.
2. Click the **Snapshots** tab and under **Snapshot Management** select **Snapshot Deletion**.

You can view the information for the scheduled tasks.

VM-001 ACTIONS

Summary Monitor Configure Permissions Datastores Networks **Snapshots** Updates

Snapshot Management

Snapshot Deletion

Multiple tasks for snapshots deletion have been scheduled in this Virtual Machine.

Snapshot Deletion	Scheduled
Virtual Machine	VM-001
Configured Tasks	Review the scheduled deletion tasks in Scheduled Tasks .
Delete snapshots older than	3 days
Scheduled	Weekly
Last Run	-
Next Run	06/10/2023, 3:41:00 PM

Enhanced vMotion Compatibility as a Virtual Machine Attribute

Enhanced vMotion Compatibility (EVC) is a cluster feature that ensures CPU compatibility between hosts in a cluster so that you can seamlessly migrate virtual machines within the EVC cluster. You can also activate, deactivate, or change the EVC mode at the virtual machine level.

The per-VM EVC feature facilitates the migration of the virtual machine beyond the cluster and across vCenter Server systems and data centers that have different processors. The per-VM EVC feature facilitates the migration of the virtual machine beyond the cluster and across vCenter Server systems and data centers that have different processors.

Starting with vSphere 7.0 Update 1, you can take advantage of the EVC feature for Virtual Shared Graphics Acceleration (vSGA). vSGA allows multiple virtual machines to share GPUs installed on ESXi hosts and leverage the 3D graphics acceleration capabilities.

The EVC mode of a virtual machine is independent from the EVC mode defined at the cluster level. The cluster-based EVC mode limits the CPU features a host exposes to virtual machines. The per-VM EVC mode determines the set of host CPU features that a virtual machine requires to power on and migrate.

By default, when you power on a newly created virtual machine, it inherits the feature set of its parent EVC cluster or host. However, you can change the EVC mode for each virtual machine separately. You can raise or lower the EVC mode of a virtual machine. Lowering the EVC mode increases the CPU compatibility of the virtual machine. You can also use the API calls to customize the EVC mode further.

Cluster-Level EVC and Per-VM EVC

There are several differences between the way the EVC feature works at the host cluster level and at the virtual machine level.

- Unlike cluster-based EVC, you can change the per-VM EVC mode only when the virtual machine is powered off.
- With cluster-based EVC, when you migrate a virtual machine out of the EVC cluster, a power cycle resets the EVC mode that the virtual machine has. With Per-VM EVC, the EVC mode becomes an attribute of the virtual machine. A power cycle does not affect the compatibility of the virtual machine with different processors.
- When you configure EVC at the virtual machine level, the per-VM EVC mode overrides cluster-based EVC. If you do not configure per-VM EVC, when you power on the virtual machine, it inherits the EVC mode of its parent EVC cluster or host.
- If a virtual machine is in an EVC cluster and the per-VM EVC is also enabled, the EVC mode of the virtual machine cannot exceed the EVC mode of the EVC cluster in which the virtual machine runs. The baseline feature set that you configure for the virtual machine cannot contain more CPU features than the baseline feature set applied to the hosts in the EVC cluster. For example, if you configure a cluster with the Intel "Merom" Generation EVC mode, you should not configure a virtual machine with any other Intel baseline feature set. All other sets contain more CPU features than the Intel "Merom" Generation feature set and as a result of such configuration, the virtual machine fails to power on.

Compatibility and Requirements

The per-VM EVC feature has the following requirements.

Compatibility	Requirement
Host compatibility	
vCenter Server compatibility	
Virtual machine compatibility	Virtual hardware version 14 or later.

To check EVC support for a specific processor or server model, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility/search.php>.

Configure the EVC Mode of a Virtual Machine

Per-VM EVC is deactivated by default. You can activate, deactivate, and change the EVC mode of a virtual machine to ensure its seamless migration across clusters, vCenter Server systems, and data centers that have different processors.

Power off the virtual machine.

1. Navigate to a virtual machine in the vCenter Server inventory.

2. On the **Configure** tab, select **VMware EVC**.

The pane shows details about the EVC mode of the virtual machine and CPUID details.

Important: For newly created virtual machines, the EVC mode that shows in the **VMware EVC** pane is deactivated.

For powered off virtual machines, the **VMware EVC** pane always shows the EVC status defined at the virtual machine level.

For powered on virtual machines with per-VM EVC enabled, the VMware EVC pane shows the EVC status defined at the virtual machine level.

For powered on virtual machines with per-VM EVC deactivated, the VMware EVC pane shows the EVC mode that the virtual machine inherits from its parent EVC cluster or host.

3. Click the **Edit** button.

The **Change EVC Mode** dialog box opens.

4. In the **Change EVC Mode** dialog box, select whether to activate or deactivate EVC.

Option	Description
Disable EVC	The EVC feature is deactivated for the virtual machine. When you power on the virtual machine, it inherits the feature set of its parent EVC cluster or host.
Enable EVC for AMD hosts	The EVC feature is enabled for AMD hosts.
Enable EVC for Intel hosts	The EVC feature is enabled for Intel hosts.
Custom	This option is visible only if you have customized the EVC mode of the virtual machine through the API calls.

5. Optional: From the **CPU Mode** drop-down menu, select a baseline CPU feature set.

Important: If the virtual machine is in an EVC cluster and the per-VM EVC mode exceeds the EVC mode for the cluster, the virtual machine fails to power on. The baseline CPU feature set for the virtual machine must not contain more CPU features than the baseline CPU feature set of the cluster.

6. Optional: From the **Graphics Mode (vSGA)** drop-down menu, select a baseline graphics feature set.

Option	Description
Baseline Graphics	Applies the Baseline Graphics feature set that includes features through Direct3D 10.1/OpenGL 3.3. Note: To configure the vSGA mode to apply the Baseline Graphics set that includes features through Direct3D 10.1/OpenGL 3.3, the virtual machine must be compatible with ESXi 7.0 Update 1 or earlier.
D3D 11.0 class features	Applies the baseline graphics feature set that includes features through Direct3D 11.0/OpenGL 4.3 Note: To configure the vSGA mode to apply the Baseline Graphics set that includes features through Direct3D 11.0/OpenGL 4.3, verify that the virtual machine is compatible with ESXi 8.0 or later and has installed VMware Tools 12.0.0 or later. To power on a virtual machine, configured with Direct3D 11.0/OpenGL 4.3, verify that the ESXi host graphics hardware is available.

7. Click **OK**.

How to Determine the EVC Mode of Your Virtual Machine

The EVC mode of a virtual machine determines the CPU and graphics features that a host must have in order for the virtual machine to migrate to that host and power on. The EVC mode of a virtual machine is independent from the EVC mode that you configure for the cluster in which the virtual machine runs.

The EVC mode of a virtual machine is determined when the virtual machine powers on. At power-on, the virtual machine also determines the EVC mode of the cluster in which it runs. If the EVC mode of a running virtual machine or the entire EVC cluster is raised, the virtual machine does not change its EVC mode until it is powered off and powered on again. This means that the virtual machine does not use any CPU features exposed by the new EVC mode until the virtual machine is powered off and powered on again.

For example, you create an EVC cluster that contains hosts with Intel processors and you set the EVC mode to Intel "Merom" Generation (Xeon Core 2). When you power on a virtual machine in this cluster, it runs in the Intel Merom Generation (Xeon Core 2) EVC mode. If you raise the EVC mode of the cluster to Intel "Penryn" Generation (Xeon 45 nm

Core 2), the virtual machine retains the lower Intel "Merom" Generation (Xeon Core 2) EVC mode. To use the feature set of the higher EVC mode, such as SSE4.1, the virtual machine must be powered off and powered on again.

1. Navigate to a cluster or a host in the vCenter Server inventory.
2. Click the **VMs > Virtual Machines** tab.
A list of all virtual machines in the selected cluster or on the selected host appears.
3. To verify the status of the CPU mode, check the **EVC CPU Mode** column.
 - a) If the column is not visible, click **Manage Columns** at the bottom left corner of the **Virtual Machines** table.

The **Show Columns** pop-up appears.

- b) To show the **EVC CPU Mode** column, locate and select the respective check box from the list.

The **EVC CPU Mode** column shows the CPU modes of all virtual machines in the cluster or on the host.

Important: For each virtual machine, the **EVC CPU Mode** column displays the EVC mode defined at the virtual machine level.

However, if you do not configure per-VM EVC for a virtual machine, the virtual machine inherits the EVC mode of its parent cluster or host. As a result, for all virtual machines that do not have per-VM EVC configured, the **EVC CPU Mode** column displays the inherited EVC mode of the parent host or cluster.

If the virtual machine is in an EVC cluster, the EVC mode that you see in the **EVC CPU Mode** column is defined in the following manner.

- When the virtual machine is powered on, the **EVC CPU Mode** column displays either the per-VM EVC mode, or the cluster-level EVC mode.

Per-VM EVC	Cluster-Level EVC	EVC Mode for the Virtual Machine
Activated	Activated	Enabled. The EVC CPU Mode column displays the EVC mode of the virtual machine.
Deactivated	Activated	Enabled. The EVC CPU Mode column displays the EVC mode of the EVC cluster.

- When the virtual machine is powered off, the **EVC CPU Mode** column displays the per-VM EVC mode. If per-VM EVC is deactivated, the **EVC CPU Mode** column for the virtual machine is empty.

When the virtual machine is not in an EVC cluster and per-VM EVC is not configured, the EVC mode that you see in the **EVC CPU Mode** column is defined in the following manner.

- When the virtual machine is powered on, the **EVC CPU Mode** column displays the EVC mode of the parent host.
- When the virtual machine is powered off, the **EVC CPU Mode** column is empty.

4. To verify the status of the graphics mode, check the **EVC Graphics Mode (vSGA)** column.
 - a) If the column is not visible, click **Manage Columns** at the bottom left corner of the **Virtual Machines** table.
The **Show Columns** pop-up appears.
 - b) To show the **EVC Graphics Mode (vSGA)** column, locate and select the respective check box from the list.
The **EVC Graphics Mode (vSGA)** column displays the baseline graphics features set. To view the baseline graphics, you must enable **3D graphics** in the virtual machine.

Virtual Machine Storage DRS Rules

The Storage DRS rules that you define at the virtual machine level function in the same way as the affinity and anti-affinity rules that you define at the datastore cluster level. Virtual machine Storage DRS rules define whether virtual machine hard disks are placed and kept on the same datastore or on different datastores within a datastore cluster.

In the vSphere Client, you can create, edit, and delete Storage DRS rules. You can also create Storage DRS rules that place and keep all virtual disks of certain virtual machines on different datastores within a datastore cluster.

VMDK Affinity Rules

By default, all virtual machine hard disks are kept together on the same datastore within a datastore cluster that has Storage DRS enabled. That is, VMDK affinity rules are enabled by default for all virtual machines that are in a datastore cluster. You can override that rule for the datastore cluster or for individual virtual machines.

Storage DRS Anti-Affinity Rules

You can also create anti-affinity rules to place certain virtual hard disks or virtual machines on different datastores and keep them separated.

- VMDK anti-affinity rules ensure that two or more virtual hard disks of a single virtual machine are placed and kept on different datastores within the datastore cluster.
- VM anti-affinity rules ensure that all the virtual hard disks of two or more virtual machines are placed and kept on different datastores within the datastore cluster.

For more information about Storage DRS, see the *vSphere Resource Management* documentaion.

Add a VMDK Affinity Rule

Create a VMDK affinity rule to place and keep all virtual disks of a virtual machine on the same datastore within a datastore cluster.

1. Navigate to a virtual machine in the inventory.
2. On the **Configure** tab, expand **Settings** and click **VM SDRS Rules**.
3. Click the **Add** button.
The **Add SDRS rule** dialog box opens.
4. From the **Type** drop-down menu, select **VMDK affinity**.
5. From the **Datastore cluster** drop-down menu, select the datastore cluster where you want to keep the virtual machine disks.
A datastore cluster is visible in the list only if the virtual machine configuration file or at least one of the virtual hard disks is placed on a datastore within the datastore cluster.
6. Optional: Deselect the **Keep VMDKs together** to create a rule that places and keeps the virtual hard disks on different datastores.
If you leave the check box selected, the rule that you create is the same as the default Storage DRS rule that operates at the datastore cluster level.
Deselecting the check box creates a Storage DRS rule that overrides the default VMDK affinity rule for the datastore cluster.
7. Click **OK**.

An Intra-VM Affinity rule is created that keeps VMDKs together. That is, all virtual hard disks of the selected virtual machine are placed and kept together on the same datastore within the datastore cluster.

Add a VMDK Anti-Affinity Rule

Create a VMDK anti-affinity rule to place and keep certain virtual hard disks of a virtual machine on different datastores within the datastore cluster.

When you create an anti-affinity rule, it applies to those of the virtual machine's hard disks that are on datastores within the selected datastore cluster. Anti-affinity Storage DRS rules operate during migrations that Storage DRS initiates or recommends, but they do not operate when the user initiates a migration.

1. Navigate to a virtual machine in the inventory.
2. On the **Configure** tab, expand **Settings** and click **VM SDRS Rules**.
3. Click the **Add** button.
The **Add SDRS rule** dialog box opens.
4. From the **Type** drop-down menu, select **VMDK anti-affinity**.
5. In the **Rule name** text box, enter a name for the rule.
6. From the **Datastore cluster** drop-down menu, select the datastore cluster where the anti-affinity rule will operate.
All virtual hard disks that are placed on datastores within the selected datastore cluster appear at the bottom of the dialog box.
7. Select the virtual hard disks to which the anti-affinity rule applies.
8. Optional: Deselect the **Enable the rule** check box.
Deselecting the **Enable the rule** check box deactivates the rule. You can still create the rule, but after creation the rule is not applied to the selected virtual hard disks.
9. Click **OK**.

A VMDK anti-affinity rule is created. If the rule is enabled, all selected virtual hard disks are placed and kept on different datastores within the datastore cluster.

Add a VM Anti-Affinity Rule

Create a VM anti-affinity rule to place and keep all virtual hard disks of selected virtual machines on different datastores within the datastore cluster.

1. Navigate to a virtual machine in the inventory.
2. On the **Configure** tab, expand **Settings** and click **VM SDRS Rules**.
3. Click the **Add** button.
The **Add SDRS rule** dialog box opens.
4. From the **Type** drop-down menu, select **VM anti-affinity**.
5. In the **Rule name** text box, enter a name for the rule.
6. From the **Datastore cluster** drop-down menu, select the datastore cluster where the anti-affinity rule will operate.
7. From the list of virtual machines at the bottom of the dialog box, select the virtual machines to which the anti-affinity rule applies.
You can add or remove virtual machines to the list.
8. Optional: Deselect the **Enable the rule** check box.
The rule is enabled by default.
Deselecting the **Enable the rule** check box deactivates the rule. If an existing rule is deactivated, it is not applied to the virtual machine.

9. Click **OK**.

A VM anti-affinity rule is created. If the rule is enabled, all virtual hard disks of the selected virtual machines are placed and kept on different datastores within the datastore cluster.

Distributing Content with GuestStore

The GuestStore feature provides an easy and flexible mechanism to maintain and distribute different content types across multiple virtual machines on multiple ESXi hosts at the same time.

By using the GuestStore framework, you can make sure that the distributed content is always consistent and you can improve the content management in your environment.

As a vSphere administrator, after you configure GuestStore on an ESXi host, the virtual machines on the host can start accessing the GuestStore content immediately.


The GuestStore content consists of:

- Binary files that contain updates for VMware Tools and VMware agents.
- Scripts that are either provided by VMware or are custom scripts.
- Configuration files of VMware Tools and VMware agents.

With GuestStore you can:

- Maintain the content on a datastore, that is called a GuestStore repository.
- Fetch and distribute VMware Tools and VMware agent updates.
- Distribute configurations for VMware Tools and different VMware agents.
- Distribute a custom content - custom scripts, agents, and configuration files.

For example, by importing VMware Tools in the GuestStore repository, you can schedule an automated VMware Tools upgrade at the next reboot of the virtual machine. From the GuestStore repository, you can also perform VMware Tools upgrade for particular virtual machines when necessary.



By using vSphere Client, you can configure a GuestStore repository on a shared datastore of an ESXi host. You perform the configuration by using ESXCLI. You manage the contents of your datastore by using the datastore file browser, or, if you use an NFS datastore, you can mount the datastore to any NFS client machine.

By using a PowerCLI script, you can configure multiple hosts or a cluster managed by a vCenter Server system.

To access the GuestStore contents, the virtual machine guest operating system uses a resource path. The path to the datastore contents must be the same as the resource path of the guest operating system. For example, to access the `/example/myrepository/bar` file in the GuestStore repository, the resource path in the guest operating system must be the same, that means `/example/myrepository/bar`.

Starting with vSphere 7.0 Update 2, you can download and extract the VMware Tools content under the GuestStore repository path. The GuestStore framework allows you, as a vSphere Client administrator, to distribute a configuration file or VMware agent to set of various guest operating systems. For information about configuring GuestStore for VMware Tools, see the *VMware Tools* documentation.

Requirements for GuestStore

To use GuestStore, your vSphere environment must meet the following requirements:

- Virtual machines with Windows guest operating systems must be running on ESXi 7.0 Update 2 and later and VMware Tools 11.2.5 and later versions.
- Virtual machines with Linux guest operating systems must be running on ESXi 7.0 Update 3 and later and VMware Tools 11.3.0 and later versions.
- The file that is distributed through GuestStore must be 512 MB or smaller.

Set the GuestStore Repository with ESXCLI

You can use ESXCLI commands to set the URL to the GuestStore repository and verify the configuration by retrieving the currently set URL.

1. Set the GuestStore repository URL.

Note: The URL must point to a datastore path where GuestStore content is stored. The datastore path must be accessible to the ESXi host. If you want to set up a common GuestStore repository across multiple ESXi hosts, you should use a shared datastore path.

```
esxcli system settings gueststore repository set --url "<datastore_path>"
```

The following example contains a possible syntax for the datastore path.

```
esxcli system settings gueststore repository set --url "ds:///vmfs/volumes/<datastore_uuid>/GuestStore"
```

2. Retrieve the GuestStore repository URL.

```
esxcli system settings gueststore repository get
```

You receive the currently set URL in the output.

```
URL: <datastore_path>
```

Clear the GuestStore Repository Setting with ESXCLI

You can use ESXCLI commands to clear the GuestStore repository URL setting and verify that the URL is not set.

1. Clear the GuestStore repository URL setting.

```
esxcli system settings gueststore repository set --url ""
```

2. Retrieve the GuestStore repository URL.

```
esxcli system settings gueststore repository get
```

You receive the URL information in the output.

```
URL: <not set>
```

Sharing Data Between the vSphere Client and a Virtual Machine Guest Operating System with DataSets

The DataSets feature provides a way to share data between the vSphere Client and a virtual machine guest operating system. By using DataSets, you can exchange data which can be used to configure the guest operating system and the software applications and agents running inside the guest operating system.


The DataSets feature is intended for information that does not change frequently. You must organize the intended information in DataSets entries. Persistent data in DataSets may be any key-value data, which then the software applications and agents can use.

You can create key-value pairs within DataSets, and thereafter list the keys, update their values, or delete key-value pairs.

You can modify DataSets by using software applications regardless of the power state of the virtual machine. To avoid conflict with other applications, each application using DataSets must store its data in at least one unique DataSets entry. You can use the DataSets feature to implement:

- Guest deployment status - when you perform guest provisioning that includes running deployment scripts. DataSets can present configuration data and deployment scripts to the guest operating system. The guest operating system then reports back status including possible errors.
- Guest agent configuration - when a management application needs to configure an in-guest agent. Through the management side you can create and populate the DataSets with configuration data, and the guest agent reads and responds to that data.
- Guest related metadata or other items.

DataSets Architecture



The DataSets feature supports VM operations such as migration of a virtual machine with vMotion, cloning a virtual machine, and taking snapshots of a virtual machine. You can use the `omitFromSnapshotAndClone` parameter to preserve the DataSets information when you clone a virtual machine, create, or revert a snapshot.

For example, if you set `omitFromSnapshotAndClone` to `true`, the DataSets information is not preserved when you revert a snapshot.

If you set `omitFromSnapshotAndClone` to `false` and take a snapshot of a virtual machine, the DataSets information is included in snapshots. When you revert to the previous snapshot, the DataSets information is reverted to the point of taking a snapshot.

Requirements for DataSets

To use DataSets, your vSphere Client environment must meet the following requirements:

- vCenter Server 8.0 or later.
- The ESXi host running in your environment must be at ESXi 8.0 or later.
- Virtual machines must be of hardware version 20 or later.
- A virtual machine supports maximum 1024 DataSets and each DataSet can have maximum 1024 entries.
- Each virtual machine can have a maximum of 100 MB DataSets data.
- Keys must be no larger than 4 KB and values no larger than 1 MB.

For more information about Data Sets, see the *VMware Guest SDK Programming Guide*.

vSphere Virtual Machines Migration

With vSphere vMotion, you can move virtual machines from one compute resource or storage location to another by using cold or hot migration.

For example, with vSphere vMotion you can move powered on virtual machines away from a host to perform maintenance, to balance loads, to colocate virtual machines that communicate with each other, to move virtual machines apart to minimize fault domain, to migrate to new server hardware, and so on.

Moving a virtual machine from one inventory folder to another folder or resource pool in the same data center is not a form of migration. Unlike migration, cloning a virtual machine or copying its virtual disks and configuration file in the same vCenter Server system are procedures that create a new virtual machine. Cloning and copying a virtual machine in the same vCenter Server system are also not forms of migration.

By using migration, you can change the compute resource that the virtual machine runs on. For example, you can move a virtual machine from one host to another host or cluster.

To migrate virtual machines with disks larger than 2 TB, the source and destination ESXi hosts must be version 6.0 and later.

Depending on the power state of the virtual machine that you migrate, migration can be cold or hot.

Cold Migration

Moving a powered off or suspended virtual machine to a new host. Optionally, you can relocate configuration and disk files for powered off or suspended virtual machines to new storage locations. You can also use cold migration to move virtual machines from one virtual switch to another, and from one data center to another. You can perform cold migration manually or you can schedule a task.

Hot Migration

Moving a powered on virtual machine to a new host. Optionally, you can also move the virtual machine disks or folder to a different datastore. Hot migration is also called live migration or vSphere vMotion. With vSphere vMotion, you migrate the virtual machine without any interruption in its availability.

Depending on the virtual machine resource type, you can perform three types of migration.

Change compute resource only

Moving a virtual machine, but not its storage, to another compute resource, such as a host, cluster, resource pool, or vApp. You can move the virtual machine to another compute resource by using cold or hot migration. If you change the compute resource of a powered on virtual machine, you use vSphere vMotion.

Change storage only

Moving a virtual machine and its storage, including virtual disks, configuration files, or a combination of these, to a new datastore on the same host. You can change the datastore of a virtual

Change both compute resource and storage

machine by using cold or hot migration. If you move a powered on virtual machine and its storage to a new datastore, you use Storage vMotion.

Moving a virtual machine to another host and at the same time moving its disk or virtual machine folder to another datastore. You can change the host and datastore simultaneously by using cold or hot migration.

In vSphere 6.0 and later, you can move virtual machines between vSphere sites by using migration between the following types of objects.

Migrate to another virtual switch

Moving the network of a virtual machine to a virtual switch of a different type. You can migrate virtual machines without reconfiguring the physical and virtual network. By using cold or hot migration, you can move the virtual machine from a standard to a standard or distributed switch, and from a distributed switch to another distributed switch. When you move a virtual machine network between distributed switches, the network configuration and policies that are associated with the network adapters of the virtual machine are transferred to the target switch.

Migrate to another data center

Moving a virtual machine to a different data center. You can change the data center of a virtual machine by using cold or hot migration. For networking in the target data center, you can select a dedicated port group on a distributed switch.

Migrate to another vCenter Server system

Moving a virtual machine to a different vCenter Server instance.

You can move a virtual machine to a vCenter Server instance that is connected to the source vCenter Server instance through vCenter Enhanced Linked Mode.

You can also move virtual machines between vCenter Server instances that are located across a long distance from each other.

Starting with vSphere 7.0 Update 1c, you can migrate workloads across vCenter Server systems by using Advanced Cross vCenter vMotion. You can initiate migration of workloads both from on-premise environments and from cloud environments. Advanced Cross vCenter vMotion does not depend on vCenter Enhanced Linked Mode or Hybrid Linked Mode and you can migrate virtual machines across vCenter Server systems in different vCenter Single Sign-On domains.

Starting with vSphere 7.0 Update 3, you can use the Advanced Cross vCenter vMotion feature to clone virtual machines across vCenter Server systems.

How do I Upgrade Virtual Machines

You can upgrade virtual machines to a higher level of compatibility and a higher version of VMware tools. After the upgrade, your VMs can take advantage of new hardware options and new features.

For a list of hardware features available to virtual machines with each ESXi hardware compatibility setting, see [Hardware Features Available with Virtual Machine Compatibility Settings](#).

VMware Tools Upgrade

The first step in upgrading virtual machines is to upgrade VMware Tools. Installing VMware Tools is part of the process of creating a new virtual machine. If you are installing VMware Tools in multiple virtual machines with Windows guest

operating systems, you can automate its installation and specify options for the components to include or exclude. For information about installing, upgrading, and configuring VMware Tools, see the *VMware Tools User Guide*.

If the virtual machines do not have VMware Tools installed, you can use the VMware Tools upgrade procedure to install VMware Tools. After you install or upgrade VMware Tools, upgrade the virtual machine compatibility.

Virtual Machine Compatibility Upgrade

VMware offers the following tools for upgrading virtual machines:

vSphere Client

Requires that you perform the virtual machine upgrade one step at a time, but does not require vSphere Lifecycle Manager.

In the vSphere Client, you can upgrade virtual machines manually, or schedule upgrades.

Manual Upgrade

Use this procedure to upgrade one or more virtual machines to the latest supported virtual hardware version immediately.

Schedule VM Upgrades

Use this procedure to schedule an upgrade of one or more virtual machines at the next reboot of the virtual machine, and choose from all supported compatibility level upgrades.

vSphere Lifecycle Manager

Automates the process of upgrading and patching virtual machines, ensuring that the steps occur in the correct order. You can use vSphere Lifecycle Manager to directly upgrade virtual machine hardware, VMware Tools, and virtual appliances. You can also patch and update third-party software running on the virtual machines and virtual appliances. See the *vSphere Lifecycle Manager* documentation.

Note: Do not use `vmware-vmupgrade.exe` to upgrade virtual machines.

Note: Upgrading virtual machine hardware is a heavyweight operation that might cause some applications or the operating system to stop working properly.

Downtime for Upgrading Virtual Machines

During a virtual machine compatibility upgrade, you must shut down the virtual machine for all guest operating systems. For VMware Tools upgrade, downtime is not required for many Linux operating systems.

Table 174: Virtual Machine Downtime by Guest Operating System

Guest Operating System	Upgrade VMware Tools	Upgrade Virtual Machine Compatibility
Microsoft Windows	Downtime to restart the guest operating system.	Downtime to shut down and power on the virtual machine.
Linux	Downtime to restart the guest operating system is required to load drivers.	Downtime to shut down and power on the virtual machine.
NetWare	No downtime.	Downtime to shut down and power on the virtual machine.
Solaris	No downtime.	Downtime to shut down and power on the virtual machine.

Guest Operating System	Upgrade VMware Tools	Upgrade Virtual Machine Compatibility
FreeBSD	No downtime.	Downtime to shut down and power on the virtual machine.
Mac OS X	No downtime.	Downtime to shut down and power on the virtual machine.

Note: For Linux guest operating systems, VMXNET3 and PVSCSI drivers are embedded in the Linux kernel. Linux virtual machines do not use VMware Tools to load VMXNET3 and PVSCSI drivers.

To load the new versions of the VMXNET driver, you must restart the virtual machine or manually reload the driver. Manual restart is not required for the Linux guest operating system using kernel version 3.10.

To verify that the drivers are configured in the Linux kernel and that the virtual hardware is available, see the VMware Knowledge Base article at: <http://kb.vmware.com/kb/2050364>.

Planning Downtime for Virtual Machines

You can stagger virtual machine downtimes to accommodate a schedule convenient to you and your customers.

For example:

- If your virtual machine users are located in diverse time zones, you can prepare by migrating virtual machines to specific hosts to serve a given time zone. This way you can arrange host upgrades so that virtual machine downtime occurs transparently outside business hours for that time zone.
- If your virtual machine users operate around the clock, you can delay downtime for their virtual machines to normally scheduled maintenance periods. You do not need to upgrade any stage within a certain time period. You can take as long as needed at any stage.

Upgrade the Compatibility of a Virtual Machine Manually

The virtual machine compatibility determines the virtual hardware available to the virtual machine, which corresponds to the physical hardware available on the host machine. You can upgrade the compatibility level to make additional hardware available to the virtual machine.

- Create a backup or snapshot of the virtual machines. See [Manage Virtual Machines With Snapshots](#).
- Upgrade VMware Tools. On Microsoft Windows VMs, the virtual machine might lose its network settings if you upgrade the compatibility before you upgrade VMware Tools.
- Verify that all virtual machines and their `.vmdk` files are stored on storage connected to the ESXi host or the client machine.
- Determine the ESXi versions that you want the virtual machines to be compatible with. See [Virtual Machine Compatibility](#).

- Check whether the guest operating systems of the virtual machines that you upgrade require a power off. For example, some Linux operating systems do not require a power-off before a virtual machine compatibility upgrade. See [Virtual Machine Compatibility](#).

Important: Upgrading virtual machine hardware might cause some applications or the operating system to stop working properly. Perform a hardware version upgrade only if you need a feature that comes with the newer hardware version.

1. In the vSphere Client, navigate to the virtual machine.
2. Optional: Right-click the virtual machine and select **Power > Power Off**.
3. Select **Actions > Compatibility > Upgrade VM Compatibility**.
4. Click **Yes** to confirm the upgrade.
5. Select a compatibility and click **OK**.

Schedule a Compatibility Upgrade for a Virtual Machine

The virtual machine compatibility determines the virtual hardware available to the virtual machine, which corresponds to the physical hardware available on the host. You can schedule a compatibility upgrade to make a virtual machine compatible with newer versions of ESXi.

- Power off the virtual machine.
- Create a backup or snapshot of the virtual machine. See [Manage Virtual Machines With Snapshots](#).
- Upgrade to the latest version of VMware Tools. If you upgrade the compatibility before you upgrade VMware Tools, the virtual machine might lose its network settings.
- Verify that all `.vmdk` files are available to the ESX/ ESXi host on a VMFS5 or NFS datastore.
- Verify that the virtual machine is stored on VMFS5 or NFS datastores.
- Verify that the compatibility settings for the virtual machine are not the latest supported version.
- Determine the ESXi version that you want the virtual machine to be compatible with. See [Virtual Machine Compatibility](#).

Use this procedure to schedule an upgrade for one virtual machine at the next reboot of the virtual machine, and choose from all supported compatibility level upgrades. To upgrade virtual machines immediately to the latest supported compatibility, see [Upgrade the Compatibility of a Virtual Machine Manually](#).

You can use this procedure to schedule an upgrade for multiple virtual machines.

For information about virtual machine hardware versions and compatibility, see [Virtual Machine Compatibility](#).

1. Navigate to a virtual machine in the inventory.
2. Right-click the virtual machine and select **Compatibility > Schedule VM Compatibility Upgrade**.
3. In the **Schedule VM Compatibility Upgrade** dialog box, confirm that you want to schedule a compatibility upgrade by clicking **Yes**.
4. From the **Compatible with** drop-down menu, select the compatibility to upgrade to.
The virtual machine compatibility is upgraded the next time you restart the virtual machine.
5. Optional: To upgrade the compatibility when you do regularly scheduled guest maintenance, select **Only upgrade after normal guest OS shutdown**.
This prevents the scheduled upgrade from occurring unless the guest operating system of the virtual machine is shut down or restarted normally.

Each of the selected virtual machines is upgraded to the compatibility that you chose at the next reboot of the virtual machine, and the Compatibility setting is updated in the Summary tab of the virtual machine.

Required vCenter Server Privileges for Common Tasks

Many tasks require permissions on multiple objects in the vSphere inventory. If the user who attempts to perform the task only has privileges on one object, the task cannot complete successfully.

The following table lists common tasks that require more than one privilege. You can add permissions to inventory objects by pairing a user with one of the predefined roles or with multiple privileges. If you expect that you must assign a set of privileges multiple times, create custom roles. To learn more about required privileges for common tasks, see [Privilege Recorder](#).

Refer to the *vSphere Web Services API Reference* documentation to learn how operations in the vSphere Client user interface map to API calls, and what privileges are required to perform operations. For example, the API documentation for the `AddHost_Task(addHost)` method specifies that the `Host.Inventory.AddHostToCluster` privilege is required to add a host to a cluster.

If the task that you want to perform is not in this table, the following rules explain where you must assign permissions to allow particular operations:

- Any operation that consumes storage space requires the **Datastore > Allocate Space** privilege on the target datastore, and the privilege to perform the operation itself. You must have these privileges, for example, when creating a virtual disk or taking a snapshot.
- Moving an object in the inventory hierarchy requires appropriate privileges on the object itself, the source parent object (such as a folder or cluster), and the destination parent object.
- Each host and cluster has its own implicit resource pool that contains all the resources of that host or cluster. Deploying a virtual machine directly to a host or cluster requires the **Resource > Assign Virtual Machine to Resource Pool** privilege.

Table 175: Required Privileges for Common Tasks

Task	Required Privileges	Applicable Role
Create a virtual machine	On the destination folder or data center: Virtual machine > Edit inventory > Create new Virtual machine > Change Configuration > Add new disk (if creating a new virtual disk) Virtual machine > Change Configuration > Add existing disk (if using an existing virtual disk) Virtual machine > Configuration > Configure Raw device (if using an RDM or SCSI pass-through device)	Administrator
	On the destination host, cluster, or resource pool: Resource > Assign virtual machine to resource pool	Resource pool administrator or Administrator
	On the destination datastore or the folder that contains the datastore: Datastore > Allocate space	Datastore Consumer or Administrator
	On the network that the virtual machine will be assigned to: Network > Assign network	Network Consumer or Administrator
Power on a virtual machine	On the data center in which the virtual machine is deployed: Virtual machine > Interaction > Power On	Virtual Machine Power User or Administrator
	On the virtual machine or folder of virtual machines: Virtual machine > Interaction > Power On	
Deploy a virtual machine from a template	On the destination folder or data center: Virtual machine > Edit inventory > Create from existing	Administrator

Task	Required Privileges	Applicable Role
	Virtual machine > Change Configuration > Add new disk	
	On a template or folder of templates: Virtual machine > Provisioning > Deploy template	Administrator
	On the destination host, cluster or resource pool: Resource > Assign virtual machine to resource pool vApp > Import	Administrator
	On the destination datastore or folder of datastores: Datastore > Allocate space	Datastore Consumer or Administrator
	On the network that the virtual machine will be assigned to: Network > Assign network	Network Consumer or Administrator
Take a virtual machine snapshot	On the virtual machine or a folder of virtual machines: Virtual machine > Snapshot management > Create snapshot	Virtual Machine Power User or Administrator
Move a virtual machine into a resource pool	On the virtual machine or folder of virtual machines: Resource > Assign virtual machine to resource pool Virtual machine > Edit inventory > Move	Administrator
	On the destination resource pool: Resource > Assign virtual machine to resource pool	Administrator
Install a guest operating system on a virtual machine	On the virtual machine or folder of virtual machines: Virtual machine > Interaction > Answer question Virtual machine > Interaction > Console interaction Virtual machine > Interaction > Device connection Virtual machine > Interaction > Power Off Virtual machine > Interaction > Power On Virtual machine > Interaction > Reset Virtual machine > Interaction > Configure CD media (if installing from a CD) Virtual machine > Interaction > Configure floppy media (if installing from a floppy disk) Virtual machine > Interaction > VMware Tools install	Virtual Machine Power User or Administrator
	On a datastore that contains the installation media ISO image: Datastore > Browse datastore (if installing from an ISO image on a datastore) On the datastore to which you upload the installation media ISO image: Datastore > Browse datastore Datastore > Low level file operations Host > Configuration > System Management	Virtual Machine Power User or Administrator
Migrate a virtual machine with vMotion	On the virtual machine or folder of virtual machines: Resource > Migrate powered on virtual machine Resource > Assign Virtual Machine to Resource Pool (if destination is a different resource pool from the source)	Resource Pool Administrator or Administrator
	On the destination host, cluster, or resource pool (if different from the source): Resource > Assign virtual machine to resource pool	Resource Pool Administrator or Administrator
Cold migrate (relocate) a virtual machine	On the virtual machine or folder of virtual machines: Resource > Migrate powered off virtual machine Resource > Assign virtual machine to resource pool (if destination is a different resource pool from the source)	Resource Pool Administrator or Administrator

Task	Required Privileges	Applicable Role
	On the destination host, cluster, or resource pool (if different from the source): Resource > Assign virtual machine to resource pool	Resource Pool Administrator or Administrator
	On the destination datastore (if different from the source): Datastore > Allocate space	Datastore Consumer or Administrator
Migrate a virtual machine with Storage vMotion	On the virtual machine or folder of virtual machines: Resource > Migrate powered on virtual machine	Resource Pool Administrator or Administrator
	On the destination datastore: Datastore > Allocate space	Datastore Consumer or Administrator
On the destination cluster: Host > Inventory > Add host to cluster Host > Inventory > Modify cluster	Administrator	Administrator
Add a single host to a data center by using the vSphere Client, or add a single host to a cluster by using PowerCLI or API (leveraging the addHost API)	On the host: Host > Inventory > Add host to cluster	
	On the cluster: Host > Inventory > Modify cluster Host > Inventory > Add host to cluster	
	On the data center: Host > Inventory > Add standalone host	
Add multiple hosts to a cluster	On the cluster: Host > Inventory > Modify cluster Host > Inventory > Add host to cluster	Administrator
	On the parent data center of the cluster (with propagate): Host > Inventory > Add standalone host Host > Inventory > Move host Host > Inventory > Modify cluster Host > Configuration > Maintenance	Administrator
Encrypt a virtual machine		Administrator
Protect a virtual machine (if using vSphere+ to protect the virtual machine)	On the data center in which the virtual machine is deployed: vSphere Tagging > Assign or Unassign vSphere Tag	Administrator

Troubleshooting Overview

vSphere Troubleshooting contains common troubleshooting scenarios and provides solutions for each of these problems. You can also find guidance here for resolving problems that have similar origins. For unique problems, consider developing and adopting a troubleshooting methodology.

The following approach for effective troubleshooting elaborates on how to gather troubleshooting information, such as identifying symptoms and defining the problem space. Troubleshooting with log files is also discussed.

Guidelines for Troubleshooting Your Implementation of vSphere

To troubleshoot your implementation of vSphere, identify the symptoms of the problem, determine which of the components are affected, and test possible solutions.

Identifying Symptoms**Defining the Problem Space****Testing Possible Solutions**

A number of potential causes might lead to the under-performance or nonperformance of your implementation. The first step in efficient troubleshooting is to identify exactly what is going wrong. After you have isolated the symptoms of the problem, you must define the problem space. Identify the software or hardware components that are affected and might be causing the problem and those components that are not involved.

When you know what the symptoms of the problem are and which components are involved, test the solutions systematically until the problem is resolved.

How Do You Identify the Symptoms

Before you attempt to resolve a problem in your vSphere implementation, you must identify precisely how it is failing.

The first step in the troubleshooting process is to gather information that defines the specific symptoms of what is happening. You might ask these questions when gathering this information:

- What is the task or expected behavior that is not occurring?
- Can the affected task be divided into subtasks that you can evaluate separately?
- Is the task ending in an error? Is an error message associated with it?
- Is the task completing but in an unacceptably long time?
- Is the failure consistent or sporadic?
- What has changed recently in the software or hardware that might be related to the failure?

How Do You Define the Problem Space

After you identify the symptoms of the problem in your vSphere implementation, determine which components in your setup are affected, which components might be causing the problem, and which components are not involved.

To define the problem space in an implementation of vSphere, be aware of the components present. In addition to VMware software, consider third-party software in use and which hardware is being used with the VMware virtual hardware.

Recognizing the characteristics of the software and hardware elements and how they can impact the problem, you can explore general problems that might be causing the symptoms.

- Misconfiguration of software settings
- Failure of physical hardware
- Incompatibility of components

Break down the process and consider each piece and the likelihood of its involvement separately. For example, a case that is related to a virtual disk on local storage is probably unrelated to third-party router configuration. However, a local disk controller setting might be contributing to the problem. If a component is unrelated to the specific symptoms, you can probably eliminate it as a candidate for solution testing.

Think about what changed in the configuration recently before the problems started. Look for what is common in the problem. If several problems started at the same time, you can probably trace all the problems to the same cause.

How Do You Test Possible Solutions

After you know the symptoms of the problem in your vSphere implementation and which software or hardware components are most likely involved, you can systematically test solutions until you resolve the problem.

With the information that you have gained about the symptoms and affected components, you can design tests for pinpointing and resolving the problem. These tips might make this process more effective.

- Generate ideas for as many potential solutions as you can.
- Verify that each solution determines unequivocally whether the problem is fixed. Test each potential solution but move on promptly if the fix does not resolve the problem.
- Develop and pursue a hierarchy of potential solutions based on likelihood. Systematically eliminate each potential problem from the most likely to the least likely until the symptoms disappear.
- When testing potential solutions, change only one thing at a time. If your setup works after many things are changed at once, you might not be able to discern which of those things made a difference.
- If the changes that you made for a solution do not help resolve the problem, return the implementation to its previous status. If you do not return the implementation to its previous status, new errors might be introduced.
- Find a similar implementation that is working and test it in parallel with the implementation that is not working properly. Make changes on both systems at the same time until few differences or only one difference remains between them.

Troubleshooting with vCenter Server Logs

You can often obtain valuable troubleshooting information by looking at the logs provided by the various services and agents that your vSphere implementation is using.

Most logs are located in `/var/log/vmware/<service_name>` for vCenter Server deployments.

Common Logs

The following logs are common to all vCenter Server deployments.

Table 176: Common Log Directories

Log Directory	Description
<code>../firstboot</code>	Stores first boot logs
<code>applmgmt</code> and <code>applmgmt-audit</code>	Stores logs related to the VMware Appliance Management Service
<code>cloudvm</code>	Stores logs for allotment and distribution of resources between services
<code>rhttpproxy</code>	Stores logs for the VMware HTTP Reverse Proxy service
<code>sca</code>	Stores logs for the VMware Service Control Agent service
<code>vapi</code>	Stores logs for the VMware vAPI Endpoint service
<code>vmafdd</code>	Stores logs for the VMware Authentication Framework - LDAP service
<code>vmldird</code>	Stores logs for the VMware Directory Service - LDAP service
<code>vmon</code>	Stores logs for the VMware Service Lifecycle Manager service

Management Node Logs

The following logs are available if a management node deployment is selected.

Table 177: Management Node Log Directories

Log Directory	Service
<code>rbd</code>	VMware vSphere Auto Deploy

Log Directory	Service
content-library	VMware Content Library Service
eam	VMware ESX Agent Manager
netdumper	VMware vSphere ESXi Dump Collector
perfcharts	VMware Performance Charts Service
vmcam	VMware vSphere Authentication Proxy
vmdird	VMware Directory Service - LDAP
vmware-sps	VMware vSphere Profile-Driven Storage Service
vpxd	VMware vCenter Server
vpostgres	VMware Postgres service
vcha	VMware vCenter High Availability service

Troubleshooting Virtual Machines

The virtual machine troubleshooting topics provide solutions to potential problems that you might encounter when using your virtual machines.

Troubleshooting USB Passthrough Devices

Information about feature behavior can help you troubleshoot or avoid potential problems when USB devices are connected to a virtual machine.

Error Message When You Try to Migrate Virtual Machine with USB Devices Attached

Migration with vMotion cannot proceed and issues a confusing error message when you connect multiple USB devices from an ESXi host to a virtual machine and one or more devices are not enabled for vMotion.

The Migrate Virtual Machine wizard runs a compatibility check before a migration operation begins. If unsupported USB devices are detected, the compatibility check fails and an error message similar to the following appears: `Currently connected device 'USB 1' uses backing 'path:1/7/1', which is not accessible.`

To successfully pass vMotion compatibility checks, you must enable all USB devices that are connected to the virtual machine from a host for vMotion. If one or more devices are not enabled for vMotion, migration will fail.

1. Make sure that the devices are not in the process of transferring data before removing them.
2. Re-add and enable vMotion for each affected USB device.

Cannot Copy Data From an ESXi Host to a USB Device That Is Connected to the Host

You can connect a USB device to an ESXi host and copy data to the device from the host. For example, you might want to gather the vm-support bundle from the host after the host loses network connectivity. To perform this task, you must stop the USB arbitrator.

If the USB arbitrator is being used for USB passthrough from an ESXi host to a virtual machine the USB device appears under `lsusb` but does not mount correctly.

This problem occurs because the nonbootable USB device is reserved for the virtual machine by default. It does not appear on the host's file system, even though `lsusb` can see the device.

1. Stop the `usbarbitrator` service: `/etc/init.d/usbarbitrator stop`
2. Physically disconnect and reconnect the USB device.
By default, the device location is `/vmfs/devices/disks/mpx.vmhbaXX:C0:T0:L0`.
3. After you reconnect the device, restart the `usbarbitrator` service: `/etc/init.d/usbarbitrator start`
4. Restart `hostd` and any running virtual machines to restore access to the passthrough devices in the virtual machine.

Recover Orphaned Virtual Machines

Virtual machines appear with `(orphaned)` appended to their names.

Virtual machines that reside on an ESXi host that vCenter Server manages might become orphaned in rare cases. Such virtual machines exist in the vCenter Server database, but the ESXi host no longer recognizes them.

Virtual machines can become orphaned if a host failover is unsuccessful, or when the virtual machine is unregistered directly on the host. If this situation occurs, move the orphaned virtual machine to another host in the data center on which the virtual machine files are stored.

1. Determine the datastore where the virtual machine configuration (`.vmx`) file is located.
 - a) Select the virtual machine in the inventory, and click the **Datastores** tab.
The datastore or datastores where the virtual machine files are stored are displayed.
 - b) If more than one datastore is displayed, select each datastore and click the **Files** tab to browse for the `.vmx` file.
 - c) To verify the location of the `.vmx` file, select the virtual machine from **Datastore**.
2. Return to the virtual machine in the inventory, right-click it, and select **Remove from Inventory**.
3. Click **Yes** to confirm the removal of the virtual machine.
4. Reregister the virtual machine with vCenter Server.
 - a) Right-click the datastore where the virtual machine file is located and select **Register VM**.
 - b) Browse to the `.vmx` file and click **OK**.
 - c) Select the location for the virtual machine and click **Next**.
 - d) Select the host on which to run the virtual machine and click **Next**.
 - e) Click **Finish**.

vSphere Host Profiles

Learn how to manage and configure Host Profiles in the vSphere Client.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

The *vSphere Host Profiles* documentation is intended for administrators who are familiar with vSphere host configuration.

What Are vSphere Host Profiles

The vSphere Host Profiles are profiles that encapsulate the host configuration and help to manage the host configuration, especially in environments where an administrator manages multiple hosts or clusters in vCenter Server.

Host Profiles provide an automated and centrally managed mechanism for host configuration and configuration compliance. Host Profiles can improve efficiency by reducing reliance upon repetitive, manual tasks. Host Profiles capture the configuration of a pre-configured and validated reference host, store the configuration as a managed object and use the catalog of parameters contained within to configure networking, storage, security, and other host-level parameters. Host Profiles can be applied to individual hosts, a cluster, or all the hosts and clusters associated to a host profile. Applying a Host Profile to a cluster affects all hosts in the cluster and results in a consistent configuration across the applied hosts.

Host Profiles can be used to validate the configuration of a host by checking compliance of a host or cluster against the Host Profile that is associated with that host or cluster.

Reference Host Independence

A dedicated reference host is not required to be available to perform host profile tasks.

When you create a host profile, you extract the configuration information from a specified ESXi reference host. In previous releases, vSphere required that the reference host was available for certain Host Profiles tasks, such as editing, importing, and exporting. From vSphere 6.0 and later, a dedicated reference host is no longer required to be available to perform these tasks. For host profile tasks that require a reference host, an ESXi host that is compatible to the host profile is assigned as a reference host.

Sometimes, a compatible host is not available to validate the host profile during these tasks. If you make small changes to the host profile, validation is skipped and you see a warning indicating that no valid reference host is associated with the profile. You can then proceed and finish the task.

Due to the introduction of this feature, users can no longer edit or change the reference host from the vSphere Client. The reference host selection occurs at runtime, without notifying users, in the vCenter Server for on-going tasks.

Reference Host Selection

To select a suitable reference host, vCenter Server performs complex version and compatibility checks on the target host.

Starting with vSphere 7.0 Update 3, a major-version vCenter Server instance can manage hosts of all ESXi Update release versions within the respective major version. For example, a vCenter Server 7.0 Update 3 instance can manage ESXi hosts and hosts profiles of version 7.0 GA, 7.0 Update 1, 7.0 Update 2, and 7.0 Update 3. This change affects the reference host selection for particular host profile operations.

Version Check

To select a reference host, vCenter Server checks whether the target host is of the same ESXi Update release version as the host profile. For some operations, the target host can run a later ESXi version than the ESXi version of the host profile. For example, if you want to edit a 7.0 Update 3 host profile, the reference host that validates the host profile cannot be running ESXi 7.0 GA, 7.0 Update 1, or 7.0 Update 2.

Compatibility Check

In addition to checking the version on the target host, vCenter Server performs an extensive compatibility check that includes a subprofile compatibility check in the existing releases. Starting with vSphere 7.0 Update 3, to appoint a reference host, vCenter Server checks the advanced options and the default roles in the target host. A reference host must support the advanced options of the host profile it validates. Also, the default roles in the reference host must match the default roles in the host profile.

Host Profile Operations and Required Reference Host Versions

A reference host is needed to validate the host profile during different host profile operations and workflows. Starting with vSphere 7.0 Update 3, depending on the operation, the reference host must be of exactly the same or later Update release version as the host profile it validates. The following table summarizes the compatible reference host versions for different host profile operations if your environment runs a vCenter Server instance 7.0 Update 3 or later.

Host Profile Operation	Compatible Reference Host Version
Edit	The same Update release version as the host profile.
Check Compliance	The same or later Update release version as the host profile.
Remediation	The same or later Update release version as the host profile.
Attach	The same or later version as the host profile.
Import/Export	The exact same version as the host profile.

Note: If the vCenter Server instance is of version earlier than 7.0 Update 3, Update release versions are not checked during the selection of a reference host. vCenter Server only checks if the reference host and the host profile are of the same major release version, for example 6.x or 7.x.

Host Profiles and vSphere Auto Deploy

You can use Host Profiles and vSphere Auto Deploy to provision physical ESXi hosts with a complete and expected configuration state for virtual switches, driver settings, boot parameters, and so on.

Because hosts that are provisioned with Auto Deploy are considered to be stateless, configuration state information is not stored on the host. Instead, you create and configure a reference host. Then, you create a host profile from the reference host. Next, you associate the host profile with an Auto Deploy rule. As a result new hosts that you provision with Auto Deploy have the host profile applied automatically.

Remediation for these hosts is the same as statefully deployed hosts. When you apply the host profile, you are prompted to customize the hosts and enter answers for the policies that you specified during the host profile creation.

Note: If you deploy ESXi with Auto Deploy, configure a `syslog` to store logs in a remote server. For more information on setting up a syslog from the Host Profiles interface, see the *vSphere Installation and Setup* documentation.

For more information on setting up an Auto Deploy reference host, see the *VMware ESXi Installation and Setup* documentation.

How Do You Manage vSphere Host Profiles in the vSphere Client

By using the vSphere Client, you can perform basic tasks for host profiles, perform and schedule compliance checks of hosts and clusters, remediate hosts which are not compliant with the attached host profiles.

Host Profiles Usage Model

The Host Profiles workflow starts with the concept of a reference host. The configuration of the reference host, which is extracted as a host profile, serves as a configuration template for configuring other hosts. The reference host does not have to be related or associated with the host profile extracted from it.

Before you begin, ensure that you have an existing vSphere environment installation with at least one properly configured ESXi host.

The sequence required to create a host profile from a reference host, apply the host profile to a host or cluster and check compliance against the host profile is as follows:

1. Set up and configure the reference host.
2. Create a host profile from the reference host.
3. Attach hosts or clusters to the host profile.
4. Check the compliance to the host profile. If all hosts are compliant with the reference host, they are correctly configured.
5. Apply (remediate).

As a licensed feature of vSphere, Host Profiles are only available when the appropriate licensing is in place. If you see errors, ensure that you have the appropriate vSphere licensing for your hosts.

If you want the Host Profile to use directory services for authentication, configure the reference host to use a directory service. See the *vSphere Security* documentation.

Hosts Provisioned with vSphere Auto Deploy

For hosts provisioned with vSphere Auto Deploy, the vSphere Client owns the entire host configuration, which is captured in a Host Profile. Usually, the Host Profile information is sufficient to store all configuration information. Sometimes the user is prompted for input when the host provisioned with Auto Deploy starts. See the *VMware ESXi Installation and Setup* documentation for more information on Auto Deploy.

How Do You Get Started with vSphere Host Profiles

Learn how to create a host profile, attach or detach hosts and clusters to a host profile, import or export a host profile, duplicate a host profile, and copy settings from a reference host to a host profile or from one host profile to another.

Note: When vSphere Configuration Profiles is enabled on a cluster, you cannot attach or edit host profiles in the cluster or in the hosts in that cluster. For more information on vSphere Configuration Profiles, see the *Managing Host and Cluster Lifecycle* guide.

How Do You Create a vSphere Host Profile?

You create a host profile by extracting the configuration from a designated reference host.

Verify that you have a working vSphere installation and at least one completely and properly configured host.

Host profiles do not capture offline or non-present devices. Any changes made to offline devices after extracting a host profile do not make a difference to the compliance check results.

Note: When vSphere Configuration Profiles is enabled on a cluster, you cannot attach or edit host profiles in the cluster or in the hosts in that cluster. For more information on vSphere Configuration Profiles, see the *Managing Host and Cluster Lifecycle* guide.

Note:

Host profiles are not supported for ESXi hosts with a Data Processing Unit (DPU).

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Click **Extract Host Profile**.
4. On the **Select host** page, select the vCenter Server instance where to create the host profile.
5. Select the host that acts as the reference host and click **Next**.
The selected host must be a valid host.
6. Enter a name and, optionally, a description for the new profile, and click **Finish**.

The new profile appears in the **Host Profiles** pane.


How Do You Attach or Detach Hosts and Clusters to a vSphere Host Profile?

After you create a host profile from a reference host, you must attach a host or a cluster to the host profile to associate that host profile with the selected host or cluster. To disassociate a configuration from an ESXi host or an entire cluster, detach the host or cluster from the host profile.

When a host profile is attached to a cluster, the host or hosts within that cluster are also attached to the host profile. However, when the host profile is detached from the entire cluster, there is no association between the host or hosts within the cluster and that host profile. Detaching a host profile from an ESXi host or a cluster does not delete that host profile. You can delete the host profile after detaching it from all the entities it is associated with.

Note: When vSphere Configuration Profiles is enabled on a cluster, you cannot attach or edit host profiles in the cluster or in the hosts in that cluster. For more information on vSphere Configuration Profiles, see the *Managing Host and Cluster Lifecycle* guide.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click the host profile and select **Attach/Detach Hosts and Clusters**.



4. Select or deselect the hosts or clusters from the list, and click **Save**.
You can use the **Filter** text box to search the hosts and clusters list.

The hosts or clusters are added or removed from the selected host profile.

How Do You Duplicate a vSphere Host Profile?

A host profile duplicate is a copy of an existing host profile. When you create the duplicate, you enter a new name and description for it and the duplicate appears in the list of host profiles.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click the host profile you want to duplicate and select **Duplicate Host Profile**.
4. Enter a different name and description for the duplicate host profile, and click **OK**.

A clone of the host profile appears in the **Host Profiles** list.

How Do You Copy Settings from a Reference Host to a vSphere Host Profile?

If the configuration of the reference host changes, you can update the host profile so that it matches the reference host's new configuration. You can make changes to the reference host and then copy the changes to the host profile. Alternatively, you can edit the host profile to match the new reference host configuration.

After you create a host profile, you can make incremental updates to the profile. When making changes to a host profile, consider the benefits and limitations of the two methods:

- Make the configuration changes to a host and copy that host's settings to the profile. The settings within the existing profile are updated to match those of the host. This method allows you to validate the configuration on a single host before rolling it to the other hosts that are attached to the profile.
- Update the profile directly by editing the host profile. This provides the ability to do more comprehensive and immediate remediation of those changes.

Note: Fixed user password, system image cache and some of the host customized settings are not present in the newly updated host profile. Edit the host profile to update these settings.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click the host profile and click **Copy Settings from Host**.
4. Select the host from which you want to copy the configuration settings and click **OK**.

How Do You Import or Export a vSphere Host Profile?

You can import a host profile from a file that is in the VMware profile format (.vpf). You can also export a host profile to a file that is in the VMware profile format (.vpf).

When you export a host profile, administrator and user profile passwords are not exported. This is a security measure and stops passwords from being exported in plain text when the profile is exported. You are prompted to re-enter the values for the password after the profile is imported and the password is applied to a host.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. If you want to import a host profile, click **Import Host Profile**.
 - a) To locate the VMware Profile Format file (.vpf) that you want to import, click **Browse**.
 - b) Enter the **Name** and **Description** for the imported Host Profile, and click **OK**.
The imported host profile appears in the **Host Profiles** list.
4. If you want to export a host profile, right-click the profile, select **Export Host Profile**, and click **Save**.

How Do You Copy Settings from a vSphere Host Profile?

After you make changes to a host profile, you can propagate those changes to other host profiles in the inventory.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click the host profile from which you want to copy settings and select **Copy Settings to Host Profiles**.
4. Select the settings you want to copy to other host profiles, and click **Next**.
5. Select the target host profile that will be overwritten with the selected settings, and click **Next**.
The differences between the host profile settings are displayed on the **Ready to complete** page.

6. Click **Finish**.

How Do You Manage Compliance Checks of Hosts and Clusters to a vSphere Host Profile

You can check the compliance of hosts and clusters to a vSphere host profile to verify if the host configuration matches the configuration of the host profile. Learn how to check compliance manually or schedule regular compliance checks of hosts and clusters, and how to remediate hosts that are not compliant with the host profile configuration.

How Do You Check Compliance of Hosts or Clusters to a vSphere Host Profile?

Learn how to check the compliance of a host or cluster to the host profile attached to them. By using the compliance check, you can see if any configuration parameters on a host are different from those specified in the host profile.

As a result of the compliance check, the host status is labeled as Compliant, Unknown, or Non-compliant.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click a host profile and click **Check Host Profile Compliance**.

Note: Host profiles do not capture offline or unpresented devices. Any changes made to offline devices after extracting a host profile do not make a difference to the compliance check results.

The compliance status is updated as Compliant, Unknown, or Non-compliant.

A non-compliant status indicates a discovered and specific inconsistency between the profile and the host. To resolve this, you should remediate the host. Any unknown status indicates that the compliance of the host cannot be verified. To resolve the issue, remediate the host through the host profile. Often the compliance check fails because the host is disconnected.

During a compliance check of an ESXi host to the host profile, two groups of advanced option parameters are not managed by the host profile and their values remain unchanged for the host:

- advanced option parameters which are disabled in the host profile.
- advanced option parameters which exist in the ESXi host you want to remediate, and do not exist in the host profile.

To see more details on compliance failures, select a host profile from the **Host Profiles** main view for which the last compliance check displayed one or more failures. To see which parameters differ between the host that failed compliance and the host profile, click the **Monitor** tab and select the **Compliance** view. Then select the failing host. The differing parameters are displayed in the Compliance window, below the host list.

How Do You Schedule a Compliance Check of Hosts or Clusters to a vSphere Host Profile?

You can schedule a regular compliance check of a host or cluster, using a standard vSphere Client scheduled tasks workflow. This automation determines whether any configuration parameters are different from those specified in the host profile.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Select the desired host profile and navigate to the **Monitor** tab.
4. From the **Monitor** tab, select **Scheduled Tasks -> New Scheduled Task -> Check Compliance**.

The **Schedule New Tasks (Check Compliance)** dialog box appears.

5. Optional: Enter a new task name.
Additional details about the task can be entered in the **Description** text box.
6. Choose a frequency how often to run the task.
You can postpone when the task is run for the first time and schedule when it ends.
7. Optional: To receive a notification after the task is finished, enter an email.
8. Click **Schedule The Task**.

The newly created task is visible on the **Scheduled Tasks** list. You can **Edit**, **Run** or **Remove** the task.

To understand more about compliance status, checks or failures, see **Check Compliance**.

How Do You Remediate a Host to Comply with a vSphere Host Profile Configuration?

When you check the compliance of a host to a host profile and a host is non-compliant, you can use the remediate option to apply the host profile settings to the host. During remediation, all host profile parameters are updated to the values defined in the host profile that is attached to the host.

Verify that a host profile is attached to the host.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click the host profile and select **Remediate**.

Note: Certain Host Profile policy configurations require that the host is rebooted after remediation. In such cases, you are prompted to place the host into maintenance mode. You might be required to place hosts into maintenance mode before remediation. Hosts that are in a fully automated DRS cluster are placed into maintenance mode during remediation. Alternatively, if the host is not placed into maintenance mode when required, the remediation process stops.

4. Select the hosts you want to remediate against the host profile.

Note:

When you remediate an ESXi host to comply with a host profile, two groups of advanced option parameters are not managed by the host profile and their values remain unchanged for the host:

- advanced option parameters which are disabled in the host profile.
- advanced option parameters which exist in the ESXi host you want to remediate, and do not exist in the host profile.

The remediation pre-check results do not contain tasks related to these two groups of advanced option parameters.

5. Optional: To specify host properties, enter your host customizations or import a host customization file.
You can modify the user input parameters for the host profiles policies by customizing the host and clicking **Next**. See [Host Profiles and vSphere Auto Deploy](#) for more information about vSphere Auto Deploy.
6. To check if the selected hosts can be remediated, click **Pre-check Remediation**.
The pre-check generates a list of tasks that will be performed on the host.
7. If needed, select the check box to reboot the hosts.
Alternatively you can reboot the hosts manually after the remediation process.
8. Review the tasks that are necessary to remediate the host profile and click **Finish**.

The compliance status is updated.

How Do You Configure vSphere Host Profiles

Learn how to edit host profile policies and components, and configure security and networking host profiles.

How Do You Manage Host Profile Policies and Policy Components

In the vSphere Client, you can view and edit host profiles policies, policy options, and parameters. You can change policy names or policy descriptions. If you do not want a host profile component to be considered during a compliance check, you can deactivate that specific component.

Edit a Host Profile

You can view and edit Host Profiles policies, select a policy to be checked for compliance, and change the policy name or description.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click the host profile that you want to edit and click **Edit Host Profile**.
4. Optional: Click the **Name and Description** tab to change the profile name and description.
5. In the **Settings** tab, expand each category to view or edit particular policies or settings.

Note: See [How Do You Edit a vSphere Host Profile Policy?](#) for detailed instructions for editing a Host Profile policy. See [How Do You Deactivate a Host Profile Component or Subprofile?](#) for detailed instructions on enabling or disabling a policy from compliance check or remediation.

6. View **All** host profile configurations or only **Favorites** configurations.
You can mark a configuration as a favorite with the star icon. Those are then added to **Favorites** configurations.
7. Optional: In the search field, filter the configuration names and values you want to view.
For example, enter **SNMP**. All configurations that relate to **SNMP** are displayed.
8. Optional: Customize the hosts.
Make any changes to the available configuration values for this profile and click **Save**.

Note: The host customization settings page only appears if you changed any settings that require host customizations.

The changes are made when the "Update Host Profile" task is completed in the Recent Tasks status. If you attempt to remediate the profile before the task is complete, the profile configuration does not contain the change.

How Do You Edit a vSphere Host Profile Policy?

A host profile policy describes how a specific configuration setting is applied. You can edit policies belonging to a specific host profile.

To see the sub-profile components that comprise the host profile, you can expand the host profile's configuration hierarchy when you edit the host profile. These components are categorized by functional group or resource class to make it easier to find a particular parameter. Each subprofile component contains one or more attributes and parameters, with the policies and compliance checks.

Each policy consists of one or more options that contain one or more parameters. Each parameter consists of a key and a value. The value can be one of a few basic types, for example, integer, string, string array, or integer array.

Note: Currently, there is no way to remove or replace policy options, policies, or sub-profiles that are deprecated in this release. To allow old host profiles to continue working, metadata is added to the deprecated policies. New host profiles are extracted with only non-deprecated parts of a host profile.

Table 178: Subset of Host Profile Subprofile Configurations

Component Categories	Configuration Settings	Notes and Examples
Advanced Configuration Settings	Advanced Options, Agent VM, DirectPath I/O, Hosts file, Power System, System Image Cache.	- Host profiles do not check advanced settings if they are the same as the default settings. vCenter Server copies only the advanced configuration settings that have changed and that differ from the default values. In addition, compliance checks are limited to the settings that are copied. - Host profiles do not support the configuration of PCI devices for virtual machine passthrough on the ESXi host.
General System Settings	Console, Core Dump, Device Alias, Host Cache, Kernel Module, Management Agent, System Resource Pool, System Swap, vFlash Host Swap Cache, CIM-XML Indication Subscriptions.	For date and time configuration: - For the time zone, enter a UTC string. For example, "America/Los_Angeles" for United States Pacific time zone. - The default time zone is set to the local time and location of the vSphere Client machine. - Configure the Network Time Protocol (NTP) correctly. You can configure the NTP settings on the host Configure tab. Click Time Configuration (under System). To configure the time settings, click Edit. Note: Enter the host cache size as a whole GB in MB. For example, 1024MB.
Networking	vSwitch, Port groups, Physical NIC speed, security and NIC teaming policies, vSphere Distributed Switch, and vSphere Distributed Switch uplink port.	When DHCPv6 is enabled in the networking subprofile, manually turn on the corresponding rule set in the firewall subprofile.
Security	Firewall, Security Settings, Service.	
Storage	Configure storage options, including Native Multi-Pathing (NMP), Pluggable Storage Architecture (PSA), iSCSI adapters, and NFS storage.	- Use the vSphere CLI to configure or modify the NMP and PSA policies on a reference host, and then extract the Host Profile from that host. If you use the Profile Editor to edit the policies, to avoid compliance failures, make sure that you understand interrelationships between the NMP and PSA policies and the consequences of changing individual policies. For information about the NMP and PSA, see the <i>vSphere Storage</i> documentation. - Add the rules that change device attributes before extracting the Host Profile from the reference host. After attaching a host to the Host Profile, if you edit the profile and change the device attributes (for example, mask device paths or adding SATP rules to mark the device as SSD) you are prompted to reboot the host to make the changes. However, after rebooting, compliance

Component Categories	Configuration Settings	Notes and Examples
		failures occur because the attributes changed. Because Host Profiles extract device attributes before rebooting, if any changes occur after the reboot, it evaluates and finds those changes, and reports it as noncompliant. - To configure or modify the SatpDeviceProfile policy after extracting the host profile, use the vSphere Client. For compliance purposes, the policy option strings must be in the following format: - For an ALUA supported array, for example, SATP_ALUA, the policy options must be separated by a semicolon (;). For example: <code>implicit_support=<on/off>; explicit_support=<on/off>; action_onRetryErrors=<on/off></code> - For an ALUA supported array with CX, for example, SATP_ALUA_CX, the policy options must be separated by a semicolon (;). For example: <code>navireg=<on/off>; implicit_support=<on/off>; action_onRetryErrors=<on/off></code> - For a CX array, for example, SATP_CX or SATP_INV, the policy options must be separated by a space. For example: <code>navireg=<on/off> ipfilter=<on/off> action_onRetryErrors=<on/off></code> Note: The policy configuration options that are marked with <code>off</code> are not present in the configuration string.

Other profile configuration categories include: user group, authentication, kernel module, DCUI keyboard, host cache settings, SFCB, resource pools, login banner, SNMP agent, power system, and CIM indication subscriptions.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Select the host profile that you want to edit and click the **Configure** tab.
4. Click **Edit Host Profile**.
5. Expand a subprofile until you reach the policy to edit.
6. Select the policy.
The policy options and parameters appear on the right side of the **Edit Host Profile** dialog box.

7. Edit the policy options or parameters.

How Do You Deactivate a Host Profile Component or Subprofile?

You can decide whether a host profile component or a subprofile is applied or considered during a compliance check. This action allows administrators to eliminate non-critical attributes from consideration during the compliance check or ignore values that, while part of the host profile, are likely to vary between hosts.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Select the host profile that you want to edit and click the **Configure** tab.
4. Click **Edit Host Profile**.
5. Expand the Host Profile Component hierarchy until you reach the desired component or component element.
6. Deselect the check box next to a component.

The check box is selected by default. If you deselect the check box, the other subprofiles that are enabled are still applied and selected.

During a host profile compliance check or host remediation, the component or component element is not considered.

Note: Sometimes, deselecting the check box might remove the component or component element from the host. This action is displayed in the task list after the pre-check remediation.

How Do You Configure a Security vSphere Host Profile

Learn how to manage role profiles, user account profiles, lockdown mode profiles, and active directory permission profiles that are grouped as part of the security host profile. Starting with vSphere 8.0, users with access to the ESXi Shell and the Administrator role can remove or grant ESXi Shell access to user accounts.

Make sure that you have the SecurityConfigProfile plugin available to validate the role, user account, and active directory permission profiles as there are dependencies between them.

You can configure the host profile options, part of the security profile.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Select the host profile that you want to edit and click the **Configure** tab.
4. Click **Edit Host Profile**.
5. Unfold the **Security and Services > Security Settings** profile category and open the **Security** folder.

The following profiles are displayed:

Role	This profile allows you to view default roles and add custom roles within the ESXi system.
User Configuration	This profile allows you to create and manage user accounts. By default, newly created user accounts have access to the ESXi Shell. Starting with vSphere 8.0, administrators can use the advanced option <code>Security.DefaultShellAccess</code> to change this default behavior. Here are some of the operations that you can perform for user accounts: • Create a user account. • Configure the password for a user account.

	• Configure the password for the root user. • Activate or deactivate user account ESXi Shell access. Note: Do not modify the shell access property for the root user account. • Configure the role for any user that is not the default one. • Assign a default or a custom role (configure permissions) for a local account. • Configure the SSH key for any user.
Active Directory Permission	This profile allows you to manage permissions for active directory users or groups. For example, you can create permissions that associate an active directory user or a group with a role. When an ESXi host joins the active directory domain, an Admin permission is created for the DOMAIN group ESX Admins. Also, when an active directory user or group is given some permissions on the ESXi host, a corresponding permission is created on that host. The Active Directory Permission profile captures that permission.
Lockdown Mode	This profile allows you to increase the security of your ESXi hosts by restricting user permissions and privileges. You can configure the following lockdown mode settings: • Normal lockdown mode: An ESXi host can be accessed from a local console and vCenter Server. DCUI service is not stopped. • Strict lockdown mode: An ESXi host can be only accessed from vCenter Server. The DCUI service is stopped. • Exception users: The users that still have their permissions regardless of the lockdown mode state.

For more information on the security profile, see the *vSphere Security* documentation.

How Do You Configure a Networking vSphere Host Profile

By default, hosts provisioned with vSphere Auto Deploy are assigned DHCP addresses by a DHCP server. You can use the vSphere Auto Deploy host customization mechanism to assign static IP addresses to hosts.

- Set up your vSphere Auto Deploy environment.

- Boot the host using vSphere Auto Deploy.
 - Extract a host profile from the host.
1. In the vSphere Client, select **Menu > Policies and Profiles**.
 2. Under **Policies and Profiles**, click **Host Profiles**.
 3. Select the host profile that you want to edit and click the **Configure** tab.
 4. Click **Edit Host Profile**.
 5. Change the default IP address settings by locating **Networking configuration > Host port group > Management Network > IP address settings**.
 6. From the **IPv4 address** drop-down menu, select **User specified IP address to be used while applying the configuration**.
 7. If the host is in a different subnet than the vCenter Server system, select **Networking Configuration > NetStack Instance > defaultTcpipStack > DNS configuration** and enter the default route in the **Default IPv4 gateway** text box.
 8. Select **Networking Configuration > NetStack Instance > defaultTcpipStack > DNS configuration**.
 9. Make sure the **Flag indicating if DHCP should be used** check box is deselected.
 10. Right-click the host and select **All vCenter Actions > Host Profiles > Attach Host Profile**.
 11. Select the profile to attach, and click **Next**.
 12. Provide the IP address and net mask, and click **Finish**.
 13. Reboot the ESXi host.

The IP address is saved as a host customization and applied to the host.

How Do You Customize Hosts by Using vSphere Host Profiles

To customize hosts with shared attributes, you can create a host profile in a reference host. To customize individual hosts, you can set up some fields in the host profile to prompt the user for input for each host. You can also edit and export host customizations.

Host Customization

Host profiles allow you to prespecify information, for example, the storage setup or Syslog setup in a reference host to and apply the host profile to a set of target hosts that share the same settings. You can also use host profiles to specify that certain settings are host dependent. If you do so, the host comes up in maintenance mode when you provision it with Auto Deploy. Remediate the host or reset the host customization to be prompted for input. The system stores your input and uses it the next time the host boots.

When the host profile is set to prompt for user input, you must specify a value in the dialog box that appears when you reset the host customization. An error results if you do not specify a value.

Table 179: Host Profile Options that Prompt for iSCSI User Input

Information to Request User Input For	Setting the Host Profile Option
When you apply a host profile on a system that includes a profile for iSCSI, you are prompted for several properties. For many of the properties, a system default is available. For some properties, you must specify a value or an error results.	1. Select Edit Host Profile, click Storage configuration, and click iSCSI Initiator Configuration. 2. Select the folder for an already enabled initiator and set up the initiator.

Information to Request User Input For		Setting the Host Profile Option
IQN name	If the iSCSI setup uses an IQN name, you are prompted when you apply the host profile. You cannot continue until you provide the name.	3. Set up the initiator. For many fields, the user is prompted as part of host customization.
CHAP information	If you set up iSCSI to require CHAP authentication, you are prompted for CHAP information including the user name and the secret when you apply the host profile. You cannot continue until you provide the name.	

Table 180: Host Profile Options that Prompt for Storage User Input

Information to Request User Input For		Setting the Host Profile Option
You are setting up the Fixed PSP configuration and want to prompt for the adapter and target IDs for the storage arrays that should use the Fixed PSP.		You can set the option only if the adapter is set up to use the Fixed PSP. 1. Select Edit Host Profile, click Storage configuration. 2. Click Native Multipathing (NMP). 3. Click Path Selection Policy (PSP) configuration. 4. In the Preferred Path window, select Prompt the user for adapter and target IDs on the host.

Table 181: Host Profile Options that Prompt for Security User Input

Information to Request User Input For		Setting the Host Profile Option
Administrator password for ESXi host when the host boots for the first time.		1. Select Edit Host Profile, and click Security and Services. 2. click Security Settings and click Security configuration. 3. In the right panel, select User Input Password to be Used to Configure Administrator Password from the Administrator password drop-down menu.
Preconfigures a user for the ESXi host but prompts for the password for that user on each host when the host boots for the first time.		You can perform this task only if a user configuration already exists. Configure the user by selecting one of the options. • Assigned fixed user configurations is available for compatibility with ESX/ESXi 4.1 system, this option displays the password in the clear. • Assign advanced fixed user configurations is for users of ESXi 5.0 and later systems. • Specify the user configuration in the profile but prompt for password during host configuration allows you to specify the information about the user but prompt for a password on each host.

Information to Request User Input For	Setting the Host Profile Option
Prompt the user for credentials when the host joins the Active Directory domain.	- Set the Authentication configuration profile to use a fixed domain. - Select Edit Host Profile, click Security and Services. - Click Security Settings, and click Authentication configuration. - Click Active Directory configuration. - In the Domain Name drop-down menu, select Configure a fixed domain name. - Set the method for joining the domain to prompt the user. - Select Edit Host Profile, click Security and Services and click Authentication configuration. - Click Active Directory configuration. - In the Join Domain Method drop-down menu, select Use user specified AD credentials to join the host to domain.

Table 182: Host Profile Options that Prompt for Networking User Input

Information to Request User Input For	Setting the Host Profile Option
Prompt the user for the MAC address for a port group. You can have the system prompt the user in all cases (User specified MAC address...) or prompt the user only if no default is available.	- Select Edit Host Profile, click Networking configuration, and click Host port group. - Click Management Network. - In the Determine how MAC address for vmknic should be decided field, select how the system manages the MAC address. User specified MAC Address to be used while applying the configuration Prompt the user for the MAC Address if no default is available
Prompt the user for the IPv4 address for each ESXi host to which the profile is applied. You can have the system prompt the user in all cases (User specified IPv4 address...) or prompt the user only if no default is available.	- Select Edit Host Profile, click Networking configuration, and click Host port group. - Click Management Network and click IP address settings. - In the IPv4 address field, select how the system manages the IPv4 address. User specified IPv4 Address to be used while applying the configuration Prompt the user for the IPv4 Address if no default is available
Prompt the user for the IPv6 address for each ESXi host to which the profile is applied. You can have the system prompt the user in all cases (User specified IPv6 address...) or prompt the user only if no default is available.	- Select Edit Host Profile, click Networking configuration, and click Host port group. - Click Management Network and click IP address settings. - In the Static IPv6 address field, select how the system manages the IPv6 address. User specified IPv6 Address to be used while applying the configuration Prompt the user for the IPv6 Address if no default is available
Prompt the user for the DNS name of the host. You can have the system prompt the user in all cases (User specified host name...) or prompt the user only if no default is available.	Select Edit Host Profile, click Networking configuration, and click DNS configuration.

Information to Request User Input For	Setting the Host Profile Option
	- In the Host name field, select how the system manages the DNS configuration. Prompt the user for host name if default is not available User specified host name to be used while applying the configuration
Prompt the user for the MAC address for a distributed switch, its port group, or one of its services. Right-click the Host virtual NIC folder icon and click the Add sub-profile icon to determine the component to which the setting is applied. You can decide to prompt the user in all cases or only if no default is available.	- Open Networking configuration. - Click Host virtual NIC. - In the Determine how MAC address for vmknics should be decided field, select how the system manages the MAC address for the distributed switch. User specified MAC address to be used while applying the configuration Prompt the user for the MAC address if no default is available
Prompt the user for the IPv4 address for a distributed switch, its port group, or one of its services. Right-click the Host virtual NIC folder icon and click the Add sub-profile icon to determine the component to which the setting is applied. You can decide to prompt the user only if no default is available or in all cases.	- Open Networking configuration. - Click Host virtual NIC. - Click IP address settings. - In the IPv4 address field, select how the system handles the IPv4 address for the distributed switch. User specified IPv4 address to be used while applying the configuration Prompt the user for IPv4 address if no default is available
Prompt the user for the IPv6 address for a distributed switch, its port group, or one of its services. Right-click the Host virtual NIC folder icon and click the Add sub-profile icon to determine the component to which the setting is applied. You can decide to prompt the user only if no default is available or in all cases.	- Open Networking configuration. - Open Host virtual NIC. - Open IP address settings. - In the Static IPv6 address field, select how the system manages the IPv6 address for the distributed switch. User specified IPv6 address to be used while applying the configuration Prompt the user for IPv6 address if no default is available

How Do You Export Host Customizations?

Learn how to export the customized attributes of a vSphere host profile as a **.CSV** file to your desktop.

For security reasons, sensitive data such as passwords are not exported.

- Navigate to **Hosts and Clusters** and select an ESXi host with an attached host profile.
- Right-click the ESXi host and select **Host Profiles > Export Host Customizations**.
- Click **Save**.

Note: Only an English version of the **.CSV** file is supported.

The file is saved as a **.CSV** file.

After the file is saved to your desktop, you can manually edit the file and save it to apply the customizations later.

How Do You Edit Host Customizations?

Read how to edit host customizations for specific hosts attached to a vSphere host profile or a cluster attached to a host profile.

1. In the vSphere Client, select **Menu > Policies and Profiles**.
2. Under **Policies and Profiles**, click **Host Profiles**.
3. Right-click the host profile and select **Edit Host Customizations**.
4. Select the host or hosts for which you want to edit the customization, and click **Next**.
5. Optional: To import host customization from your desktop, you can select **Import Host Customization .csv** in the **Customize hosts** page.

If you want to edit host customizations by using a .csv file, you can export the host customizations in .csv file format, edit the values in the answer file fields, and then import the file. For more information on how to export host customizations, see [How Do You Export Host Customizations?](#)

Note: User input in the **Customize hosts** pane is overwritten by the import process and the values from the .csv file.

After importing the .csv file, the text boxes are updated with the information from the file.

6. Edit the host configuration values and click **Finish**.

Answer File Field and vSphere Host Profile Extraction

When a new profile is extracted from an ESXi host or post upgrade, host profiles policy options and host customization values are changed. You can get more familiar with the changes by referring to the Answer File Field and Host Profile Extraction table.

After you extract a new host profile from an ESXi host, the values for some of the policy options change. To keep these values, you have to save them manually before extracting the host profile. The *Answer File Field and Host Profile Extraction* table lists the policy options with changed values and the options that keep their values the same in the newly extracted host profile.

To locate the policy options, navigate to **Host Profiles** main view, right-click the extracted host profile, and select **Edit Host Customizations**.

Table 183: Answer File Field and Host Profile Extraction

Policy option path	Policy option	Variations in the policy options value in the newly extracted profile
Security and Services -> Security Settings -> Security	UserInputPasswordConfigOption	For default users, it leaves the password unchanged (root). For non-default users, it always extracts user input.
	UserInputPasswordConfigOption	If host profile version is earlier than 6.5, this option is not extracted. You explicitly need to create a UserAccount profile in the extracted host profile.
	UserInputAdminPasswordOption	Changes to fixed. Does not keep user input value.
Storage configuration -> iSCSI Initiator Configuration	AdvancedUserInputUserConfigOption	Changes to fixed. Does not keep user input value.
	UserInputIqn	Keeps user input value.
	UserInputAlias	Keeps user input value.
	UserInputChapName	Changes to fixed. Does not keep user input value.
	UserInputChapSecret	Changes to fixed. Does not keep user input value.

Policy option path	Policy option	Variations in the policy options value in the newly extracted profile
General System Settings -> Management Agent Configuration -> SNMP Agent Configuration	UserInputFirstBurstLength	Changes to fixed. Does not keep user input value.
	UserInputMaxBurstLength	Changes to fixed. Does not keep user input value.
	UserInputMaxReceiveSegmentLength	Changes to fixed. Does not keep user input value.
	UserInputUsers	Changes to fixed. Does not keep user input value.
	UserInputV3Targets	Changes to fixed. Does not keep user input value.
	UserInputEngineId	Changes to fixed. Does not keep user input value.
Security and Services - > Security Settings -> Security -> Active Directory Permission	UserInputContact	Changes to fixed. Does not keep user input value.
	UserInputLocation	Changes to fixed. Does not keep user input value.
Networking Configuration	UserInputUserName	Keeps user input value.
	UserInputPassword	Keeps user input value.
	UserInputMacAddress	Conditional, prompt user if no default value is available.
Advanced Configuration Settings -> Advanced Options	UserInputHostName	Conditional, prompt user if no default value is available.
	UserInputIPAddress	Conditional, prompt user if no default value is available.
	UserInputAdvancedConfigOption	Changes to fixed. Does not keep user input value.
Advanced Configuration Settings -> System Image Cache Configuration	Enable Stateless Caching	You must explicitly select the policy option.
	Enable Stateful Install	You must explicitly select the policy option.

Note: This table does not list all policy options. You must be aware of the policy options you modified before extracting a new host profile.

Recommended Upgrade Workflows for vSphere Host Profiles

Each workflow is the recommended upgrade path while upgrading your vCenter Server 6.5, 6.7 and 7.0 environment, that contains one or more host profiles with version 6.5, 6.7 and 7.0.

Upgrade Support Matrix

When you upgrade your vCenter Server from version 6.7 or 7.0 to version 8.0, you can keep your current ESXi host version and host profiles version. To ensure the best possible performance, when you upgrade the ESXi host, also upgrade the host profile. The host profile version must match the ESXi host version.

Note: vCenter Server instances and ESXi hosts of versions 6.5 or earlier are not supported in vCenter Server 8.0. You can use the *Recommended Host Profiles Upgrade Workflows* in the *vSphere Host Profiles 6.7* guide as intermediate upgrade steps.

Table 184: Upgrade Support Matrix

Upgrade Workflow	vCenter Server	ESXi host	Host Profiles
earlier than 6.5 → 7.0	Requires an upgrade	Requires an upgrade	Requires an upgrade

Upgrade Workflow	vCenter Server	ESXi host	Host Profiles
Note: You must first upgrade your vCenter Server and ESXi hosts to versions 6.5 or 6.7. You can use the <i>Recommended Host Profiles Upgrade Workflows</i> in the <i>vSphere Host Profiles 6.7</i> guide as intermediate upgrade steps.			
6.5 → 7.0	Requires an upgrade	An optional upgrade	An optional upgrade
6.7 → 7.0	Requires an upgrade	An optional upgrade	An optional upgrade
6.5/6.7/7.0 → 7.0 Update 3 Note: If you want to validate or edit a 7.0 GA host profile, you might need to keep an ESXi host of version 7.0 GA.	Requires an upgrade	An optional upgrade	An optional upgrade
6.5 → 8.0 Note: You must first upgrade your vCenter Server and ESXi hosts to versions 6.7 or 7.0. You can use the <i>Recommended Host Profiles Upgrade Workflows</i> in the <i>vSphere Host Profiles 7.0</i> guide as intermediate upgrade steps.	Requires an upgrade	Requires an upgrade	Requires an upgrade
6.7/7.0 → 8.0	Requires an upgrade	An optional upgrade	An optional upgrade

vCenter Server and ESXi Version Interoperability

Starting with vSphere 7.0 Update 3, a major-version vCenter Server instance supports any ESXi Update release version of the corresponding major release. That is, a vCenter Server 8.0 instance can manage ESXi 6.7, ESXi 7.0, ESXi 8.0, and any Update release of ESXi 8.0.

For more information about vCenter Server and ESXi Version Interoperability, you can refer to the [VMware Product Interoperability Matrix](#).

How Do You Upgrade Host Profiles and vCenter Server from 6.5 to 8.0

Use the upgrade workflows to resolve version compliance errors when upgrading vCenter Server with version 6.5 to version 8.0.

ESXi hosts of versions 5.5 and 6.0 are not supported in vCenter Server 8.0. You can use the *Recommended Host Profiles Upgrade Workflows* in the *vSphere Host Profiles 6.7* guide as intermediate upgrade steps.

How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi 6.5 Hosts Only?

You can use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 6.5 to version 8.0, if your cluster contains stateful ESXi 6.5 hosts.

- Verify that your cluster contains ESXi 6.5 hosts only.
- Verify that a host profile with version 6.5 is attached to the cluster.
- Verify that your vCenter Server is with version 6.5.

1. Upgrade all ESXi hosts in the cluster to version 7.0.

For more information, see the *VMware ESXi Upgrade* documentation.

2. Remediate the upgraded hosts against the host profile with version 6.5.
During remediation, all host profile settings are applied to the hosts in the cluster.
3. Extract a new host profile from an ESXi 7.0 host in the cluster.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
4. Attach the host profile with version 7.0 to the cluster.
5. Remediate the cluster against the host profile with version 7.0.
6. Upgrade your vCenter Server instance to version 8.0.
The vCenter Server configuration is not changed during the upgrade.
For more information, see the *vCenter Server Upgrade* documentation.
7. Upgrade all ESXi hosts in the cluster to version 8.0.
Edit host customization and host profile edit operations are not available, while compliance check, attach host profile, and remediate hosts operations are available. For more information, see [KB 2150534](#)
8. Remediate the cluster against the host profile with version 7.0.
During remediation, all host profile settings are applied.
9. Extract a new host profile from an ESXi8.0 host.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
10. Attach the host profile with version 8.0 to the cluster.
Edit host customization to validate host customization data.
All host profile operations become available.

How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateless ESXi 6.5 Hosts Only?

You can use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 6.5 to version 8.0, if your cluster contains stateless ESXi 6.5 hosts.

- Verify that your cluster contains ESXi 6.5 hosts only.
 - Verify that a host profile with version 6.5 is attached to the cluster.
 - Verify that your vCenter Server is with version 6.5.
1. Upgrade your vCenter Server instance to version 7.0.
The vCenter Server configuration is not changed during the upgrade.
For more information, see the *vCenter Server Upgrade* documentation.
 2. Create a new rule or edit the existing rule with an image profile version 7.0 and select the cluster.
 3. Remediate all ESXi hosts in the cluster by clicking **Remediate Host Associations**.
 4. Reboot all ESXi hosts in the cluster.
All hosts are compliant with the host profile version 6.5, but the host profile cannot be edited. Remediation and check compliance operations are available.

5. To be able to customize your host profile, follow these steps.
 - a) Extract a new host profile from an ESXi 7.0 host.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
 - b) Attach the host profile with version 7.0 to the cluster.
Edit host customization to validate host customization data.
All host profile operations become available.
6. Upgrade your vCenter Server instance to version 8.0.
The vCenter Server configuration is not changed during the upgrade.
For more information, see the *vCenter Server Upgrade* documentation.
7. Create a new rule or edit the existing rule with an image profile version 8.0 and select the cluster.
8. Remediate all ESXi hosts in the cluster by clicking **Remediate Host Associations**.
9. Reboot all ESXi hosts in the cluster.
All hosts are compliant with the host profile version 6.5, but the host profile cannot be edited. Remediation and check compliance operations are available.
10. To be able to customize your host profile, follow these steps.
 - a) Extract a new host profile from an ESXi 8.0 host.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
 - b) Attach the host profile with version 8.0 to the cluster.
Edit host customization to validate host customization data.
All host profile operations become available.

How Do You Upgrade vSphere Host Profiles and vCenter Server from 6.7 to 8.0

Use the upgrade workflows to resolve version compliance errors when upgrading vCenter Server with version 6.7 to version 8.0.

How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi6.5 and 6.7 Hosts?

You use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 6.7 to version 8.0, if your cluster contains stateful ESXi6.5 and 6.7 hosts.

Prerequisites

- Verify that your cluster contains ESXi6.5 and ESXi6.7 hosts.
- Verify that a host profile with version 6.5 is attached to the cluster.
- Verify that your vCenter Server is with version 6.7.
- Upgrade all ESXi hosts and host profiles of version 6.5 to version 6.7.
You can use the *Recommended Host Profiles Upgrade Workflows* in the *vSphere Host Profiles 6.7* guide as intermediate upgrade steps.

Procedure

1. Upgrade your vCenter Server instance to version 8.0.

The vCenter Server configuration is not changed during the upgrade.

For more information, see the *vCenter Server Upgrade* documentation.

2. Upgrade all ESXi hosts in the cluster to version 8.0.

Edit host customization and host profile edit operations are not available, while compliance check, attach host profile, and remediate hosts operations are available. For more information, see [KB 2150534](#)

3. Remediate the cluster against the host profile with version 6.7.

During remediation, all host profile settings are applied.

4. Extract a new host profile from an ESXi8.0 host.

During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).

5. Attach the host profile with version 8.0 to the cluster.

The host customization data updates automatically.

All host profile operations become available.

How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi6.7 Hosts Only?

You use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 6.7 to version 8.0, if your cluster contains stateful ESXi6.7 hosts.

Prerequisites

- Verify that your cluster contains ESXi6.7 hosts only.
- Verify that a host profile with version 6.7 is attached to the cluster.
- Verify that your vCenter Server is with version 6.7.

Procedure

1. Upgrade your vCenter Server instance to version 8.0.

The vCenter Server configuration is not changed during the upgrade.

For more information, see the *vCenter Server Upgrade* documentation.

2. Upgrade all ESXi hosts in the cluster to version 8.0.

Edit host customization and host profile edit operations are not available, while compliance check, attach host profile, and remediate hosts operations are available. For more information, see [KB 2150534](#)

3. Remediate the cluster against the host profile with version 6.7.

During remediation, all host profile settings are applied.

4. Extract a new host profile from an ESXi8.0 host.

During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).

5. Attach the host profile with version 8.0 to the cluster.

The host customization data updates automatically.

All host profile operations become available.

How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateless ESXi6.7 Hosts Only?

You use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 6.7 to version 8.0, if your cluster contains stateless ESXi6.7 hosts.

Prerequisites

- Verify that your cluster contains ESXi6.7 hosts only.
- Verify that a host profile with version 6.7 is attached to the cluster.
- Verify that your vCenter Server is with version 6.7.

Procedure

1. Upgrade your vCenter Server instance to version 8.0.
The vCenter Server configuration is not changed during the upgrade.
For more information, see the *vCenter Server Upgrade* documentation.
2. Create a new rule or edit the existing rule with an image profile version 8.0 and select the cluster.
3. Remediate all ESXi hosts in the cluster by clicking **Remediate Host Associations**.
4. Reboot all ESXi hosts in the cluster.
All hosts are compliant with the host profile version 6.7, but the host profile cannot be edited. Remediation and check compliance operations are available.

What to do next

To be able to customize your host profile, follow these steps.

1. Extract a new host profile from an ESXi8.0 host.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
2. Attach the host profile with version 8.0 to the cluster.
The host customization data updates automatically.
All host profile operations become available.

How Do You Upgrade vSphere Host Profiles and vCenter Server from 7.0 to 8.0

Use the upgrade workflows to resolve version compliance errors when upgrading vCenter Server with version 7.0 to version 8.0.

How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi Hosts of version 6.7 and 7.0?

You use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 7.0 to version 8.0, if your cluster contains stateful ESXi 6.7 and 7.0 hosts.

Prerequisites

- Verify that your cluster contains ESXi6.7 and ESXi7.0 hosts.
- Verify that a host profile with version 6.7 is attached to the cluster.
- Verify that your vCenter Server is with version 7.0.

Procedure

1. Upgrade your vCenter Server instance to version 8.0.
The vCenter Server configuration is not changed during the upgrade.
For more information, see the *vCenter Server Upgrade* documentation.

2. Upgrade all ESXi hosts in the cluster to version 8.0.
Edit host customization and host profile edit operations are not available, while compliance check, attach host profile, and remediate hosts operations are available. For more information, see [KB 2150534](#)
3. Remediate the cluster against the host profile with version 6.7.
During remediation, all host profile settings are applied.
4. Extract a new host profile from an ESXi8.0 host.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
5. Attach the host profile with version 8.0 to the cluster.
The host customization data updates automatically.
All host profile operations become available.

How Do You Upgrade Host Profiles in a vCenter Server Instance with Stateful ESXi7.0 Hosts Only?

You use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 7.0 to version 8.0, if your cluster contains stateful ESXi7.0 hosts.

Prerequisites

- Verify that your cluster contains ESXi7.0 hosts only.
- Verify that a host profile with version 7.0 is attached to the cluster.
- Verify that your vCenter Server is with version 7.0.

Procedure

1. Upgrade your vCenter Server instance to version 8.0.
The vCenter Server configuration is not changed during the upgrade.
For more information, see the *vCenter Server Upgrade* documentation.
2. Upgrade all ESXi hosts in the cluster to version 8.0.
Edit host customization and host profile edit operations are not available, while compliance check, attach host profile, and remediate hosts operations are available. For more information, see [KB 2150534](#)
3. Remediate the cluster against the host profile with version 7.0.
During remediation, all host profile settings are applied.
4. Extract a new host profile from an ESXi8.0 host.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
5. Attach the host profile with version 8.0 to the cluster.
The host customization data updates automatically.
All host profile operations become available.

How Do You Upgrade Host Profiles in a vCenter Server Instance With Stateless ESXi7.0 Hosts Only?

You use the present workflow to resolve version compliance errors when upgrading vCenter Server with version 7.0 to version 8.0, if your cluster contains stateless ESXi7.0 hosts.

Prerequisites

- Verify that your cluster contains ESXi7.0 hosts only.

- Verify that a host profile with version 7.0 is attached to the cluster.
- Verify that your vCenter Server is with version 7.0.

Procedure

1. Upgrade your vCenter Server instance to version 8.0.
The vCenter Server configuration is not changed during the upgrade.
For more information, see the *vCenter Server Upgrade* documentation.
2. Create a new rule or edit the existing rule with an image profile version 8.0 and select the cluster.
3. Remediate all ESXi hosts in the cluster by clicking **Remediate Host Associations**.
4. Reboot all ESXi hosts in the cluster.
All hosts are compliant with the host profile version 7.0, but the host profile cannot be edited. Remediation and check compliance operations are available.

What to do next

To be able to customize your host profile, follow these steps.

1. Extract a new host profile from an ESXi8.0 host.
During host profile extraction, some parameters of the configured host profile policy change. For more information, see [Answer File Field and vSphere Host Profile Extraction](#).
2. Attach the host profile with version 8.0 to the cluster.
The host customization data updates automatically.
All host profile operations become available.

Using vSphere Host Profiles to Migrate an NSX Virtual Distributed Switch to a vSphere Distributed Switch

For clusters enabled with VMware NSX®, you can use the regular ESXi upgrade workflow to migrate the NSX-managed Virtual Distributed Switches of the hosts to vCenter Server-managed vSphere Distributed Switches. By using host profiles, you can perform the upgrade and the host switch migration in an environment where all hosts are provisioned with Auto Deploy.

Starting with vSphere 7.0, the vSphere Distributed Switch supports the NSX functionality. Migrating your host switch to vSphere Distributed Switch 7.0 ensures optimal pNIC usage, and lets you manage the networking for your NSX hosts from vCenter Server.

Requirements

- Verify that the ESXi hosts to upgrade are version 6.7 or later.
- vCenter Server 7.0 Update 2
- VMware NSX® 3.1.1

As a good practice, contact VMware support to assess the impact of migrating to vSphere Distributed Switch 7.0.

Workflow

1. Upgrade vCenter Server to version 7.0 Update 2.
For more information about upgrading , see the *vSphere Upgrade* documentation.
2. Upload a software depot that contains an ESXi 7.0 Update 2 image. to vCenter Server.
For more information, see the *VMware ESXi Installation and Setup* documentation.

3. In NSX Manager, run the migration readiness pre-check, address any configuration issues, review the recommended topology, and create a vSphere Distributed Switch.

Note: At this point, do not make any configuration changes to VMware NSX®.

For more information about the steps you need to perform in NSX Manager, see "Migrate Host Switch to vSphere Distributed Switch" in the *NSX Administration* documentation.

4. Create a new image rule or update the existing one with the uploaded ESXi 7.0 Update 2 image.
For more information, see the *VMware ESXi Installation and Setup* documentation.
5. Reboot the reference host.
The reboot triggers the migration of the host switches from NSX Virtual Distributed Switches to vSphere Distributed Switches.
After the reboot, the ESXi version on the reference host is upgraded to 7.0 Update 2. A new TNP with vSphere Distributed Switch is attached to the cluster.
6. Update the host profile by copying the settings from the reference host.
The updated host profile is configured with vSphere Distributed Switch instead of NSX Virtual Distributed Switch.
For more information, see [How Do You Copy Settings from a Reference Host to a vSphere Host Profile?](#).
7. Reboot all ESXi hosts in the cluster.
After the reboot, all hosts in the cluster use the new host profile, which is configured with vSphere Distributed Switch instead of NSX Virtual Distributed Switch.

Troubleshooting vSphere Host Profiles

The Host Profiles troubleshooting topics provide solutions to problems that you might encounter when performing tasks for vSphere Host Profiles in vCenter Server .

If upgrading your vCenter Server to version 8.0 results with a failure due to an existing host profile in the inventory, see *vCenter Server Upgrade* for Upgrade Issues with vCenter Server Containing Host Profiles.

Host Customization Data Is Missing

Deploying a stateless host results with a failure in exiting maintenance mode.

When you are in the process of deploying a stateless host, the host fails in exiting maintenance mode.

In an environment where the following takes place for a host profile:

- The host profile is attached to a stateless host.
- The host profile has host-specific settings.

When the stateless host is rebooting, applying the host profile fails because the host does not exit from a maintenance mode.

- Provide a host customization as described in [How Do You Edit Host Customizations?](#).
- Remediate the host profile to the host again.
- Reboot the host, if needed.
- Refer to the host profiles command-line utility in [Host Profiles CLI](#) for customizing stateless hosts.

Reference Host is Unavailable

Creating a vSphere Host Profile in the vCenter Server inventory results with a failure.

When you perform Host Profiles operations in the vCenter Server inventory such as:

- Edit a Host Profile.
- Import a Host Profile.
- Export a Host Profile.

The process of creating a host profile fails with an error.

```
There is no suitable host in the inventory as reference host for the profile Host Profile. The profile does not have any associated reference host.
```

The vCenter Server inventory does not have a compatible host that acts as a reference host for the host profile you are trying to create.

Add a host that acts as a reference host in the vCenter Server inventory with the same version as the host profile you are trying to create.

Compliance Failures with vSphere Storage Host Profiles

When you use storage devices, that are not shared across a cluster, but the vSphere storage stack cannot detect them as local (for example, some SAS devices), applying a host profile might result with compliance failures.

To resolve the compliance failures caused by using local storage devices, use the upgraded Pluggable Storage Architecture (PSA) host profile policies.

To find out whether a device is detected as local or not, you can check the **Is Local** setting for a device by running the command `esxcli storage core device list -d naa.xxxx` in the ESXi shell. For more information on this command and identifying disks or LUNs, see [KB 1014953](#).

Note: ESXi diagnostic data that you obtain by running the `vm-support` command contains host profiles information which includes storage host profiles, PSA, NMP, and Virtual Volumes data. No sensitive information, such as passwords, is collected.

Compliance Failures with Serial Attached SCSI (SAS) Devices

This procedure resolves compliance failures for SAS devices that are not detected as local by the vSphere storage stack.

Extract a host profile from a reference host. See [How Do You Create a vSphere Host Profile?](#) for instructions.

1. Navigate to **Host Profiles** main view and select the host profile that you want to edit.
2. Click the **Configure** tab, and click **Edit Host Profile**.
3. Select **Storage configuration > Pluggable Storage Architecture configuration > PSA device sharing > *name of device***
4. For each device not shared across the cluster, deactivate **Device is shared clusterwide**.

Note: By default, the **Device is shared clusterwide** setting is deactivated for devices detected as local and enabled for non-local devices. This setting allows storage host profiles to ignore local devices during compliance checks.

The **Device is Shared Clusterwide** setting for PSA devices helps you determine which devices in the cluster should be configured by a host profile. Correctly setting this value for devices in the cluster eliminates compliance errors due to non-shared devices.

5. Remediate the host profile to the reference host for the changes to take effect.

Compliance failures with SAN Boot LUN Devices

This procedure resolves compliance failures for SAN boot LUN devices shared across the cluster, but logically local to the ESXi host.

Extract a host profile from a reference host. See [How Do You Create a vSphere Host Profile?](#) for instructions.

To avoid this compliance failure from occurring, set the **Device is shared clusterwide** value to **False** before extracting the host profile from the reference host. You can deactivate the value by using the command `esxcli storage core`

`device setconfig -d naa.xxxx --shared-clusterwide=false`, where `naa.xxxx` is the unique device identifier generated in naa format.

1. Navigate to **Host Profiles** main view and select the host profile that you want to edit.
2. Click the **Configure** tab, and click **Edit Host Profile**.
3. Select **Storage configuration > Pluggable Storage Architecture configuration > PSA device sharing > name of device**
4. Deactivate the **Device is shared clusterwide** in the host profile by deselecting the check box.

When applying the host profile to the target host, the boot device settings for the remote boot LUN device are copied from the reference host into the target host.

5. Optional: Select **Storage configuration > Pluggable Storage Architecture configuration > Host boot device configuration** and verify that the boot LUN ID is correctly captured.
6. Remediate the host profile to the reference host for the changes to take effect.
If you must re-extract the host profile (for example, if you attach more shared SAN boot LUNs to your cluster), you do not need to reconfigure the **Device is Shared Clusterwide** setting for devices that you previously configured through ESXCLI.

Edit Settings for vSphere Host Profiles is Failing

Editing a vSphere host profile results in an error.

In a vCenter Server inventory upgraded to version 6.5, if you try to edit the settings for a host profile with version 5.1, one of the following errors might appear:

```
Unexpected status code: 503
```

or

```
There are no hosts available in the inventory at the version for the selected Host Profile
```

vSphere 6.5 supports only hosts and host profiles with version 5.5 and later.

- Extract a Host Profile from an already upgraded host to version 5.5. or later.
- Use the **Copy Settings from Host** option to copy the configuration settings from a host already upgraded to a version 5.5. or later.
Note: Using the **Copy Settings from Host** option overwrites the settings within the existing Host Profile or sets them to the default values.

vSphere Host Profile without an NFS Datastore

A host profile compliance check fails for NFS datastore.

When you confirm the compliance of your ESXi host to its attached Host Profile, a non-compliant status is present for the NFS datastore. The compliance check detects the mounted NFS datastore to the ESXi host as an additional datastore.

The remediate or compliance check failure occurs when the following two conditions take place:

- The extracted Host Profile does not have NFS storage (NasStorageProfile).
- The ESXi host, attached to the Host Profile, has an already mounted NFS storage.

1. Create a dummy NFS datastore in the Host Profile.
2. Deactivate the dummy NFS datastore so that the NFS storage profile is fully deactivated.

Upgraded Host is Not Compliant with a Legacy Host Profile

An upgraded ESXi host with attached legacy vSphere host profile results in a non-compliant status after reboot.

When you upgrade a host to the latest ESXi version and remediate the attached legacy host profile to that host, the compliance status might be updated to non-compliant after the host reboots.

You are using a legacy host profile, which is attached to an ESXi host version 6.7 or earlier and contains the following sub-profiles:

- `Misc.LogPort`
- `host/vim/vmvisor/plugins/vmware_*`
- Make the ESXi host compliant with the legacy host profile.
 - a. From the **Host Profiles** main view, right-click the host profile that you want to edit and select **Edit Host Profile**.
 - b. In the **Edit host profile** page, deactivate the sub-profile(s).
 - c. **Check Host Profile Compliance**.
- Upgrade the legacy host profile to the same version as the ESXi host.
 - a. Right-click the legacy host profile and select **Copy Settings from Host**.
 - b. Select the upgraded ESXi host from which you want to copy the configuration settings and click **OK**.

vSphere Networking

vSphere Networking provides information about configuring networking for VMware vSphere®, including how to create vSphere distributed switches and vSphere standard switches.

vSphere Networking also provides information on monitoring networks, managing network resources, and networking best practices.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

The information presented is written for experienced Windows or Linux system administrators who are familiar with network configuration and virtual machine technology.

Introduction to vSphere Networking

Get to know the basic concepts of the vSphere networking and learn to set up and configure a network in a vSphere environment.

Networking Concepts

A few concepts are essential for a thorough understanding of virtual networking. If you are new to vSphere, it is helpful to review these concepts.

Networking Concept	Description
Physical Network	A network of physical machines that are connected so that they can send data to and receive data from each other. VMware ESXi runs on a physical machine.
Virtual Network	A network of virtual machines running on a physical machine that are connected logically to each other so that they can send data to and receive data from each other. Virtual machines can be connected to the virtual networks that you create when you add a network.
Opaque Network	An opaque network is a network created and managed by a separate entity outside of vSphere. For example, logical networks that are created and managed by VMware NSX® appear in vCenter Server as opaque networks of the type nsx.LogicalSwitch. Select an opaque network as the backing for a VM network adapter. To manage an opaque network, use the management tools associated with the opaque network, such as VMware NSX VMware NSX® Manager™ or the VMware NSX API VMware NSX® API management tools. Note: With VMware NSX® 3.0, it is now possible to run NSX directly on vSphere Distributed Switch (vDS) version 7.0 or later. Such networks are not opaque, and are identified as NSX logical segments running on vDS 7.0. For more information, see Knowledge Base article KB #79872.

Networking Concept	Description
Physical Ethernet Switch	A physical ethernet switch manages network traffic between machines on the physical network. A switch has multiple ports, each of which can be connected to a single machine or another switch on the network. Each port can be configured to behave in certain ways depending on the needs of the machine connected to it. The switch learns which hosts are connected to which of its ports and uses that information to forward traffic to the correct physical machines. Switches are the core of a physical network. Multiple switches can be connected together to form larger networks.
vSphere Standard Switch	It works much like a physical Ethernet switch. It detects which virtual machines are logically connected to each of its virtual ports and uses that information to forward traffic to the correct virtual machines. A vSphere standard switch can be connected to physical switches by using physical Ethernet adapters, also referred to as uplink adapters, to join virtual networks with physical networks. This type of connection is similar to connecting physical switches together to create a larger network. Even though a vSphere standard switch works much like a physical switch, it does not have some of the advanced functionality of a physical switch.
vSphere Distributed Switch	A vSphere distributed switch acts as a single switch across all associated hosts in a data center to provide centralized provisioning, administration, and monitoring of virtual networks. You configure a vSphere distributed switch on the vCenter Server system and the configuration is propagated to all hosts that are associated with the switch. This lets virtual machines maintain consistent network configuration as they migrate across multiple hosts.
Host Proxy Switch	A hidden standard switch that resides on every host that is associated with a vSphere distributed switch. The host proxy switch replicates the networking configuration set on the vSphere distributed switch to the particular host.
Standard Port Group	Network services connect to standard switches through port groups. Port groups define how a connection is made through the switch to the network. Typically, a single standard switch is associated with one or more port groups. A port group specifies port configuration options such as bandwidth limitations and VLAN tagging policies for each member port.
Distributed Port	A port on a vSphere distributed switch that connects to a host's VMkernel or to a virtual machine's network adapter.
Distributed Port Group	A port group associated with a vSphere distributed switch that specifies port configuration options for each member port. Distributed port groups define how a connection is made through the vSphere distributed switch to the network.

Networking Concept	Description
NSX Distributed Port Group	A port group associated with a vSphere distributed switch that specifies port configuration options for each member port. To distinguish between vSphere distributed port groups and NSX port groups, in the vSphere Client the NSX virtual distributed switch, and its associated port group, is identified with the  icon. NSX appears as an opaque network in vCenter Server, and you cannot configure NSX settings in vCenter Server. The NSX settings displayed are read only. You configure NSX distributed port groups using VMware NSX® Manager or the VMware NSX API management tools. To learn about configuring NSX, see the <i>NSX Data Center for vSphere</i> documentation.
NSX Manager Appliance Plug-in	NSX Manager Appliance can be installed from the vSphere Web Client. You do not need to perform any installation operations from the NSX Manager UI. After the NSX Manager is installed, NSX appears as a plug-in in the vCenter Server that is ready to install VMware NSX for virtual networking or security use cases. For more information, see the <i>VMware NSX</i> documentation. This functionality is available on VMware vSphere 7.0 Update 3 and later and VMware NSX® 3.2 and later.
NIC Teaming	NIC teaming occurs when multiple uplink adapters are associated with a single switch to form a team. A team can either share the load of traffic between physical and virtual networks among some or all of its members, or provide passive failover if there is a hardware failure or a network outage.
VLAN	VLAN enable a single physical LAN segment to be further segmented so that groups of ports are isolated from one another as if they were on physically different segments. The standard is 802.1Q.
VMkernel TCP/IP Networking Layer	The VMkernel networking layer provides connectivity to hosts and handles the standard infrastructure traffic of vSphere vMotion, IP storage, Fault Tolerance, and vSAN.
IP Storage	Any form of storage that uses TCP/IP network communication as its foundation. iSCSI and NFS can be used as virtual machine datastores and for direct mounting of .ISO files, which are presented as CD-ROMs to virtual machines.
TCP Segmentation Offload	TCP Segmentation Offload, TSO, allows a TCP/IP stack to emit large frames (up to 64 KB) even though the maximum transmission unit (MTU) of the interface is smaller. The network adapter then separates the large frame into MTU-sized frames and prepends an adjusted copy of the initial TCP/IP headers.

Network Services

A virtual network provides several services to the host and virtual machines. You can enable two types of network services in ESXi:

- Connecting virtual machines to the physical network and to each other.
- Connecting VMkernel services (such as NFS, iSCSI, or vMotion) to the physical network.

VMware ESXi Dump Collector Support

The ESXi Dump Collector sends the state of the VMkernel memory, that is, a core dump to a network server when the system encounters a critical failure. The ESXi Dump Collector in ESXi supports both vSphere standard and distributed Switches. The ESXi Dump Collector can also use any active uplink adapter from the team of the port group that handles the VMkernel adapter for the collector.

Changes to the IP address for the ESXi Dump Collector interface are automatically updated if the IP addresses for the configured VMkernel adapter changes. The ESXi Dump Collector also adjusts its default gateway if the gateway configuration of the VMkernel adapter changes.

If you try to delete the VMkernel network adapter used by the ESXi Dump Collector, the operation fails and a warning message appears. To delete the VMkernel network adapter, deactivate dump collection and delete the adapter.

There is no authentication or encryption in the file transfer session from a crashed host to the ESXi Dump Collector. You can configure the ESXi Dump Collector on a separate VLAN when possible to isolate the ESXi core dump from regular network traffic.


For information about installing and configuring the ESXi Dump Collector, see the *vCenter Server Installation and Setup* documentation.

How do I Set Up Networking with vSphere Standard Switches

vSphere standard switches handle network traffic at the host level in a vSphere deployment. You can create abstracted network devices called vSphere Standard Switches. Learn how to use standard switches to provide network connectivity to hosts and virtual machines. A standard switch can bridge traffic internally between virtual machines in the same VLAN and link to external networks.

Standard Switch Overview

To provide network connectivity to hosts and virtual machines, you connect the physical NICs of the hosts to uplink ports on the standard switch. Virtual machines have network adapters (vNICs) that you connect to port groups on the standard switch. Every port group can use one or more physical NICs to handle their network traffic. If a port group does not have a physical NIC connected to it, virtual machines on the same port group can only communicate with each other but not with the external network.

Figure 67: vSphere Standard Switch architecture

A vSphere Standard Switch is very similar to a physical Ethernet switch. Virtual machine network adapters and physical NICs on the host use the logical ports on the switch as each adapter uses one port. Each logical port on the standard switch is a member of a single port group. For information about maximum allowed ports and port groups, see the *Configuration Maximums* documentation.

Standard Port Groups

Each port group on a standard switch is identified by a network label, which must be unique to the current host. You can use network labels to make the networking configuration of virtual machines portable across hosts. You should give the same label to the port groups in a data center that use physical NICs connected to one broadcast domain on the physical network. Conversely, if two port groups are connected to physical NICs on different broadcast domains, the port groups should have distinct labels.

For example, you can create *Production* and *Test environment* port groups as virtual machine networks on the hosts that share the same broadcast domain on the physical network.

A VLAN ID, which restricts port group traffic to a logical Ethernet segment within the physical network, is optional. For port groups to receive the traffic that the same host sees, but from more than one VLAN, the VLAN ID must be set to VGT (VLAN 4095).

Number of Standard Ports

To ensure efficient use of host resources on ESXi hosts, the number of ports of standard switches are dynamically scaled up and down. A standard switch on such a host can expand up to the maximum number of ports supported on the host.

Create a vSphere Standard Switch

Create a vSphere Standard Switch to provide network connectivity for hosts, virtual machines, and to handle VMkernel traffic. Depending on the connection type that you want to create, you can create a new vSphere Standard Switch with a VMkernel adapter, only connect physical network adapters to the new switch, or create the switch with a virtual machine port group.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual switches**.
3. Click **Add networking**.
4. Select a connection type for which you want to use the new standard switch and click **Next**.

Option	Description
VMkernel Network Adapter	Create a new VMkernel adapter to handle host management traffic, vMotion, network storage, fault tolerance, or vSAN traffic.
Physical Network Adapter	Add physical network adapters to an existing or a new standard switch.
Virtual Machine Port Group for a Standard Switch	Create a new port group for virtual machine networking.

5. Select **New standard switch** and click **Next**.
6. Add physical network adapters to the new standard switch.
 - a) Under Assigned adapters, click **Add adapters**.
 - b) Select one or more physical network adapters from the list and click **OK**.
For higher throughput and to provide redundancy, configure at least two physical network adapters in the Active list.
 - c) Optional: Use the **Move up** and **Move down** arrows in the **Assigned adapters** list to change the position of the adapter.
 - d) Click **Next**.
7. If you create the new standard switch with a VMkernel adapter or virtual machine port group, enter connection settings for the adapter or the port group.

Option	Description
VMkernel adapter	1. Enter a label that indicates the traffic type for the VMkernel adapter, for example <code>vMotion</code>. 2. Set a VLAN ID to identify the VLAN that the network traffic of the VMkernel adapter will use. 3. Select IPv4, Ipv6 or both. 4. Select an option from the drop-down menu to set the MTU size. If you select Custom, enter a value for the MTU size. You can enable jumbo frames by setting an MTU value greater than 1500. You cannot set an MTU size greater than 9000 bytes. 5. Select a TCP/IP stack. After you set a TCP/IP stack for the VMkernel adapter, you cannot change it later. If you select the vMotion or the Provisioning TCP/IP stack, you will be able to use only this stack to handle vMotion or Provisioning traffic on the host. 6. If you use the default TCP/IP stack, select from the available services.

Option	Description
	7. Configure IPv4 and IPv6 settings.
Virtual machine port group	1. Enter a network Label or the port group, or accept the generated label. 2. Set the VLAN ID to configure VLAN handling in the port group.

8. On the Ready to Complete page, click **Finish**.

- You might need to change the teaming and failover policy of the new standard switch. For example, if the host is connected to an Etherchannel on the physical switch, you must configure the vSphere Standard Switch with Route based on IP hash as a load balancing algorithm. See [What is Teaming and Failover Policy](#) for more information.
- If you create the new standard switch with a port group for virtual machine networking, connect virtual machines to the port group.

Port Group Configuration for Virtual Machines

Learn how to add or modify a virtual machine port group to set up traffic management on a set of virtual machines.

The **Add Networking** wizard in the vSphere Client guides you through the process to create a virtual network to which virtual machines can connect, including creating a vSphere Standard Switch and configuring settings for a network label.

When you set up virtual machine networks, consider whether you want to migrate the virtual machines in the network between hosts. If so, ensure that the same broadcast domain is accessible by both hosts—that is, the same Layer 2 subnet.

ESXi does not support virtual machine migration between hosts in different broadcast domains because the migrated virtual machine might require systems and resources that it would no longer have access to in the new network. Even if your network configuration is set up as a high-availability environment or includes intelligent switches that can resolve the virtual machine's needs across different networks, you might experience lag times as the Address Resolution Protocol (ARP) table updates and resumes network traffic for the virtual machines.

Virtual machines reach physical networks through uplink adapters. A vSphere Standard Switch can transfer data to external networks only when one or more network adapters are attached to it. When two or more adapters are attached to a single standard switch, they are transparently teamed.

Add a Virtual Machine Port Group

Learn how to add VM port groups on a vSphere Standard Switch to provide connectivity and common network configuration for virtual machines.

1. In the vSphere Client, navigate to the host.
2. Right-click the host and select **Add Networking**.
3. In **Select connection type**, select **Virtual Machine Port Group for a Standard Switch** and click **Next**.
4. In **Select target device**, select an existing standard switch or create a new standard switch.
5. If the new port group is for an existing standard switch, navigate to the switch.
 - a) Click **Browse**.
 - b) Select a standard switch from the list and click **OK**.
 - c) Click **Next** and go to [Step 7](#).
6. Optional: If you choose to create a new standard switch, enter a value for the MTU size and click **Next**.

You can create a standard switch with or without adapters.

If you create a standard switch without physical network adapters, all traffic on that switch is confined to that switch. No other hosts on the physical network or virtual machines on other standard switches can send or receive traffic over this standard switch. You might create a standard switch without physical network adapters if you want a group of virtual machines to be able to communicate with each other, but not with other hosts or with virtual machines outside the group.

- a) Click **Add adapters**.
 - b) Select an adapter from the **Network Adapters** list and click **OK**.
 - c) Optional: Use the up and down arrows in the **Assigned adapters** list to change the position of the adapter if needed.
 - d) Click **Next**.
7. On the Connection settings page, identify traffic through the ports of the group.
 - a) Type a **Network label** for the port group, or accept the generated label.

Note: Port group names must not contain a colon character (:).
 - b) Set the **VLAN ID** to configure VLAN handling in the port group.

The VLAN ID also reflects the VLAN tagging mode in the port group.

VLAN Tagging Mode	VLAN ID	Description
External Switch Tagging (EST)	0	The virtual switch does not pass traffic associated with a VLAN.
Virtual Guest Tagging (VGT)	4095	Virtual machines handle VLANs. The virtual switch passes traffic from any VLAN.

- c) Click **Next**.
8. Review the port group settings in the Ready to complete page, and click **Finish**.
Click **Back** if you want to change any settings.

Edit a Standard Switch Port Group

Learn how to edit the name and VLAN ID of a standard switch port group, and override networking policies at the port group level.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select a standard switch from the list.
The topology diagram of the switch appears.
4. In the topology diagram of the switch, click the name of the port group.
5. Next to the the topology diagram title, click the horizontal elipsis icon and select **Edit settings**.
6. On the Properties page, rename the port group in the **Network label** text field.
7. Configure VLAN tagging in the **VLAN ID** drop-down menu.

VLAN Tagging Mode	VLAN ID	Description
External Switch Tagging (EST)	0	The virtual switch does not pass traffic associated with a VLAN.
Virtual Guest Tagging (VGT)	4095	Virtual machines handle VLANs. The virtual switch passes traffic from any VLAN.

8. On the Security page, override the switch settings for protection against MAC address changes, forged transmits and for running virtual machines in promiscuous mode.
9. On the Traffic shaping page, override at the port group level the size of average and peak bandwidth and of bursts.
10. On the Teaming and failover page, override the teaming and failover settings inherited from the standard switch.
You can configure traffic distribution and rerouting between the physical adapters associated with the port group. You can also change the order in which host physical adapters are used upon failure.
11. Click **OK**.

Remove a Port Group from a vSphere Standard Switch

Learn how to remove port groups from vSphere Standard Switches in case you no longer need the associated labeled networks.

Verify that there are no powered-on virtual machines connected to the port group that you want to remove.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the standard switch.
4. From the topology diagram of the switch, select the port group that you want to remove by clicking its label.
5. From the toolbar in the switch topology, click the **Remove** action icon .

vSphere Standard Switch Properties

vSphere Standard Switch settings control switch-wide defaults for ports, which can be overridden by port group settings for each standard switch. Learn how to edit standard switch properties, such as the uplink configuration and the number of available ports.

Number of Ports on ESXi Hosts

To ensure efficient use of host resources on ESXi hosts, the ports of virtual switches are dynamically scaled up and down. A switch on such a host can expand up to the maximum number of ports supported on the host. The port limit is determined based on the maximum number of virtual machines that the host can handle.

Change the Size of the MTU on a vSphere Standard Switch

Learn how to change the size of the maximum transmission unit (MTU) on a vSphere Standard Switch to improve the networking efficiency by increasing the amount of payload data transmitted with a single packet, that is, enabling jumbo frames.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select a standard switch from the table and click **Edit settings**.
4. Change the **MTU (Bytes)** value for the standard switch.
You can enable jumbo frames by setting an MTU value greater than 1500. You cannot set an MTU size greater than 9000 bytes.
5. Click **OK**.

Change the Speed of a Physical Adapter

Learn how to change the connection speed and duplex of a physical adapter to transfer data in compliance with the traffic rate.

If the physical adapter supports SR-IOV, you can enable it and configure the number of virtual functions to use for virtual machine networking.

1. In the vSphere Client, navigate to a host.
2. On the **Configure** tab, expand **Networking** and select **Physical adapters**.
The physical network adapters of the host appear in a table that contains details for each physical network adapter.
3. Select the physical network adapter from the list and click the **Edit adapter settings** icon.
4. Select speed and duplex mode of the physical network adapter from the drop-down menu.
The speed and duplex settings you select must match those configured on the physical switch. Speed is the rate of the interface, usually listed in megabits per second (Mbps). Duplex refers to how data flows on the interface.

5. Click **OK**.

Assign Physical Adapters to a vSphere Standard Switch

Learn how to assign a physical adapter to a standard switch to provide connectivity to virtual machines and VMkernel adapters on the host. You can form a team of NICs to distribute traffic load and to configure failover.

NIC teaming combines multiple network connections to increase throughput and provide redundancy should a link fail. To create a team, you associate multiple physical adapters to a single vSphere Standard Switch.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the standard switch you want to add a physical adapter to.
4. Click **Manage Physical Adapters**.
5. Add one or more available physical network adapters to the switch.
 - a) Click **Add adapters**, select one or more network adapters from the list and click **OK**.
The selected adapters appear in the failover group list under the Assigned Adapters list.
 - b) Optional: Use the up and down arrows to change the position of an adapter in the failover groups.
The failover group determines the role of the adapter for exchanging data with the external network, that is, active, standby or unused. By default, the adapters are added as active to the standard switch.
6. Click **OK** to apply the physical adapter configuration.

Topology of a vSphere Standard Switch

View the structure and components of a vSphere Standard Switch by using its topology diagram.

The topology diagram of a standard switch provides a visual representation of the adapters and port groups connected to the switch.

From the diagram you can edit the settings of a selected port group and of a selected adapter.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the standard switch from the list.

The diagram appears under the list of virtual switches on the host.


Diagram of a Standard Switch That Connects the VMkernel and Virtual Machines to the Network

In your virtual environment, a vSphere Standard Switch handles VMkernel adapters for vSphere vMotion and for the management network, and virtual machines grouped. You can use the central topology diagram to examine whether a virtual machine or VMkernel adapter is connected to the external network and to identify the physical adapter that carries the data.

Figure 68: Topology Diagram of a vSphere Standard Switch

How do I Set Up Networking with vSphere Distributed Switches

With vSphere distributed switches you can set up and configure networking in a vSphere environment.

Figure 69: vSphere Distributed Switch Architecture

A network switch in vSphere consists of two logical sections that are the data plane and the management plane. The data plane implements the packet switching, filtering, tagging, and so on. The management plane is the control structure that you use to configure the data plane functionality. A vSphere Standard Switch contains both data and management planes, and you configure and maintain each standard switch individually.

A vSphere Distributed Switch separates the data plane and the management plane. The management functionality of the distributed switch resides on the vCenter Server system that lets you administer the networking configuration of your environment on a data center level. The data plane remains locally on every host that is associated with the distributed switch. The data plane section of the distributed switch is called a host proxy switch. The networking configuration that you create on vCenter Server (the management plane) is automatically pushed down to all host proxy switches (the data plane).

The vSphere Distributed Switch introduces two abstractions that you use to create consistent networking configuration for physical NICs, virtual machines, and VMkernel services.

Uplink port group

An uplink port group or dvuplink port group is defined during the creation of the distributed switch and can have one or

more uplinks. An uplink is a template that you use to configure physical connections of hosts as well as failover and load balancing policies. You map physical NICs of hosts to uplinks on the distributed switch. At the host level, each physical NIC is connected to an uplink port with a particular ID. You set failover and load balancing policies over uplinks and the policies are automatically propagated to the host proxy switches, or the data plane. In this way you can apply consistent failover and load balancing configuration for the physical NICs of all hosts that are associated with the distributed switch.

Distributed port group


Distributed port groups provide network connectivity to virtual machines and accommodate VMkernel traffic. You identify each distributed port group by using a network label, which must be unique to the current data center. You configure NIC teaming, failover, load balancing, VLAN, security, traffic shaping, and other policies on distributed port groups. The virtual ports that are connected to a distributed port group share the same properties that are configured to the distributed port group. As with uplink port groups, the configuration that you set on distributed port groups on vCenter Server (the management plane) is automatically propagated to all hosts on the distributed switch through their host proxy switches (the data plane). In this way you can configure a group of virtual machines to share the same networking configuration by associating the virtual machines to the same distributed port group.

For example, suppose that you create a vSphere Distributed Switch on your data center and associate two hosts with it. You configure three uplinks to the uplink port group and connect a physical NIC from each host to an uplink. Each uplink has two physical NICs, one from each host which is mapped to it. For example, Uplink 1 is configured with vmnic0 from Host 1 and Host 2. Next you create the Production and the VMkernel network distributed port groups for virtual machine networking and VMkernel services. Respectively, a representation of the Production and the VMkernel network port groups is also created on Host 1 and Host 2. All policies that you set to the Production and the VMkernel network port groups are propagated to their representations on Host 1 and Host 2.

To ensure efficient use of host resources, the number of distributed ports of proxy switches is dynamically scaled up and down. A proxy switch on such a host can expand up to the maximum number of ports supported on the host. The port limit is determined based on the maximum number of virtual machines that the host can handle.


vSphere Distributed Switch Data Flow

The data flow from the virtual machines and VMkernel adapters down to the physical network depends on the NIC teaming and load balancing policies that are set to the distributed port groups. The data flow also depends on the port allocation on the distributed switch.

Figure 70: NIC Teaming and Port Allocation on a vSphere Distributed Switch

For example, suppose that you create the VM network and the VMkernel network distributed port groups, respectively with 3 and 2 distributed ports. The distributed switch allocates ports with IDs from 0 to 4 in the order that you create the distributed port groups. Next, you associate Host 1 and Host 2 with the distributed switch. The distributed switch allocates ports for every physical NIC on the hosts, as the numbering of the ports continues from 5 in the order that you add the hosts. To provide network connectivity on each host, you map vmnic0 to Uplink 1, vmnic1 to Uplink 2, and vmnic2 to Uplink 3.

To provide connectivity to virtual machines and to accommodate VMkernel traffic, you configure teaming and failover to the VM network and to the VMkernel network port groups. Uplink 1 and Uplink 2 handle the traffic for the VM network port group, and Uplink 3 handles the traffic for the VMkernel network port group.

Figure 71: Packet Flow on the Host Proxy Switch

On the host side, the packet flow from virtual machines and VMkernel services passes through particular ports to reach the physical network. For example, a packet sent from VM1 on Host 1 first reaches port 0 on the VM network distributed port group. Because Uplink 1 and Uplink 2 handle the traffic for the VM network port group, the packet can continue from uplink port 5 or uplink port 6. If the packet goes through uplink port 5, it continues to vmnic0, and if the packet goes to uplink port 6, it continues to vmnic1.

What is Network Offloads Compatibility

Starting with vSphere 8.0, vSphere Distributed Services Engine (vDSE) introduces virtual infrastructure as a distributed architecture with the addition of data processing units (DPUs) also known as SmartNIC that enable offloading infrastructure functions from the host or server CPUs to data processing units (DPUs).

The network offloads compatibility allows you to offload networking operations to the DPU device. You can offload the networking functionality from the ESXi host to DPU for better performance. vSphere Distributed Switch backed by ESXi on DPU supports the following modes:

- Non-offloading mode before is enabled: The DPU is used as a traditional NIC.
- Offloading mode after is enabled: Traffic forwarding logic is offloaded from the ESXi host to the vSphere Distributed Switch backed by the DPU.

Hosts backed by DPU are associated with the vSphere Distributed Switch. It is configured during the creation of the distributed switch. Network offloads compatibility cannot be modified after associating hosts to the distributed switch. You can only add hosts that are backed by DPU to these distributed switches. ESXi on DPU is used as a traditional NIC until transport node is configured. The vSphere Distributed Switch on vCenter Server indicates if network offloading is permitted when is enabled.

Features supported by vSphere Distributed Switch backed by the DPU.

- Creation and deletion of the vSphere distributed switch.

- Configuration management.
- vSphere distributed switch health check.
- Link aggregation control protocol (LACP).
- Port mirroring.
- Private LANs.
- Link layer discovery protocol.

Note: Features not supported by vSphere Distributed Switch backed by DPU.

- Network I/O control.
- Traffic shaping policies.
- DV filter.

Dual DPU

Starting vSphere 8.0 Update 3, you can use the vDSE with 2 data processing units (DPUs) on high availability mode. For more information on dual DPUs, see [Introducing VMware vSphere® Distributed Services Engine™ and Networking Acceleration by Using DPUs](#).

Dual DPUs can be consumed in High Availability (HA) and Non High Availability (Non-HA) mode.

- HA mode: In this mode, each DPU is consumed by a single offloaded Distributed Virtual Switch (vDS). For instance, if DPU 1 is designated as active then the DPU 2 acts as a stand-by. Stand-by DPU is designated as a backup DPU. If the active DPU fails, then the active networking offload is switched to the stand-by DPU. This provides high availability for the DPU. This switchover minimizes the risk of failure on the active workloads.
If dual DPUs are connected to the same network switch at the same time, only one of them processes data packets. The other DPU is on stand-by mode. However, shadow switches and ports are created on the stand-by DPU. Networking policies are also applied on the DPU but the shadow switch does not process any packets. When ESXi detects an active DPU failure, it initiates a fail over to the stand-by DPU and sends a signal to the shadow switch to enable packet processing.
- Non-HA mode: In this mode, there is no High Availability (HA), but each DPU is allowed to be consumed by a separate offloaded vDS. This mode allows both the DPUs to be consumed for active networking datapath offload.

Enable Network Offloads

To enable the network offloads you have to perform multiple steps in the vCenter Server and in .

Step	Solution
Create a vSphere Distributed Switch	Create a vSphere Distributed Switch
Associate hosts to vSphere Distributed Switch	Add Hosts to a vSphere Distributed Switch
Configure host transport node	Configure NSX host transport node on DPU-enabled vSphere Lifecycle Manager cluster
View the topology of the vSphere Distributed Switch with Network Offloads	View the Topology of Network Offloads Switch

Create a vSphere Distributed Switch


Learn how to create a vSphere distributed switch on a data center to handle the networking configuration of multiple hosts at a time from a central place.

1. In the vSphere Client, right-click a data center from the inventory tree.
2. Select **Distributed Switch > New Distributed Switch**.
3. On the **Name and location** page, enter a name for the new distributed switch, or accept the generated name, and click **Next**.
4. On the **Select version** page, select a distributed switch version and click **Next**.

Option	Description
Distributed Switch: 8.0.3	Compatible with ESXi 8.0.3 and later. Features released with later vSphere Distributed Switch versions are not supported.
Distributed Switch: 8.0.0	Compatible with ESXi 8.0 and later. Features released with later vSphere Distributed Switch versions are not supported.
Distributed Switch: 7.0.3	Compatible with ESXi 7.0.3 and later. Features released with later vSphere Distributed Switch versions are not supported.
Distributed Switch: 7.0.2	Compatible with ESXi 7.0.2 and later. Features released with later vSphere Distributed Switch versions are not supported.
Distributed Switch: 7.0.0	Compatible with ESXi 7.0 and later. Features released with later vSphere Distributed Switch versions are not supported.
Distributed Switch: 6.6.0	Compatible with ESXi 6.7, ESXi 7.0 and later. Features released with later vSphere Distributed Switch versions are not supported.

5. On the **Configure settings** page, configure the distributed switch settings.
 - a) Use the drop-down menu to select the type of **Network Offloads Compatibility**.

Figure 72: Network Offloads Compatibility



By using network offloads compatibility you can offload network and security functions to the DPU device. DPU is a network card that has compute capability embedded in it. You can offload the networking functionality from the ESXi host to DPU for better performance.

- **None:** If you select **None**, Network Offloads compatibility is not supported.
- **Pensando:** If you select **Pensando**, Network Offloads compatibility is supported. **Network I/O Control** is disabled.
- **NVIDIA BlueField:** If you select **NVIDIA BlueField**, Network Offloads compatibility is supported. **Network I/O Control** is disabled.

Note: You can configure Network Offloads compatibility when you use vSphere Distributed Switch 8.0.0 and later.

- b) Use the arrow buttons to select the **Number of uplinks**.

Uplink ports connect the distributed switch to physical NICs on associated hosts. The number of uplink ports is the maximum number of allowed physical connections to the distributed switch per host.

- c) Use the drop-down menu to enable or disable **Network I/O Control**.

By using Network I/O Control you can prioritize the access to network resources for certain types of infrastructure and workload traffic according to the requirements of your deployment. Network I/O Control continuously monitors the I/O load over the network and dynamically allocates available resources.

- d) Optional: Select the **Create a default port group** check box to create a new distributed port group with default settings for this switch. Enter a **Port group name**, or accept the generated name.

If your system has custom port group requirements, create distributed port groups that meet those requirements after you add the distributed switch.

6. In the **DPU Failover Configuration** page, specify how the uplinks in a team are used when a failover occurs by configuring the failover order list. If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the **Move Up** and **Move Down** to move uplinks into different groups.

In the **DPU Failover Configuration** page, you can specify how DPU mapped uplinks can be consumed. When multiple DPUs are present on the host, then they can be attached to same distributed switch in high availability mode, that is one DPU will run the active traffic and other DPU will be available as standby.

- a) Select the type of DPU from the **Preset** drop-down menu.

- **Single DPU:** If you select **Single DPU**, then only one DPU connects to the switch at a time as an active DPU.
- **Dual DPU:** If you select **Dual DPU**, then both the DPUs connect to the switch at the same time. Only one of them is active.
- **Custom:** If you select **Custom**, then you can specify the uplinks.

Option	Description
Active uplinks	Continue to use the uplink if the DPU adapter connectivity is up and active.
Standby uplinks	Use this uplink if the active DPU adapter is down.

7. On the **Ready to complete** page, review the settings you selected and click **Finish**.

Use the **Back** button to edit any settings.

A distributed switch is created in the data center. You can view the features supported on the distributed switch and other details by navigating to the new distributed switch and clicking the **Summary** tab.

Add hosts to the distributed switch and configure their network adapters on the switch.

Related Links

[Add Hosts to a vSphere Distributed Switch on page 2217](#)

Learn to manage the networking of your vSphere environment by using a vSphere Distributed Switch, you must associate hosts with the switch. You connect the physical NICs, VMkernel adapters, and virtual machine network adapters of the hosts to the distributed switch.

[Configure Physical Network Adapters on a vSphere Distributed Switch on page 2220](#)

Learn how to assign physical NICs to uplinks on the switch for the hosts that are associated with a distributed switch. You can configure physical NICs on the distributed switch for multiple hosts at a time.

[Migrate VMkernel Adapters to a vSphere Distributed Switch on page 2221](#)

Learn how to migrate VMkernel adapters to a distributed switch if you want to handle the traffic for VMkernel services by using only this switch when you no longer need the adapters on other standard or distributed switches.

[Create a VMkernel Adapter on a vSphere Distributed Switch on page 2222](#)

Learn how to create a VMkernel adapter on hosts associated with a distributed switch to provide network connectivity to the hosts and to handle the traffic for vSphere vMotion, IP storage, Fault Tolerance logging, and vSAN.

[Migrate Virtual Machine Networking to the vSphere Distributed Switch on page 2224](#)

Learn how to manage virtual machine networking by using a distributed switch, migrate virtual machine network adapters to labeled networks on the switch.

[Use a Host as a Template to Create a Networking Configuration on a vDS](#)

Upgrade a vSphere Distributed Switch to a Later Version

Learn how to upgrade vSphere Distributed Switch version 6.x to a later version. The upgrade lets the distributed switch take advantage of features that are available only in the later version.

- Upgrade vCenter Server to version 8.0.
- Upgrade all hosts connected to the distributed switch to ESXi 8.0

The upgrade of a distributed switch causes the hosts and virtual machines attached to the switch to experience a brief downtime.

Note: If your current VDS version is 6.5, then you might experience a brief downtime while upgrading your switches to a later version. If your current VDS version is 6.6 or later, then you might not experience any downtime while upgrading the switches to a later version.

For more information, see [KB 52621](#).

Note:

To be able to restore the connectivity of the virtual machines and VMkernel adapters if the upgrade fails, back up the configuration of the distributed switch.

If the upgrade is not successful, to recreate the switch with its port groups and connected hosts, you can import the switch configuration file. See [Export vSphere Distributed Switch Configurations](#) and [Import a vSphere Distributed Switch Configuration](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Upgrade > Upgrade Distributed Switch**.
3. Select the vSphere Distributed Switch version that you want to upgrade the switch to and click **Next**.

Option	Description
Distributed Switch: 8.0.3	Compatible with ESXi 8.0.3 and later. New features and enhancements • Dual DPU Support. See Network Offloads Compatibility.
Distributed Switch: 8.0.0	Compatible with ESXi 8.0.0 and later. New features and enhancements. See Network Offloads Compatibility .
Distributed Switch: 7.0.3	Compatible with ESXi 7.0.3 and later. New features and enhancements • NVMe over TCP. See Create a VMkernel Adapter on a vSphere Standard Switch.
Distributed Switch: 7.0.2	Compatible with ESXi 7.0.2 and later. New features and enhancements.

Option	Description
	LACP fast mode. See LACP Support on a vSphere Distributed Switch.
Distributed Switch: 7.0.0	Compatible with ESXi 7.0 and later. New features and enhancements. NSX distributed port group. See Networking Concepts Overview.
Distributed Switch: 6.6.0	Compatible with ESXi 6.7 and later. Features released with later vSphere distributed switch versions are not supported. New features and enhancements. MAC Learning. See What is MAC Learning Policy

- Review host compatibility and click **Next**.

Some ESXi instances that are connected to the distributed switch might be incompatible with the selected target version. Upgrade or remove the incompatible hosts, or select another upgrade version for the distributed switch.

- Complete the upgrade configuration and click **Finish**.

CAUTION: After you upgrade the vSphere Distributed Switch, you cannot revert it to an earlier version. You also cannot add ESXi hosts that are running an earlier version than the new version of the switch.

Edit General and Advanced vSphere Distributed Switch Settings

Learn how use basic settings for a vSphere Distributed Switch with setting the switch name and number of uplinks. You can also learn how to use advanced settings for a distributed switch including Cisco Discovery Protocol and the maximum MTU for the switch.

- In the vSphere Client Home page, click **Networking** and select the distributed switch.
- On the **Configure** tab, expand **Settings** and select **Properties**.
- Click **Edit**.
- Click **General** to edit the vSphere Distributed Switch settings.

Option	Description
Name	Enter the name for the distributed switch.
Network I/O Control	To activate or deactivate Network I/O control , use the drop-down menu. Note: This menu option is deactivated if the Network Offloads compatibility is activated.
Network offloads compatibility	To select the type of network offloads compatibility, use the drop-down menu. None: If you select None, network offloads compatibility is not activated. Pensando: If you select Pensando the Network I/O control is deactivated. NVIDIA BlueField: If you select NVIDIA BlueField the Network I/O control is deactivated.
Description	Add or modify a description of the distributed switch settings.

5. Click **Advanced** to edit the vSphere Distributed Switch settings.

Option	Description
MTU (Bytes)	Maximum MTU size for the vSphere Distributed Switch. To enable jumbo frames, set a value greater than 1500 bytes.
Multicast filtering mode	• Basic. The distributed switch forwards traffic that is related to a multicast group based on a MAC address generated from the last 23 bits of the IPv4 address of the group. • IGMP/MLD snooping. The distributed switch forwards multicast traffic to virtual machines according to the IPv4 and IPv6 addresses of subscribed multicast groups by using membership messages defined by the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery protocol.
Discovery Protocol	1. Select Cisco Discovery Protocol, Link Layer Discovery Protocol, or (disabled) from the Type drop-down menu. 2. Set Operation to Listen, Advertise, or Both. For information about Discovery Protocol, see Switch Discovery Protocol.
Administrator Contact	Enter the name and other details of the administrator for the distributed switch.

6. Click **Uplinks** to edit the vSphere Distributed Switch settings.

Option	Description
Number of uplinks	Select Add to add uplink ports for the distributed switch. Edit uplink names to modify the names of the uplinks.

7. Click **OK**.

Manage Networking on Multiple Hosts on a vSphere Distributed Switch

Learn how to create and manage virtual networks on a vSphere Distributed Switch by adding hosts to the switch and connecting their network adapters to the switch. To create uniform networking configuration throughout multiple hosts on the distributed switch, you can use a host as a template and apply its configuration to other hosts.

Tasks for Managing Host Networking on a vSphere Distributed Switch

You can add new hosts to a vSphere Distributed Switch, connect network adapters to the switch, and remove hosts from the switch. In a production environment, you might need to keep the network connectivity up for virtual machines and VMkernel services while you manage host networking on the distributed switch.

Adding Hosts to a vSphere Distributed Switch

Consider preparing your environment before you add new hosts to a distributed switch.

- Create distributed port groups for virtual machine networking.
- Create distributed port groups for VMkernel services. For example, create distributed port groups for management network, vMotion, and Fault Tolerance.
- Configure enough uplinks on the distributed switch for all physical NICs that you want to connect to the switch. For example, if the hosts that you want to connect to the distributed switch have eight physical NICs each, configure eight uplinks on the distributed switch.
- Make sure that the configuration of the distributed switch is prepared for services with specific networking requirements. For example, iSCSI has specific requirements for the teaming and failover configuration of the distributed port group where you connect the iSCSI VMkernel adapter.

You can use the **Add and Manage Hosts** wizard in the to add multiple hosts at a time.

Managing Network Adapters on a vSphere Distributed Switch

After you add hosts to a distributed switch, you can connect physical NICs to uplinks on the switch, configure virtual machine network adapters, and manage VMkernel networking.

If some hosts on a distributed switch are associated to other switches in your data center, you can migrate network adapters to or from the distributed switch.

If you migrate virtual machine network adapters or VMkernel adapters, make sure that the destination distributed port groups have at least one active uplink, and the uplink is connected to a physical NIC on the hosts. Another approach is to migrate physical NICs, virtual network adapters, and VMkernel adapters simultaneously.

If you migrate physical NICs, leave at least one active NIC that handles the traffic of port groups. For example, if *vmnic0* and *vmnic1* handle the traffic of the *VM Network* port group, migrate *vmnic0* and leave *vmnic1* connected to the group.

Watch the video about migrating VMkernel interfaces and physical NICs to a vSphere Distributed Switch.

Removing Hosts from a vSphere Distributed Switch

Before you remove hosts from a distributed switch, you must migrate the network adapters that are in use to a different switch.

- To add hosts to a different distributed switch, you can use the **Add and Manage Hosts** wizard to migrate the network adapters on the hosts to the new switch all together. You can then remove the hosts safely from their current distributed switch.
- To migrate host networking to standard switches, you must migrate the network adapters in stages. For example, remove physical NICs on the hosts from the distributed switch by leaving one physical NIC on every host connected to the switch to keep the network connectivity up. Next, attach the physical NICs to the standard switches and migrate VMkernel adapters and virtual machine network adapters to the switches. Lastly, migrate the physical NIC that you left connected to the distributed switch to the standard switches.

Tasks for Managing Host Networking on a vSphere Distributed Switch

You can add new hosts to a vSphere Distributed Switch, connect network adapters to the switch, and remove hosts from the switch. In a production environment, you might need to keep the network connectivity up for virtual machines and VMkernel services while you manage host networking on the distributed switch.

Adding Hosts to a vSphere Distributed Switch

Consider preparing your environment before you add new hosts to a distributed switch.

- Create distributed port groups for virtual machine networking.
- Create distributed port groups for VMkernel services. For example, create distributed port groups for management network, vMotion, and Fault Tolerance.
- Configure enough uplinks on the distributed switch for all physical NICs that you want to connect to the switch. For example, if the hosts that you want to connect to the distributed switch have eight physical NICs each, configure eight uplinks on the distributed switch.
- Make sure that the configuration of the distributed switch is prepared for services with specific networking requirements. For example, iSCSI has specific requirements for the teaming and failover configuration of the distributed port group where you connect the iSCSI VMkernel adapter.

You can use the **Add and Manage Hosts** wizard in the to add multiple hosts at a time.

Managing Network Adapters on a vSphere Distributed Switch

After you add hosts to a distributed switch, you can connect physical NICs to uplinks on the switch, configure virtual machine network adapters, and manage VMkernel networking.

If some hosts on a distributed switch are associated to other switches in your data center, you can migrate network adapters to or from the distributed switch.

If you migrate virtual machine network adapters or VMkernel adapters, make sure that the destination distributed port groups have at least one active uplink, and the uplink is connected to a physical NIC on the hosts. Another approach is to migrate physical NICs, virtual network adapters, and VMkernel adapters simultaneously.

If you migrate physical NICs, leave at least one active NIC that handles the traffic of port groups. For example, if *vmnic0* and *vmnic1* handle the traffic of the *VM Network* port group, migrate *vmnic0* and leave *vmnic1* connected to the group.

Removing Hosts from a vSphere Distributed Switch

Before you remove hosts from a distributed switch, you must migrate the network adapters that are in use to a different switch.

- To add hosts to a different distributed switch, you can use the **Add and Manage Hosts** wizard to migrate the network adapters on the hosts to the new switch all together. You can then remove the hosts safely from their current distributed switch.
- To migrate host networking to standard switches, you must migrate the network adapters in stages. For example, remove physical NICs on the hosts from the distributed switch by leaving one physical NIC on every host connected to the switch to keep the network connectivity up. Next, attach the physical NICs to the standard switches and migrate VMkernel adapters and virtual machine network adapters to the switches. Lastly, migrate the physical NIC that you left connected to the distributed switch to the standard switches.

Add Hosts to a vSphere Distributed Switch

Learn to manage the networking of your vSphere environment by using a vSphere Distributed Switch, you must associate hosts with the switch. You connect the physical NICs, VMkernel adapters, and virtual machine network adapters of the hosts to the distributed switch.

- Verify that enough uplinks are available on the distributed switch to assign to the physical NICs that you want to connect to the switch.
- Verify that there is at least one distributed port group on the distributed switch.
- Verify that the distributed port group have active uplinks configured in its teaming and failover policy.

If you migrate or create VMkernel adapters for iSCSI, verify that the teaming and failover policy of the target distributed port group meets the requirements for iSCSI:

- Verify that only one uplink is active, the standby list is empty, and the rest of the uplinks are unused.

- Verify that only one physical NIC per host is assigned to the active uplink.
1. On the vSphere Client home page, click **Networking** and navigate to the distributed switch.
 2. Right-click the distributed switch and select **Add and Manage Hosts**.
 3. On the Select task page, select **Add hosts**, and click **Next**.
 4. On the Select hosts page,
 - a) Select the host(s) from the list of available hosts under **All hosts**.
 - b) To view the selected hosts, click **Selected**.
The selected hosts are displayed.
 - c) To filter the hosts based on their compatibility, click **Compatibility**.
Note: While adding hosts to the vSphere Distributed Switch with Network Offloads compatibility, you can only add compatible adapters that are backed by compatible DPU.
 - d) To select all the available hosts, click **Select All**.
 5. Click **Next**.
 6. On the **Manage physical adapters** page, you can add or remove network adapters to the distributed switch by assigning or unassigning an uplink.
 7. To manage adapters on all hosts that have the same physical network adapter, select **Adapters on all hosts**.
 - a) Click **Select All** to select all hosts.
 - b) To assign an uplink to the host, select an uplink from the drop-down menu.
 - c) To unassign an uplink from the hosts, select **None** from the drop-down menu.
 - d) To see more details about the hosts, expand the network adapter listed under **Physical network adapters**.
 - e) You can view the switches that use this VMkernel adapter in **In use by switch**.
For instance, if you assign *uplink1* to *vmnic1*, it is assigned to all the hosts that have *vmnic1* as its physical network adapter.
 8. To manage adapters per host, select **Adapters per host**.
 - a) Select the individual host from the list.
 - b) To assign an uplink to the host, select an uplink from the drop-down menu.
 - c) To unassign an uplink from the host, select **None** from the drop-down menu.

If you select physical NICs that are assigned to other standard or distributed switches, the NICs are migrated to the current distributed switch.

For consistent network configuration, you can connect one and the same physical NIC on every host to the same uplink on the distributed switch.

For instance, if you are adding two hosts connect *vmnic1* on each host to *Uplink1* on the distributed switch.
 9. Click **Next**.
Note: If a host does not have an assigned physical network adapter, then a warning appears.

10. On the **Manage VMkernel adapters** page, you can manage VMkernel adapters to the distributed switch.
11. To manage VMkernel adapters on all hosts that have the same VMkernel adapter, select **Adapters on all hosts**.
 - a) To select all the hosts, click **Select All**.
 - b) Click **Assign Port Group**.
You can see all the available port groups.
 - c) To assign a port group, click **Assign**.
 - d) To unassign a port group, click **Unassign**.
 - e) You can view the switches that use this VMkernel adapter in **In use by switch**.
 - f) To see more details about the hosts, expand the VMkernel adapter listed under **Name**.
For instance, if you assign *DPortGroup1* to *vmk0*, the port group is assigned to all the hosts that has *vmk0* as its VMkernel network adapter.
12. To manage VMkernel adapter per host, select **Adapters per host**.
 - a) Select the individual host from the list.
 - b) Click **Assign Port Group**.
You can see all the available port groups.
 - c) To assign a port group, click **Assign**.
 - d) To unassign a port group, click **Unassign**.
13. Click **Next**.
14. On the **Migrate VM networking** page, select the check box **Migrate virtual machine networking** to migrate virtual machines to a distributed switch.
15. To configure per network adapter, click **Assign port group**.
 - a) To assign a port group, click **Assign**.
For instance, the port group is assigned to all the virtual machines that have the same network adapter.
 - b) To unassign port groups, click **Unassign**.
16. To configure per virtual machine, click **Assign port group**.
 - a) To assign a port group, click **Assign**.
 - b) To unassign a port group, click **Unassign**.
17. Click **Next**.
18. On the **Ready to Complete** page of the **Add and Manage Hosts** wizard, review the settings for the virtual machine.
19. Click **Finish**.

You have now successfully added a host to the vSphere distributed switch.

Having hosts associated with the distributed switch, you can manage physical network adapters, VMkernel adapters, and virtual machine network adapters.

Configure Physical Network Adapters on a vSphere Distributed Switch

Learn how to assign physical NICs to uplinks on the switch for the hosts that are associated with a distributed switch. You can configure physical NICs on the distributed switch for multiple hosts at a time.

For consistent networking configuration throughout all hosts, you can assign the same physical NIC on every host to the same uplink on the distributed switch. For example, you can assign *vmnic1* from hosts *ESXi A* and *ESXi B* to *Uplink 1*.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Add and Manage Hosts**.
3. On the Select task page, select **Manage host networking** and click **Next**.
4. On the Select hosts page, under **Member Hosts**, select from the hosts that are associated with the distributed switch.
5. Click **Next**.
6. On the **Manage physical adapters** page, you can add or remove network adapters to the distributed switch by assigning or unassigning an uplink.
7. To manage adapters on all hosts that have the same physical network adapter, select **Adapters on all hosts**.

The user can see -- while selecting an option in the uplinks drop down menu. This option indicates that not all the hosts with the same network adapter can be assigned to the same uplink.

For example, if *vmnic 0* is assigned to *Uplink 1* in *ESXi A*, and to *Uplink 2* in *ESXi B* then the -- option will be selected in the drop down menu. Selecting this option and moving to the next operation does not change any configuration.

- a) To assign an uplink to the host, select an uplink from the drop-down menu.
 - b) To unassign an uplink from the hosts, select **None** from the drop-down menu
 - c) You can view the switches that use this VMkernel adapter in **In use by switch**.
 - d) To see more details about the hosts, expand the network adapter listed under **Physical network adapters**.
8. To manage adapters per host, select **Adapters per host**.
 - a) Select the individual host from the list.
 - b) To assign an uplink to the host, select an uplink from the drop-down menu.
 - c) To unassign an uplink from the host, select **None** from the drop-down menu.

Note: While adding hosts to the vSphere Distributed Switch with network offloads compatibility, you can only add a compatible adapter that is supported by a compatible DPU.

If you select physical NICs that are already assigned to other standard or distributed switches, the NICs are migrated to the current distributed switch.

For consistent network configuration, you can connect one and the same physical NIC on every host to the same uplink on the distributed switch.

For example, if you are adding two hosts connect *vmnic1* on each host to *Uplink1* on the distributed switch.

9. Click **Next**.

Note: If a host does not have an assigned physical network adapter, then a warning appears.

10. On the **Manage VMkernel adapters** page, you can add VMkernel adapters to the distributed switch.
11. On the **Migrate VM networking** page, select the check box **Migrate virtual machine networking** to migrate virtual machines to a distributed switch.
12. Click **Next**.
13. On the **Ready to Complete** page of the **Add and Manage Hosts** wizard, review the settings for the virtual machine.
14. Click **Finish**.

Migrate VMkernel Adapters to a vSphere Distributed Switch

Learn how to migrate VMkernel adapters to a distributed switch if you want to handle the traffic for VMkernel services by using only this switch when you no longer need the adapters on other standard or distributed switches.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Add and Manage Hosts**.
3. On the Select task page, select **Manage host networking** and click **Next**.
4. On the Select hosts page, under **Member Hosts**, select from the hosts that are associated with the distributed switch.
5. Click **Next**.
6. On the **Manage physical adapters** page, you can add or remove network adapters to the distributed switch by assigning or unassigning an uplink.
7. Click **Next**.

Note: If a host does not have an assigned physical network adapter, then a warning appears.

8. On the **Manage VMkernel adapters** page, you can manage VMkernel adapters to the distributed switch.
9. To manage VMkernel adapters on all hosts that have the same VMkernel adapter, select **Adapters on all hosts**.
 - a) Click **Assign Port Group**
You can see all the available port groups.
 - b) To assign a port group, click **Assign**.
 - c) To unassign a port group, click **Unassign**.

For instance, if you assign *DPortGroup1* to *vmk0*, the port group is assigned to all the hosts that has *vmk0* as its VMkernel network adapter.

10. To see more details about the hosts, expand the VMkernel adapter listed under **Name**.
11. You can view the switches that use this VMkernel adapter in **In use by switch**.
12. To manager VMkernel adapter per host, select **Adapters per host**.
 - a) Select the individual host from the list.
 - b) Click **Assign Port Group**.
You can see all the available port groups.
 - c) To assign a port group, click **Assign**.
 - d) To unassign a port group, click **Unassign**.
13. Click **Next**.
14. On the **Migrate VM networking** page, select the check box **Migrate virtual machine networking** to migrate virtual machines to a distributed switch.
15. Click **Next**.
16. On the **Ready to Complete** page of the **Add and Manage Hosts** wizard, review the settings for the virtual machine.
17. Click **Finish**.

Create a VMkernel Adapter on a vSphere Distributed Switch

Learn how to create a VMkernel adapter on hosts associated with a distributed switch to provide network connectivity to the hosts and to handle the traffic for vSphere vMotion, IP storage, Fault Tolerance logging, and vSAN.

You should dedicate one distributed port group for each VMkernel adapter. One VMkernel adapter should handle only one traffic type.

1. On the vSphere Client Home page, click **Networking** and navigate to a distributed port group.
2. From the **Actions** menu, select **Add VMkernel Adapters**.
3. On the Select hosts page, click **Attached hosts**, select from the hosts that are associated with the distributed switch and click **OK**.
4. Click **Next**.
5. On the Configure VMkernel adapter page, configure the settings for the VMkernel adapter.

Option	Description
Network label	The network label is inherited from the label of the distributed port group.
IP settings	Select IPv4, IPv6, or both. Note: The IPv6 option does not appear on hosts that do not have IPv6 enabled.
MTU	Choose whether to get MTU for the network adapter from the switch or to set a custom size. You cannot set the MTU size to a value greater than 9000 bytes.
TCP/IP stack	Select a TCP/IP stack from the list. Once you set a TCP/IP stack for the VMkernel adapter, you cannot change it later. If you select the vMotion or the Provisioning TCP/IP stack, you will be able to use only these stacks to handle vMotion or Provisioning traffic on the host. All VMkernel adapters for vMotion on the default TCP/IP stack are deactivated for future vMotion sessions. If you set the Provisioning TCP/IP stack, VMkernel adapters on the default TCP/IP stack are deactivated

Option	Description
	for operations that include Provisioning traffic, such as virtual machine cold migration, cloning, and snapshot migration.
Available services	You can enable services for the default TCP/IP stack on the host. Select from the available services: • vMotion. Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. The migration with vMotion to the selected host is not possible if the vMotion service is not enabled for any VMkernel adapter on the default TCP/IP stack, or there are no adapters using the vMotion TCP/IP stack. • Provisioning. Handles the data transferred for virtual machine cold migration, cloning, and snapshot migration. • Fault Tolerance logging. Enables Fault Tolerance logging on the host. You can use only one VMkernel adapter for FT traffic per host. • Management. Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software is installed. You can create another VMkernel adapter for management traffic on the host to provide redundancy. • vSphere Replication. Handles the outgoing replication data that is sent from the source ESXi host to the vSphere Replication server. • vSphere Replication NFC. Handles the incoming replication data on the target replication site. • vSAN. Enables the vSAN traffic on the host. Every host that is part of a vSAN cluster must have such a VMkernel adapter. • vSphere Backup NFC. VMkernel port setting for dedicated backup NFC traffic. NFC traffic goes through the VMkernel Adapter when vSphereBackup NFC service is enabled. • NVMe over TCP. VMkernel port setting for dedicated NVMe over TCP storage traffic. NVMe over TCP storage traffic goes through the VMkernel Adapter when NVMe over TCP adapter is enabled. For more information, see <i>vSphere Storage Guide</i>. • NVMe over RDMA. VMkernel port setting for dedicated NVMe over RDMA storage traffic. NVMe over RDMA storage traffic goes through the VMkernel Adapter when NVMe over RDMA adapter is enabled. For more information, see <i>vSphere Storage Guide</i>.

6. On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IPv4 settings automatically	Use DHCP to obtain IP settings. A DHCP server must be present on the network.
Use static IPv4 settings	Enter the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack. To change the VMkernel default gateway, select Configure on VMkernel adapters or Configure on TCP/IP stack and enter a gateway address.

7. On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses. In ESXi 6.5 and later router advertisement is enabled by default and supports the M and O flags in accordance with RFC 4861.
Static IPv6 addresses	1. Enter the IPv6 address and subnet prefix length. 2. To change the VMkernel default gateway, select Configure on VMkernel adapters or Configure on TCP/IP stack and enter a gateway address. The VMkernel Default Gateway address for IPv6 is obtained from the selected TCP/IP stack.

8. Review your settings selections on the Ready to complete page and click **Finish**.

Migrate Virtual Machine Networking to the vSphere Distributed Switch

Learn how to manage virtual machine networking by using a distributed switch, migrate virtual machine network adapters to labeled networks on the switch.

Verify that at least one distributed port group intended for virtual machine networking exists on the distributed switch.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the Select task page, select **Manage host networking** and click **Next**.
3. On the Select hosts page, under **Member Hosts**, select from the hosts that are associated with the distributed switch.
4. Click **Next**.
5. On the **Manage physical adapters** page, you can add or remove network adapters to the distributed switch by assigning or unassigning an uplink.
6. Click **Next**.

Note: If a host does not have an assigned physical network adapter, then a warning appears.

7. On the **Manage VMkernel adapters** page, you can add VMkernel adapters to the distributed switch.
8. On the **Migrate VM networking** page, select the check box **Migrate virtual machine networking** to migrate virtual machines to a distributed switch.
9. To configure per network adapter, click **Assign port group**.
 - a) To assign a port group, click **Assign**.

For instance, the port group is assigned to all the virtual machines that have the same network adapter.
 - b) To unassign a port group, click **Unassign**.
10. To configure per virtual machine, click **Assign port group**.
 - a) To assign a port group, click **Assign**.
 - b) To unassign a port group, click **Unassign**.
11. Click **Next**.
12. On the **Ready to Complete** page of the **Add and Manage Hosts** wizard, review the settings for the virtual machine.
13. Click **Finish**.

Remove Hosts from a vSphere Distributed Switch

Learn how to remove hosts from a vSphere distributed switch if you have configured a different switch for the hosts.

- Verify that physical NICs on the target hosts are migrated to a different switch.
- Verify that VMkernel adapters on the hosts are migrated to a different switch.
- Verify that virtual machine network adapters are migrated to a different switch.

For details about migrating network adapters to different switches, see [Tasks for Managing Host Networking on a vSphere Distributed Switch](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. From the **Actions** menu, select **Add and Manage Hosts**.
3. On the Select task page, select **Remove hosts** and click **Next**.
4. To remove an individual host from the distributed switch, select the host.
5. To remove all the hosts from the distributed switch, click **Select All**.
6. Click **Next**.
7. Click **Finish**.

Managing Networking on Host Proxy Switches

Learn how to change the configuration of the proxy switch on every host that is associated with a vSphere distributed switch. You can manage physical NICs, VMkernel adapters, and virtual machine network adapters.

For details about setting up VMkernel networking on host proxy switches, see [Create a VMkernel Adapter on a vSphere Distributed Switch](#).

Migrate Network Adapters on a Host to a vSphere Distributed Switch

Learn how to migrate network adapters from a standard switch to the distributed switch for hosts associated with a distributed switch. You can migrate physical NICs, VMkernel adapters, and virtual machine network adapters at the same time.

To migrate virtual machine network adapters or VMkernel adapters, make sure that the destination distributed port groups have at least one active uplink, and the uplink is connected to a physical NIC on this host. Alternatively, migrate physical NICs, virtual network adapters, and VMkernel adapters at once.

To migrate physical NICs, make sure that the source port groups on the standard switch have at least one physical NIC to handle their traffic. For example, if you migrate a physical NIC that is assigned to a port group for virtual machine networking, make sure that the port group is connected to at least one physical NIC. Otherwise the virtual machines on same VLAN on the standard switch will have connectivity between each other but not to the external network.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the destination distributed switch and click the horizontal ellipsis icon next to **Manage physical adapters**.
4. Select **Migrate Networking**.
5. Configure physical NICs.
 - a) From the **On other switches/unclaimed** list, select a physical NIC and click **Assign uplink**.
 - b) Select an uplink and click **OK**.
 - c) Click **Next**.
6. Configure VMkernel adapters.
 - a) Select an adapter and click **Assign port group**.
 - b) Select a distributed port group and click **OK**.
You should connect one VMkernel adapter to one distributed port group at a time.
 - c) Click **Next**.
7. Configure virtual machine network adapters.
 - a) Select the check box **Migrate virtual machine networking**.
 - b) Select a virtual machine or a virtual machine network adapter and click **Assign port group**.
If you select a virtual machine, you migrate all network adapters on the virtual machine. If you select a network adapter, you migrate only this network adapter.
 - c) Select a distributed port group from the list and click **OK**.
 - d) Click **Next**.
8. On the Ready to complete page, review the new networking configuration and click **Finish**.

Migrate a VMkernel Adapter on a Host to a vSphere Standard Switch

Learn how to migrate VMkernel adapters from the distributed to a standard switch if a host is associated with a distributed switch.

Verify that the destination standard switch has at least one physical NIC.

For details about creating VMkernel adapters on a vSphere distributed switch, see [Create a VMkernel Adapter on a vSphere Distributed Switch](#).

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the destination standard switch from the list.
4. Click **Migrate VMkernel Adapter**.
5. On the Select VMkernel adapter page, select the virtual network adapter to migrate to the standard switch from the list.
6. On the Configure settings page, edit the **Network label** and **VLAN ID** for the network adapter.
7. On the Ready to complete page, review the migration details and click **Finish**.
Click **Back** to edit settings.

Assign a Physical NIC of a Host to a vSphere Distributed Switch

Learn how to assign physical NICs of a host that is associated with a distributed switch to uplink port on the host proxy switch.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select a distributed switch from the list.
4. Click **Manage Physical Adapters**.
5. Select a free uplink from the list and click **Add adapters**.
6. Select a physical NIC and click **OK**.

Remove a Physical NIC from a vSphere Distributed Switch

Learn how to remove a physical NIC of a host from an uplink on a vSphere distributed switch.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the distributed switch.
4. Click **Manage Physical Adapters**.
5. Select an uplink and click **Remove selected**.
6. Click **OK**.

When you remove physical NICs from active virtual machines, you might see the NICs you removed reported in the . See [Removing NICs from Active Virtual Machines](#).

Removing NICs from Active Virtual Machines

Learn how to remove NICs from active virtual machines, and see the NICs you have removed in the vSphere Client.

Removing NICs from an Active Virtual Machine without a Guest Operating System Installed

You cannot remove NICs from an active virtual machine on which no operating system is installed.

The vSphere Client might report that the NIC has been removed, but you continue to see it attached to the virtual machine.

Removing NICs from an Active Virtual Machine with a Guest Operating System Installed

You can remove a NIC from an active virtual machine, but it might not be reported to the vSphere Client for some time. If you click **Edit Settings** for the virtual machine, you might see the removed NIC listed even after the task is complete. The Edit Settings dialog box for the virtual machine does not immediately display the removed NIC.

You might also see the NIC attached to the virtual machine if the guest operating system of the virtual machine does not support hot removal of NICs.

Distributed Port Groups

Learn how a distributed port group specifies port configuration options for each member port on a vSphere distributed switch. Distributed port groups define how a connection is made to a network.

Add a Distributed Port Group

Learn how to create a distributed switch network for your virtual machines, and to associate VMkernel adapters, and add a distributed port group to a vSphere Distributed Switch.

Related to adding a port group, is applying VLAN tagging globally on all distributed ports. Using the VLAN options you can select VLAN tags. To learn more, see [Configure VLAN Tagging on a Distributed Port Group or Distributed Port](#)

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Distributed port group > New distributed port group**.
3. On the Name and location page, enter the name of the new distributed port group, or accept the generated name, and click **Next**.


Figure 73: New Distributed Port Group - Name and Location

4. On the Configure settings page, set the general properties for the new distributed port group.

Setting	Description
Port binding	Select the ports that are assigned to virtual machines connected to this distributed port group. • Static binding: Assign a port to a virtual machine when the virtual machine connects to the distributed port group. • Ephemeral - no binding: No port binding. You can assign a virtual machine to a distributed port group with ephemeral port binding also when connected to the host.

Setting	Description
Port allocation	Elastic: The default number of ports is eight. When all the ports are assigned, a new set of eight ports is created. Fixed: The default number of ports is set to eight. When all the ports are assigned, no additional ports are created.
Number of ports	Enter the number of ports on the distributed port group.
Network resource pool	To assign the new distributed port group to a user-defined network resource pool, use the drop-down menu. If you have not created a network resource pool, this menu is empty. Note: You cannot assign Network Resource Pool if Network Offloads is enabled.
VLAN	Use the VLAN type drop-down menu to specify the type of VLAN traffic filtering and marking: None: Do not use VLAN. Select None if you are using External Switch Tagging. VLAN: In the VLAN ID text box, enter a number between 1 and 4094 for Virtual Switch Tagging. VLAN trunking: Enter a VLAN trunk range. Pass VLAN traffic with an ID to the guest OS. You can set multiple ranges and individual VLANs by using a comma-separated list. For example: 1702–1705, 1848–1849. Use this option for Virtual Guest Tagging. Private VLAN: Associate the traffic with a private VLAN created on the distributed switch. If you did not create any private VLANs, this menu is empty.
Advanced	To customize the policy configurations for the new distributed port group, select this check box.

Figure 74: New Distributed Port Group - Configure settings




- Click **Next**.
- Optional: On the Security page, edit the security exceptions and click **Next**.

Setting	Description
Promiscuous mode	Reject: Placing an adapter in promiscuous mode from the guest operating system does not result in receiving frames for other virtual machines. Accept: If an adapter is placed in promiscuous mode from the guest operating system, the switch allows the guest adapter to receive all frames passed on the switch in

Setting	Description
	compliance with the active VLAN policy for the port where the adapter is connected. Firewalls, port scanners, intrusion detection systems, and so on, must run in promiscuous mode.
MAC address changes	The MAC address change feature allows a VM to change its MAC address. A VM connected to a port can run an administrative command to change the MAC address of its vNIC and still send and receive traffic on that vNIC. • Reject: If the option is set to Reject and the guest OS changes the MAC address of the adapter to a value different from the address in the <code>.vmx</code> configuration file, then the switch drops all inbound frames to the virtual machine adapter. - If the guest OS changes the MAC address back, the virtual machine receives frames again. • Accept: If the guest OS changes the MAC address of a network adapter, the adapter receives frames to its new address.
Forged Transmits	• Reject: The switch drops any outbound frame with a source MAC address that is different from the one in the <code>.vmx</code> configuration file. • Accept: The switch does not perform filtering and permits all outbound frames.

Figure 75: New Distributed Port Group - Security



7. Optional: On the Security page, edit the MAC Learning policy and click **Next**.

Setting	Description
Status	Enable or disable the MAC learning feature. The default is disabled.
Allow unicast flooding	When a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, if MAC learning is enabled.
MAC Limit	The number of MAC addresses that can be learned is configurable. The maximum value is 4096 per port, which is the default.

Setting	Description
MAC Limit Policy	The policy for when the MAC limit is reached. The options are: • Drop - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled. • Allow - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

8. Optional: On the Traffic shaping page, enable or disable Ingress or Egress traffic shaping and click **Next**.

Setting	Description
Status	If you enable either Ingress traffic shaping or Egress traffic shaping, you are setting limits on the amount of networking bandwidth allocated for each virtual adapter associated with this particular port group. If you disable the policy, services have a free, clear connection to the physical network by default. Note: You cannot assign traffic shaping policies if Network Offloads Compatibility enabled.
Average bandwidth	This feature establishes the number of bits per second to allow across a port, averaged over time. It is the allowed average load.
Peak bandwidth	The maximum number of bits per second to allow across a port when it is sending and receiving a burst of traffic. It tops the bandwidth used by a port whenever it is using its burst bonus.
Burst size	The maximum number of bytes to allow in a burst. If this parameter is set, a port can gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by Average bandwidth , it can temporarily transmit data at a faster speed if a burst bonus is available. This parameter tops the number of bytes that can be

Setting	Description
	accumulated in the burst bonus and as a result transferred at a faster speed.

Figure 76: New Distributed Port Group - Traffic shaping

9. Optional: On the Teaming and failover page, edit the settings and click **Next**.

Setting	Description
Load balancing	Specify the way an uplink is selected. • Route based on originating virtual port: Select an uplink based on the virtual port where the traffic entered the distributed switch. • Route based on IP hash: Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash. • Route based on source MAC hash: Select an uplink based on a hash of the source Ethernet. • Route based on physical NIC load: Select an uplink based on the current loads of physical NICs. • Use explicit failover order: Always use the highest order uplink from the list of Active adapters which passes failover detection criteria. Note: IP-based teaming requires that the physical switch is configured with EtherChannel. For all other options, disable EtherChannel.
Network failure detection	Specify the method to use for failover detection. • Link status only: Relies solely on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch. • Beacon probing: Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. This detects many of the failures previously mentioned that are not detected by link status alone. Note: Do not use beacon probing with IP-hash load-balancing.

Setting	Description
Notify switches	Select Yes or No to notify switches in case of failover. If you select Yes, whenever a virtual NIC is connected to the distributed switch or whenever that virtual NIC's traffic can be routed over a different physical NIC in the team because of a failover event, a notification is sent out over the network to update the lookup tables on physical switches. In almost all cases, this process is desirable for the lowest latency of failover occurrences and migrations with vMotion. If the Notify switches is set to Yes, then all the connected ports, port groups, and distributed switches are connected back to the host when vCenter Server reconnects with ESXi hosts. Note: Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode. No such issue exists with NLB running in multicast mode.
Failback	Select Yes or No to disable or enable failback. This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to Yes (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. If failback is set to No, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.
Failover order	Specify how to distribute the workload for uplinks. To use some uplinks but reserve others for emergencies if the uplinks in use fail, set this condition by moving them into different groups: • Active uplinks: Continue to use the uplink when the network adapter connectivity is up and active. • Standby uplinks : Use this uplink if one of the active adapters' connectivity is down. • Unused uplinks : Do not use this uplink.

Setting	Description
	Note: When using IP-hash load-balancing, do not configure standby uplinks.

Figure 77: New Distributed Port Group - Teaming and failover

New Distributed Port Group

Teaming and failover
Controls load balancing, network failure detection, switches notification, failback, and uplink failover order.

Load balancing
Route based on IP hash

Network failure detection
Link status only

Notify switches
Yes

Failback
Yes

Failover order

Active uplinks

- Uplink-1
- Uplink-2
- Uplink-3
- Uplink-4

Standby uplinks

Unused uplinks


CANCEL BACK NEXT

10. Optional: On the Monitoring page, enable or disable NetFlow and click **Next**.

Setting	Description
Disabled	NetFlow is disabled on the distributed port group.

Setting	Description
Enabled	NetFlow is enabled on the distributed port group. NetFlow settings can be configured at the vSphere Distributed Switch level.


Figure 78: New Distributed Port Group - Monitoring



11. Optional: On the Miscellaneous page, select **Yes** or **No** and click **Next**.

Selecting **Yes** shuts down all ports in the port group. This action can disrupt the normal network operations of the hosts or virtual machines using the ports.

Figure 79: New Distributed Port Group - Miscellaneous



12. On the Ready to complete page, review your settings and click **Finish**.
To change any settings, click the **Back** button.

Figure 80: New Distributed Port Group - Ready to complete

New Distributed Port Group

Ready to complete

Review the changes before proceeding.

Property	Value
Distributed port group name	DPortGroup 1
Port binding	Static binding
Number of ports	8
Port allocation	Elastic (default)
Network resource pool	(default)
VLAN ID	---

CANCEL BACK FINISH

Edit General Distributed Port Group Settings

Learn to edit general distributed port group settings such as the distributed port group name, port settings and network resource pool.

- Locate a distributed port group in the vSphere Client.
 - Select a distributed switch and click the **Networks** tab.
 - Click **Distributed Port Groups**.
- Right-click the distributed port group and select **Edit Settings**.
- Select **General** to edit the following distributed port group settings.

Option	Description
Name	The name of distributed port group. You can edit the name in the text field.
Port binding	Choose when ports are assigned to virtual machines connected to this distributed port group. Static binding: Assign a port to a virtual machine when the virtual machine connects to the distributed port group. Ephemeral: No port binding. You can also assign a virtual machine to a distributed port group with ephemeral port binding when connected to the host.
Port allocation	Elastic: The default number of ports is set to eight. When all ports are assigned, a new set of eight ports is created. This is the default. Fixed: The default number of ports is set to eight. No additional ports are created when all ports are assigned.
Number of ports	Enter the number of ports on the distributed port group.
Network resource pool	Use the drop-down menu to assign the new distributed port group to a user-defined network resource pool. If you have not created a network resource pool, this menu is empty.
Description	Enter any information about the distributed port group in the description field.

- Click **OK**.

Remove a Distributed Port Group

Learn how to remove a distributed port group when you no longer need the corresponding labeled network to provide connectivity and configure connection settings for virtual machines or VMkernel networking.

- Verify that all virtual machines connected to the corresponding labeled network are migrated to a different labeled network.
 - Verify that all VMkernel adapters connected to the distributed port group are migrated to a different port group, or are deleted.
1. Locate a distributed port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups**.
 2. Select the distributed port group.
 3. From the **Actions** menu, select **Delete**.

Working with Distributed Ports

A distributed port is a port on a vSphere distributed switch that connects to the VMkernel or to a virtual machine's network adapter.

Default distributed port configuration is determined by the distributed port group settings, but some settings for individual distributed ports can be overridden.

Monitor the State of Distributed Ports in the vSphere Client

Learn how vSphere monitors distributed ports and provides information about the current state and runtime statistics of each port. Starting vSphere 8.0 Update 3, all the ports are displayed, you can filter all the columns in the ports list.

1. In the vSphere Client Home page, click **Networking** and select the distributed switch.
2. Click the **Ports** tab and select a port from the list.
All the ports are displayed on the grid.
The ports table for the distributed switch displays runtime statistics for each distributed port.
3. Click the funnel icon next to **Port ID** to filter ports by Port ID.
4. Click the funnel icon next to **Name** to filter ports by the name of the port.
5. Click the funnel icon next to **Connectee** column to filter based on the connectee for each distributed port.

Option	Description
Physical Adapter	Search is based on the connected physical adapter.
VMkernel Adapter	Search is based on the connected host VmkVnic.
Virtual Machine	Search is based on the connected VM.
Host	Search is based on the connected host.
--	Search based on the connectee of this distributed port is currently unavailable.

6. Click the funnel icon next to **Runtime MAC Address** to filter ports by the MAC address of distributed port.
7. Click the funnel icon next to **Port Group** to filter ports by the name of the port group.
8. Click the funnel icon next to **State** column to filter based on the uplink connection for each distributed port.

The **State** column displays the current state for each distributed port.

Option	Description
Link Up	The link for this distributed port is up.
Link Down	The link for this distributed port is down.
Blocked	This distributed port is blocked.
--	The state of this distributed port is currently unavailable.

9. Click the funnel icon next to **VLAN ID** column to filter based on the connectee for each distributed port.

Option	Description
VLAN	Search results are based on the VLAN ID.
Private VLAN	Search results are based on the Private VLAN that is identified by its primary VLAN ID
VLAN Trunk	Search results are based on the VLAN traffic with ID within the VLAN trunk range
--	Search is based on the VLAN ID of this distributed port that is currently unavailable.

10. Click the funnel icon next to **VIF ID** to filter ports by the virtual interface identifier.

Configure Distributed Port Settings in vSphere Client

Learn how to configure general distributed port settings such as the port name and description.

- Locate a distributed port group in the vSphere Client.
 - Select a distributed switch and click the **Networks** tab.
 - Click **Distributed Port Groups**.
- Click a distributed port group from the list.
- Click the **Ports** tab, and select a distributed port from the table.
Information about the distributed port appears at the bottom of the screen.
- Click the **Edit settings** icon.
- On the Properties page and policy pages, edit information about the distributed port and click **OK**.
If overrides are not allowed, the policy options are disabled.

You can allow overrides at the port level by changing the **Advanced** settings of the distributed port group. See [Configure Overriding Networking Policies on Port Level](#).

Configure Virtual Machine Networking on a vSphere Distributed Switch

Connect virtual machines to a vSphere distributed switch either by configuring an individual virtual machine NIC or migrating groups of virtual machines from the vSphere distributed switch itself.

Connect virtual machines to vSphere distributed switches by connecting their associated virtual network adapters to distributed port groups. You can do this either for an individual virtual machine by modifying the virtual machine's network adapter configuration, or for a group of virtual machines by migrating virtual machines from an existing virtual network to a vSphere distributed switch.

Migrate Virtual Machines to or from a vSphere Distributed Switch

In addition to connecting virtual machines to a distributed switch at the individual virtual machine level, learn to migrate a group of virtual machines between a vSphere Distributed Switch network and a vSphere Standard Switch network.

- In the vSphere Client, navigate to a data center.
- Right-click the data center in the navigator and select **Migrate VMs to Another Network**.
- Select a source network.
 - Select **Specific network** and use the **Browse** button to select a specific source network.
 - Select **No network** to migrate all virtual machine network adapters that are not connected to any other network.

4. Use **Browse** to select a destination network and click **Next**.
5. Select virtual machines from the list to migrate from the source network to the destination network and click **Next**.
6. Review your selections and click **Finish**.
Click **Back** to edit any selections.

Connect an Individual Virtual Machine to a Distributed Port Group

Learn how to connect an individual virtual machine to a vSphere Distributed Switch by modifying the NIC configuration of the virtual machine.

1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
2. From the Actions menu, select Edit Settings.
3. Expand the **Network adapter** section and select **Browse** from the **Network adapter** drop-down menu.
4. In the Select Network dialog box, select a distributed port group and click **OK**.
5. Click **OK**.

Topology of a vSphere Distributed Switch

The topology diagrams of a vSphere Distributed Switch in the vSphere Client show the structure of virtual machine adapters, VMkernel adapters, and physical adapters in the switch.

You can examine the components, arranged in port groups, whose traffic is handled by the switch, and the connections between them. The diagram displays information about the physical adapter that connects the virtual adapters to the external network.

You can view the components that are running on the entire distributed switch and on each host participating in it.

Central Topology Diagram

You can use the central topology diagram of the switch to locate and edit the settings for distributed port groups and uplink groups associated with multiple hosts. You can initiate migration of virtual machine adapters from a port group to a destination on the same or different switch. You can also reorganize the hosts and their networking on the switch by using the **Add and Manage Hosts** wizard.

Topology Diagram of a Host Proxy Switch

The topology diagram of a host proxy switch shows the adapters attached to the switch ports on the host. You can edit the settings of the VMkernel and physical adapters.

Topology Diagram of a Network Offloads Switch

The topology diagram of a network offloads switch shows the adapters attached to the switch ports on the host. You can edit the settings of the VMkernel and physical adapters.

Diagram Filters

You can use diagram filters to limit the information displayed in topology diagrams. The default filter limits the topology diagram to display 64 port groups, 96 hosts, and 1024 virtual machines.

You can change the scope of the diagram by using no filters or by applying custom filters. By using a custom filter, you can view information only about a set of virtual machines, a set of port groups on certain hosts, or a port. You can create filters from the central topology diagram of the distributed switch.

View the Topology of a vSphere Distributed Switch

Examine the organization of components that are connected to the distributed switch across the hosts in a vCenter Server.


1. Navigate to the vSphere distributed switch in the vSphere Client.
2. On the **Configure** tab, expand **Settings** and select **Topology**.

By default the diagram shows up to 32 distributed port groups, 32 hosts, and 1024 virtual machines.

Diagram of a Distributed Switch That Connects the VMkernel and Virtual Machines to the Network

In your virtual environment, a vSphere Distributed Switch handles VMkernel adapters for vSphere vMotion and for the management network, and virtual machines grouped. You can use the central topology diagram to examine whether a virtual machine or VMkernel adapter is connected to the external network and to identify the physical adapter that carries the data.

Figure 81: Topology Diagram of a Distributed Switch That Handles VMkernel and Virtual Machine Networking



You can perform the following common tasks in the topology of the distributed switch:

- Use filters to view the networking components only for selected port groups on certain hosts, for selected virtual machines, or for a port.
- Locate, configure and migrate virtual machine networking components across host and port groups by using the **Migrate Virtual Machine Networking** wizard.
- Detect the virtual machine adapters that have no network assigned and move them to the selected port group by using the **Migrate Virtual Machine Networking** wizard.
- Handle networking components on multiple hosts by using the **Add and Manage Hosts** wizard.
- View the physical NIC or NIC team that carries the traffic related to a selected virtual machine adapter or VMkernel adapter.

In this way you can also view the host on which a selected VMkernel adapter resides. Select the adapter, trace the route to the associated physical NIC, and view the IP address or domain name next to the NIC.

- Determine the VLAN mode and ID for a port group. For information about VLAN modes, see [VLAN Configuration](#).

View the Topology of a Host Proxy Switch

Examine and reorganize the networking of the VMkernel and virtual machines that the vSphere Distributed Switch handles on a host.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the distributed switch from the list.

The topology of the host proxy switch appears under the list.

View the Topology of Network Offloads Switch

Examine the organization of the distributed switch with network offloads.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. The  icon indicates that network offloads is supported for this distributed switch.

How do I Set Up VMkernel Networking

Learn how to set up VMkernel adapters to provide network connectivity to hosts and to accommodate system traffic of vMotion, IP storage, Fault Tolerance logging, vSAN, and so on.

VMkernel Networking Layer

The VMkernel networking layer provides connectivity to hosts and handles the standard system traffic of vSphere vMotion, IP storage, Fault Tolerance, vSAN, and others. You can also create VMkernel adapters on the source and target vSphere Replication hosts to isolate the replication data traffic.

Multihoming in the VMkernel networking is defined as multiple VMkernel adapters in a TCP/IP stack.

Note: If there is a single vmnic with multiple TCP/IP stacks, it cannot be defined as multihoming.

Multihoming is supported if the vmknics are in different IP subnets. The admin must ensure that the route entries reflect the intent. The ESXi does not run any routing protocol and it does not support or add dynamic routes. For any specific route entry, it must be added statically from the vCenter or the ESXi host. Configurations with more than one vmknic interface on the same IP subnet in a single network stack instance is not supported. Deploying such a configuration might lead to unexpected results like connectivity issues, low throughput and asymmetric routing. However, there are a few exceptions to this rule. For example, iSCSI, Multi-NIC vMotion, and NSX VTEPS. For more information, see <http://kb.vmware.com/kb/2010877>.

TCP/IP Stacks at the VMkernel Level

Default TCP/IP stack

Provides networking support for the management traffic between vCenter Server and ESXi hosts, and for system traffic such as vMotion, IP storage, Fault Tolerance, and so on.

vMotion TCP/IP stack

Supports the traffic for live migration of virtual machines. Use the vMotion TCP/IP to provide better isolation for the vMotion traffic. After you create a VMkernel adapter on the vMotion TCP/IP stack, you can use only this stack for vMotion on this host. The VMkernel adapters on the default TCP/IP stack are disabled for the vMotion service. If a live migration uses the default TCP/IP stack while you configure VMkernel adapters with the vMotion TCP/IP stack,

Provisioning TCP/IP stack

the migration completes successfully. However, the involved VMkernel adapters on the default TCP/IP stack are disabled for future vMotion sessions.

Supports the traffic for virtual machine cold migration, cloning, and snapshot migration. You can use the provisioning TCP/IP to handle Network File Copy (NFC) traffic during long-distance vMotion. NFC provides a file-specific FTP service for vSphere. ESXi uses NFC for copying and moving data between datastores. VMkernel adapters configured with the provisioning TCP/IP stack handle the traffic from cloning the virtual disks of the migrated virtual machines in long-distance vMotion. By using the provisioning TCP/IP stack, you can isolate the traffic from the cloning operations on a separate gateway. After you configure a VMkernel adapter with the provisioning TCP/IP stack, all adapters on the default TCP/IP stack are disabled for the Provisioning traffic.

Custom TCP/IP stacks

You can add custom TCP/IP stacks at the VMkernel level to handle networking traffic of custom applications.

Mirror TCP/IP stack

You can create a vmknic on mirror TCP/IP stack when you select mirror stack for ERSPAN.

Securing System Traffic

Take appropriate security measures to prevent unauthorized access to the management and system traffic in your vSphere environment. For example, isolate the vMotion traffic in a separate network that includes only the ESXi hosts that participate in the migration. Isolate the management traffic in a network that only network and security administrators can access. For more information, see *vSphere Security* and *vSphere Installation and Setup*.

System Traffic Types

Dedicate a separate VMkernel adapter for every traffic type . For distributed switches, dedicate a separate distributed port group for each VMkernel adapter.

Management traffic

Carries the configuration and management communication for ESXi hosts, vCenter Server, and host-to-host High Availability traffic. By default, when you install the ESXi software, a vSphere Standard switch is created on the host together with a VMkernel adapter for management traffic. To provide redundancy, you can connect two or more physical NICs to a VMkernel adapter for management traffic.

vMotion traffic

Accommodates vMotion. A VMkernel adapter for vMotion is required both on the source and the target hosts. Configure The VMkernel adapters for vMotion to handle only the vMotion traffic. For better performance, you can configure multiple NIC vMotion. To have multi-NIC vMotion, you can dedicate two or more port groups to the vMotion traffic, respectively every port group must have a vMotion VMkernel adapter associated with it. Then you can connect one or more physical NICs to every port group. In this way, multiple physical NICs are used for vMotion, which results in greater bandwidth .

Note: vMotion network traffic is not encrypted. You should provision secure private networks for use by vMotion only.

Provisioning traffic	Handles the data that is transferred for virtual machine cold migration, cloning, and snapshot migration.
IP storage traffic and discovery	Handles the connection for storage types that use standard TCP/IP networks and depend on the VMkernel networking. Such storage types are software iSCSI, dependent hardware iSCSI, and NFS. If you have two or more physical NICs for iSCSI, you can configure iSCSI multipathing. ESXi hosts support NFS 3 and 4.1.
Fault Tolerance traffic	Handles the data that the primary fault tolerant virtual machine sends to the secondary fault tolerant virtual machine over the VMkernel networking layer. A separate VMkernel adapter for Fault Tolerance logging is required on every host that is part of a vSphere HA cluster.
vSphere Replication traffic	Handles the outgoing replication data that the source ESXi host transfers to the vSphere Replication server. Dedicate a VMkernel adapter on the source site to isolate the outgoing replication traffic.
vSphere Replication NFC traffic	Handles the incoming replication data on the target replication site.
vSAN traffic	Every host that participates in a vSAN cluster must have a VMkernel adapter to handle the vSAN traffic.
vSphere Backup NFC	VMkernel port setting for dedicated backup NFC traffic. NFC traffic goes through the VMkernel Adapter when vSphereBackup NFC service is enabled.
NVMe over TCP	VMkernel port setting for dedicated NVMe over TCP storage traffic. NVMe over TCP storage traffic goes through the VMkernel Adapter when NVMe over TCP adapter is enabled. For more information see <i>vSphere Storage Guide</i> .
NVMe over RDMA	VMkernel port setting for dedicated NVMe over RDMA storage traffic. NVMe over RDMA storage traffic goes through the VMkernel Adapter when NVMe over RDMA adapter is enabled. For more information, see <i>vSphere Storage Guide</i> .

View Information About VMkernel Adapters on a Host

You can view each VMkernel adapter's assigned services, associated switch, port settings, IP settings, TCP/IP stack, VLAN ID, and policies.

1. In the vSphere Client, navigate to the host.
2. Click the **Configure** tab and expand the **Networking** menu.
3. To view information about all VMkernel adapters on the host, select **VMkernel adapters**.
4. Select an adapter from the VMkernel adapters list to view its settings.

Tab	Description
All	Displays all configuration information about the VMkernel adapter. This information includes port and NIC settings, IPv4 and IPv6 settings, traffic shaping, teaming and failover, and security policies.
Properties	Displays the port properties and NIC settings of the VMkernel adapter. The port properties include the port group (network label) to which the adapter is associated, the VLAN ID, and the enabled services. The NIC settings include MAC address and the configured MTU size.

Tab	Description
IP Settings	Displays all IPv4 and IPv6 settings for the VMkernel adapter. IPv6 information is not displayed if IPv6 has not been enabled on the host.
Policies	Displays the configured traffic shaping, teaming and failover, and security policies that apply for the port group to which the VMkernel adapter is connected.

Create a VMkernel Adapter on a vSphere Standard Switch

Learn how to create a VMkernel network adapter on a vSphere standard switch to provide network connectivity for hosts and to handle the system traffic for vSphere vMotion, IP storage, Fault Tolerance logging, vSAN, and so on. You can also create VMkernel adapters on the source and target vSphere Replication hosts to isolate the replication data traffic. Dedicate a VMkernel adapter to only one traffic type.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **VMkernel adapters**.
3. Click **Add networking**.
4. On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.
5. On the Select target device page, select either an existing standard switch or select **New standard switch**.
6. Optional: On the Create a Standard Switch page, assign physical NICs to the switch.
You can create the standard switch without physical NICs and configure them later. During the time that no physical NICs are attached to the host, the host does not have network connectivity to the other hosts on the physical network. The virtual machines on the host are able to communicate with each other.
 - a) Click **Add adapters** and select as many physical NICs as you need.
 - b) Use the up and down arrows to configure the active and standby NICs.
7. On the Port properties page, configure the settings for the VMkernel adapter.

Option	Description
Network label	The network label is inherited from the label of the distributed port group.
IP settings	Select IPv4, IPv6, or both. Note: The IPv6 option does not appear on hosts that do not have IPv6 enabled.
MTU	Choose whether to get MTU for the network adapter from the switch or to set a custom size. You cannot set the MTU size to a value greater than 9000 bytes.
TCP/IP stack	Select a TCP/IP stack from the list. Once you set a TCP/IP stack for the VMkernel adapter, you cannot change it later. If you select the vMotion or the Provisioning TCP/IP stack, you will be able to use only these stacks to handle vMotion or Provisioning traffic on the host. All VMkernel adapters for vMotion on the default TCP/IP stack are deactivated for future vMotion sessions. If you set the Provisioning TCP/IP stack, VMkernel adapters on the default TCP/IP stack are deactivated for operations that include Provisioning traffic, such as virtual machine cold migration, cloning, and snapshot migration.
Available services	You can enable services for the default TCP/IP stack on the host. Select from the available services:

Option	Description
	• vMotion. Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. The migration with vMotion to the selected host is not possible if the vMotion service is not enabled for any VMkernel adapter on the default TCP/IP stack, or there are no adapters using the vMotion TCP/IP stack. • Provisioning. Handles the data transferred for virtual machine cold migration, cloning, and snapshot migration. • Fault Tolerance logging. Enables Fault Tolerance logging on the host. You can use only one VMkernel adapter for FT traffic per host. • Management. Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software is installed. You can create another VMkernel adapter for management traffic on the host to provide redundancy. • vSphere Replication. Handles the outgoing replication data that is sent from the source ESXi host to the vSphere Replication server. • vSphere Replication NFC. Handles the incoming replication data on the target replication site. • vSAN. Enables the vSAN traffic on the host. Every host that is part of a vSAN cluster must have such a VMkernel adapter. • vSphere Backup NFC. VMkernel port setting for dedicated backup NFC traffic. NFC traffic goes through the VMkernel Adapter when vSphereBackup NFC service is enabled. • NVMe over TCP. VMkernel port setting for dedicated NVMe over TCP storage traffic. NVMe over TCP storage traffic goes through the VMkernel Adapter when NVMe over TCP adapter is enabled. For more information, see <i>vSphere Storage Guide</i>. • NVMe over RDMA. VMkernel port setting for dedicated NVMe over RDMA storage traffic. NVMe over RDMA storage traffic goes through the VMkernel Adapter when NVMe over RDMA adapter is enabled. For more information, see <i>vSphere Storage Guide</i>.

8. Optional: On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IPv4 settings automatically	Use DHCP to obtain IP settings. A DHCP server must be present on the network.
Use static IPv4 settings	Enter the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack. Select the Override default gateway for this adapter check box and enter a gateway address, if you want to specify a different gateway for the VMkernel adapter. Note: Selected netstack should have the default gateway explicitly on each host.

9. Optional: On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses. In ESXi 6.5 and later router advertisement is enabled by default and supports the M and O flags in accordance with RFC 4861.
Static IPv6 addresses	1. Click Add IPv6 address to add a new IPv6 address. 2. Enter the IPv6 address and subnet prefix length, and click OK. 3. To change the VMkernel default gateway, click Override default gateway for this adapter. The VMkernel Default Gateway address for IPv6 is obtained from the selected TCP/IP stack.

10. Review your settings selections on the Ready to complete page and click **Finish**.

Create a VMkernel Adapter on a Host Associated with a vSphere Distributed Switch

Learn how to create a VMkernel adapter on a host that is associated with a distributed switch to provide network connectivity to the host and to handle the traffic for vSphere vMotion, IP storage, Fault Tolerance logging, vSAN, and others. You can set up VMkernel adapters for the standard system traffic on vSphere standard switches and on vSphere distributed switches.

You should dedicate a single distributed port group per VMkernel adapter. For better isolation, you should configure one VMkernel adapter with one traffic type.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **VMkernel adapters**.
3. Click **Add networking**.
4. On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.
5. From the **Select an existing network** option, select a distributed port group and click **Next**.
6. On the Port properties page, configure the settings for the VMkernel adapter.

Option	Description
Network label	The network label is inherited from the label of the distributed port group.
IP settings	Select IPv4, IPv6, or both. Note: The IPv6 option does not appear on hosts that do not have IPv6 enabled.
MTU	Choose whether to get MTU for the network adapter from the switch or to set a custom size. You cannot set the MTU size to a value greater than 9000 bytes.
TCP/IP stack	Select a TCP/IP stack from the list. Once you set a TCP/IP stack for the VMkernel adapter, you cannot change it later. If you select the vMotion or the Provisioning TCP/IP stack, you will be able to use only these stacks to handle vMotion or Provisioning traffic on the host. All VMkernel adapters for vMotion on the default TCP/IP stack are deactivated for future

Option	Description
	vMotion sessions. If you set the Provisioning TCP/IP stack, VMkernel adapters on the default TCP/IP stack are deactivated for operations that include Provisioning traffic, such as virtual machine cold migration, cloning, and snapshot migration.
Available services	You can enable services for the default TCP/IP stack on the host. Select from the available services: • vMotion. Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. The migration with vMotion to the selected host is not possible if the vMotion service is not enabled for any VMkernel adapter on the default TCP/IP stack, or there are no adapters using the vMotion TCP/IP stack. • Provisioning. Handles the data transferred for virtual machine cold migration, cloning, and snapshot migration. • Fault Tolerance logging. Enables Fault Tolerance logging on the host. You can use only one VMkernel adapter for FT traffic per host. • Management. Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software is installed. You can create another VMkernel adapter for management traffic on the host to provide redundancy. • vSphere Replication. Handles the outgoing replication data that is sent from the source ESXi host to the vSphere Replication server. • vSphere Replication NFC. Handles the incoming replication data on the target replication site. • vSAN. Enables the vSAN traffic on the host. Every host that is part of a vSAN cluster must have such a VMkernel adapter. • vSphere Backup NFC. VMkernel port setting for dedicated backup NFC traffic. NFC traffic goes through the VMkernel Adapter when vSphereBackup NFC service is enabled. • NVMe over TCP. VMkernel port setting for dedicated NVMe over TCP storage traffic. NVMe over TCP storage traffic goes through the VMkernel Adapter when NVMe over TCP adapter is enabled. For more information, see <i>vSphere Storage Guide</i>. • NVMe over RDMA. VMkernel port setting for dedicated NVMe over RDMA storage traffic. NVMe over RDMA storage traffic goes through the VMkernel Adapter when NVMe over RDMA adapter is enabled. For more information, see <i>vSphere Storage Guide</i>.

7. Optional: On the IPv4 settings page, select an option for obtaining IP addresses.

Option	Description
Obtain IPv4 settings automatically	Use DHCP to obtain IP settings. A DHCP server must be present on the network.
Use static IPv4 settings	Enter the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack.

Option	Description
	Select the Override default gateway for this adapter check box and enter a gateway address, if you want to specify a different gateway for the VMkernel adapter. Note: Selected netstack should have the default gateway explicitly on each host.

8. Optional: On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses. In ESXi 6.5 and later router advertisement is enabled by default and supports the M and O flags in accordance with RFC 4861.
Static IPv6 addresses	1. Click Add IPv6 address to add a new IPv6 address. 2. Enter the IPv6 address and subnet prefix length, and click OK. 3. To change the VMkernel default gateway, click Override default gateway for this adapter. The VMkernel Default Gateway address for IPv6 is obtained from the selected TCP/IP stack.

9. Review your settings selections on the Ready to complete page and click **Finish**.

Edit a VMkernel Adapter Configuration

Learn how to modify the supported traffic type for a VMkernel adapter, or the way IPv4 or IPv6 addresses are obtained.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **VMkernel adapters**.
3. Select the VMkernel adapter that resides on the target distributed or standard switch and click **Edit**.
4. On the Port properties page, edit the settings for the VMkernel adapter.

Option	Description
MTU	Choose whether to get MTU for the network adapter from the switch or to set a custom size. You cannot set the MTU size to a value greater than 9000 bytes.
TCP/IP stack	Select a TCP/IP stack from the list. Once you set a TCP/IP stack for the VMkernel adapter, you cannot change it later. If you select the vMotion or the Provisioning TCP/IP stack, you will be able to use only these stacks to handle vMotion or Provisioning traffic on the host. All VMkernel adapters for vMotion on the default TCP/IP stack are disabled for future vMotion sessions. If you set the Provisioning TCP/IP stack, VMkernel adapters on the default TCP/IP stack are disabled for operations that include Provisioning traffic, such as virtual machine cold migration, cloning, and snapshot migration.
Available services	You can enable services for the default TCP/IP stack on the host. Select from the available services: • vMotion. Enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion

Option	Description
	traffic is sent. The migration with vMotion to the selected host is not possible if the vMotion service is not enabled for any VMkernel adapter on the default TCP/IP stack, or there are no adapters using the vMotion TCP/IP stack. • Provisioning. Handles the data transferred for virtual machine cold migration, cloning, and snapshot migration. • Fault Tolerance logging. Enables Fault Tolerance logging on the host. You can use only one VMkernel adapter for FT traffic per host. • Management. Enables the management traffic for the host and vCenter Server. Typically, hosts have such a VMkernel adapter created when the ESXi software is installed. You can create another VMkernel adapter for management traffic on the host to provide redundancy. • vSphere Replication. Handles the outgoing replication data that is sent from the source ESXi host to the vSphere Replication server. • vSphere Replication NFC. Handles the incoming replication data on the target replication site. • vSAN. Enables the vSAN traffic on the host. Every host that is part of a vSAN cluster must have such a VMkernel adapter. • vSphere Backup NFC. VMKernel port setting for dedicated backup NFC traffic. NFC traffic goes through the VMKernel Adapter when vSphereBackup NFC service is enabled. • NVMe over TCP. VMKernel port setting for dedicated NVMe over TCP storage traffic. NVMe over TCP storage traffic goes through the VMKernel Adapter when NVMe over TCP adapter is enabled. For more information, see <i>vSphere Storage Guide</i>. • NVMe over RDMA. VMKernel port setting for dedicated NVMe over RDMA storage traffic. NVMe over RDMA storage traffic goes through the VMKernel Adapter when NVMe over RDMA adapter is enabled. For more information, see <i>vSphere Storage Guide</i>.

5. Optional: On the IPv4 settings page, select the method by which IP addresses are obtained.

Option	Description
Obtain IPv4 settings automatically	Use DHCP to obtain IP settings. A DHCP server must be present on the network.
Use static IPv4 settings	Enter the IPv4 IP address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack. Select the Override default gateway for this adapter check box and enter a gateway address, if you want to specify a different gateway for the VMkernel adapter. Note: Selected netstack should have the default gateway explicitly on each host.

6. Optional: On the IPv6 settings page, select an option for obtaining IPv6 addresses.

Note: The IPv6 option does not appear on hosts that do not have IPv6 enabled.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses. In ESXi 6.5 and later router advertisement is enabled by default and supports the M and O flags in accordance with RFC 4861.
Static IPv6 addresses	1. Click Add IPv6 address to add a new IPv6 address. 2. Enter the IPv6 address and subnet prefix length, and click OK. 3. To change the VMkernel default gateway, click Override default gateway for this adapter. The VMkernel Default Gateway address for IPv6 is obtained from the selected TCP/IP stack.

7. Click **OK**.

Override the Default Gateway of a VMkernel Adapter

Learn how to override the default gateway for a VMkernel adapter to provide a different gateway for vSphere vMotion.

Each TCP/IP stack on a host can have only one default gateway. This default gateway is part of the routing table and all services that operate on the TCP/IP stack use it.

For example, the VMkernel adapters vmk0 and vmk1 can be configured on a host.

- vmk0 is used for management traffic on the 10.162.10.0/24 subnet, with default gateway 10.162.10.1
- vmk1 is used for vMotion traffic on the 172.16.1.0/24 subnet

If you set 172.16.1.1 as the default gateway for vmk1, vMotion uses vmk1 as its egress interface with the gateway 172.16.1.1. The 172.16.1.1 gateway is a part of the vmk1 configuration and is not in the routing table. Only the services that specify vmk1 as an egress interface use this gateway. This provides additional Layer 3 connectivity options for services that need multiple gateways.

You can use the vSphere Client or an ESXCLI command to configure the default gateway of a VMkernel adapter.

See [Create a VMkernel Adapter on a vSphere Standard Switch](#), [Create a VMkernel Adapter on a Host Associated with a vSphere Distributed Switch](#), and [Configure the VMkernel Adapter Gateway by Using esxcli Commands](#).

Configure the VMkernel Adapter Gateway by Using esxcli Commands

Learn how to override the default gateway of a VMkernel adapter to provide a different gateway for vSphere vMotion by using esxcli commands.

1. Open an SSH connection to the host.
2. Log in as the root user.
3. Run the following command.

Option	Description
IPv4	<code>esxcli network ip interface ipv4 set -i vmknic -t static -g IPv4 gateway -I IPv4 address -N mask</code>
IPv6	Important: You must turn off DHCPv6 or Router Advertisement before you can set the IPv6 vmknic gateway.

Option	Description
	<pre>esxcli network ip interface ipv6 set -i vmknic -d off -r off</pre> To add a static IPv6 address: <pre>esxcli network ip interface ipv6 address add -i vmknic -I IPv6 address</pre> To set the IPv6 vmknic gateway: <pre>esxcli network ip interface ipv6 set -i vmknic -g IPv6 gateway</pre>

Where *vmknic* is the name of the VMkernel adapter, *gateway* is the IP address of the gateway, *IP address* is the address of the VMkernel adapter, and *mask* is the network mask.

Configure the resolv.conf File by Using esxcli Commands

The `resolv.conf` file is used to configure centrally managed DNS servers. You can configure entries in the `/etc/resolv.conf` file using `esxcli` commands, which allows the modification to remain when rebooting the ESXi host. The property can be set explicitly if DHCP is not enabled.

1. Open an SSH connection to the host.
2. Log in as the root user.
3. Run the following command.

Option	Description
Add DNS server	To add a new DNS server to the end of the list of DNS servers to use for this ESXi host. <pre>esxcli network ip dns server add</pre> Command Options <pre>-N --netstack=<str></pre> The network stack instance. If it is not specified, then use the default netstack instance. <pre>-s --server=<str></pre> The IPV4 or the IPV6 address of the DNS server that you want to add to the DNS server list. Note: This command is required. For example: <code>esxcli network ip dns server add -N vmotion -s xx.xx.xx.xx</code>
Add DNS search	To add a search domain to the list of domains to be searched when trying to resolve a host name on the ESXi host. <pre>esxcli network ip dns search add</pre> Command Options <pre>-d --domain=<str></pre> The string name of a domain that you want to add to the list of search domains. Note: This command is required. <pre>-N --netstack=<str></pre> The network stack instance. If it is not specified, then use the default netstack instance.
Remove DNS server	To remove a DNS server from the list of DNS servers to use for this ESXi host. <pre>esxcli network ip dns server remove IP address</pre> Command Options

Option	Description
	<pre>-a --all</pre> <pre>-N --netstack=<str></pre> The network stack instance. If it is not specified, then use the default netstack instance. <pre>-s --server=<str></pre>
Remove DNS search	To remove a search domain to the list of domains to be searched when trying to resolve an host name on the ESXi host. <pre>esxcli network ip dns search remove</pre> Command Options <pre>-d --domain=<str></pre> The string name of a domain that you want to remove from the list of search domains. Note: This command is required. <pre>-N --netstack=<str></pre> The network stack instance. If it is not specified, then use the default netstack instance.
List DNS search	To print a list the DNS server currently configured on the system in the order in which they will be used. <pre>esxcli network ip dns server list</pre> Command Options <pre>-N --netstack=<str></pre> The network stack instance. If it is not specified, then use the default netstack instance. For example: <code>esxcli network ip dns server list -N vmotion</code>
List DNS search domain	To list the search domains currently configured on the ESXi host in the order in which they will be used while searching. <pre>esxcli network ip dns search list -N</pre> Command Options <pre>--N --netstack=<str></pre> The network stack instance. If it is not specified, then use the default netstack instance. For example: <code>esxcli network ip dns search list -N vmotion</code>

Configure the DNS Hosts File by Using ESXCLI Commands

The DNS hosts file is used for mapping hostnames or domain names to IP addresses. You can configure entries in the `/etc/hosts` file using `esxcli` commands, which allows the modifications to remain unchanged while rebooting the ESXi host.

1. Open an SSH connection to the host.
2. Log in as the root user.
3. Run the following command.

Option	Description
add	To add the hostname and IP address mapping, run the following command. <pre>esxcli network ip hosts add --ip --hostname</pre>

Option	Description
	For example: <code>esxcli network ip hosts add --hostname www.samplehostname.com --ip xx.xx.xx.xx</code>
remove	To remove a hostname mapping from the configuration, run the following command. <code>esxcli network ip hosts remove --hostname</code> For example: <code>esxcli network ip hosts remove --hostname www.samplehostname.com --ip xx.xx.xx.xx</code>
list	To list the IP addresses and their associated DNS entries, run the following command. <code>esxcli network ip hosts list</code>

Where *IP address* is the address of the VMkernel adapter, *hostname* is the DNS entry you want to associate with the IP address, *aliases* is any aliases you want to associate, and *comment* is a line about this entry.

Note: ESXCLI network IP hosts add values only to the other hosts. The ESXCLI network IP host does not add values to the current host. The information about the current host is managed through the host name under the **Network Settings**. If you try to change the current hostname you might see an error similar to the following:

```
Error adding etc host item: User can not set item for management IP: X.X.X.X
```

For information on changing the host name of the current host, see [Change the Configuration of a TCP/IP Stack on a Host](#).

View TCP/IP Stack Configuration on a Host

You can view the DNS and routing configuration of a TCP/IP stack on a host. You can also view the IPv4 and IPv6 routing tables, the congestion control algorithm, and the maximum number of allowed connections.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **TCP/IP configuration**.
3. Select a stack from the TCP/IP Stacks table.
If no custom TCP/IP stacks are configured on the host, you view the default, vMotion, and Provisioning TCP/IP stacks on the host.

DNS and routing details about the selected TCP/IP stack appear below the TCP/IP Stacks table. You can view the IPv4 and IPv6 routing tables, and the DNS and routing configuration for the stack.

Note: The IPv6 routing table is only visible if IPv6 is enabled on the host.

The **Advanced** tab contains information about the configured congestion control algorithm and the maximum number of allowed connections to the stack.

Change the Configuration of a TCP/IP Stack on a Host

Learn how to modify the DNS and default gateway configuration of a TCP/IP stack on a host. You can also change the congestion control algorithm, the maximum number of connections, and the name of custom TCP/IP stacks.

Note: You can change the DNS and default gateway configuration of the default TCP/IP stack only. Multiple DNS and gateways configurations are supported when using separate TCP/IP stacks.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **TCP/IP configuration**.
3. Select a stack from the table, click **Edit** and make the appropriate changes.

Page	Option
DNS Configuration	Select a method of obtaining the DNS server. • Select Obtain settings automatically from a VMkernel network adapter and select a network adapter from the VMKernel network adapter drop-down menu • Select Enter settings manually and edit the DNS configuration settings. a. Edit the Host name. b. Edit the Domain name. c. Type a preferred DNS server IP address. d. Type an alternate DNS server IP address. e. Use the Search domains text box to specify DNS suffixes to use in DNS search when resolving unqualified domain names.
Routing	Edit the VMkernel gateway information. Note: Removing the default gateway might cause the client to lose connectivity with the host.
Name	Change the name of a custom TCP/IP stack
Advanced	Edit the maximum number of connections and the congestion control algorithm of the stack

4. Click **OK** to apply your changes.

You can add static routes to additional gateways by using CLI commands. For more information, see <http://kb.vmware.com/kb/2001426>

Explicit Congestion Notification

Explicit Congestion Notification (ECN) allows a TCP sender to reduce the transmission rate to avoid packet drops. ECN is specified in RFC 3168. vSphere 7.0 and later supports ECN and it is enabled by default.

You can get the ECN status of any netstack using esxcli command.

1. In the ESXi Shell on the host, use the following command.

```
esxcli network ip netstack set -N <NetStack-Name> --ecn=<str>
```

2. You can set the status of ECN. The setting has the following possible values in ESXi.

```
...
--ecn=<str>          The status of ECN (Explicit Congestion Notification).
                      disabled: Completely disable ECN functionality
                      echo-only: Only echoing ECN, won't initiate
```

enabled: Fully enable ECN functionality

The default value of this parameter is enabled. ESXi is likely to use ECN if the value of this parameter is enabled. If a router or a network appliance in the environment cannot handle IP packets with ECN bits accurately, then the ECN capability can be set to disabled.

Create a Custom TCP/IP Stack

Learn how to create a custom TCP/IP stack on a host to forward networking traffic through a custom application.

Once you set a TCP/IP stack for the VMkernel adapter, you cannot change it later. If you select the vMotion or the Provisioning TCP/IP stack, you can use only these stacks to handle vMotion or Provisioning traffic on the host. All VMkernel adapters for vMotion on the default TCP/IP stack are disabled for future vMotion sessions. If you set the Provisioning TCP/IP stack, VMkernel adapters on the default TCP/IP stack are disabled for operations that include Provisioning traffic, such as virtual machine cold migration, cloning, and snapshot migration.

If you must change the TCP/IP stack configuration, delete the existing VMkernel adapter and create a new one. You can then create a TCP/IP stack for that adapter.

1. Open an SSH connection to the host.
2. Log in as the root user.
3. Run the ESXCLI command.

```
esxcli network ip netstack add -N="stack_name"
```

The custom TCP/IP stack is created on the host. You can assign VMkernel adapters to the stack.


Remove a VMkernel Adapter

Learn how to remove a VMkernel adapter from a vSphere distributed or a standard switch when you no longer need the adapter. Make sure that you leave at least one VMkernel adapter for management traffic on the host to keep the network connectivity up.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **VMkernel adapters**.
3. Select a VMkernel adapter from the list, and click the **Remove selected network adapter** icon.
4. Click **Remove**.

LACP Support on a vSphere Distributed Switch

With link aggregation control protocol (LACP) support on a vSphere Distributed Switch, you can connect ESXi hosts to physical switches by using dynamic link aggregation. You can create multiple link aggregation groups (LAGs) on a distributed switch to aggregate the bandwidth of physical NICs on ESXi hosts that are connected to LACP port channels.

Figure 82: Enhanced LACP Support on a vSphere Distributed Switch

LACP Configuration on the Distributed Switch

You configure a LAG with two or more ports and connect physical NICs to the ports. LAG ports are teamed within the LAG, and the network traffic is load balanced between the ports through an LACP hashing algorithm. You can use a LAG to handle the traffic of distributed port groups to provide increased network bandwidth, redundancy, and load balancing to the port groups.

When you create a LAG on a distributed switch, a LAG object is also created on the proxy switch of every host that is connected to the distributed switch. For example, if you create LAG1 with two ports, LAG1 with the same number of ports is created on every host that is connected to the distributed switch.

On a host proxy switch, you can connect one physical NIC to only one LAG port. On the distributed switch, one LAG port can have multiple physical NICs from different hosts connected to it. The physical NICs on a host that you connect to the LAG ports must be connected to links that participate in an LACP port channel on the physical switch.

You can create up to 64 LAGs on a distributed switch. A host can support up to 32 LAGs. However, the number of LAGs that you can actually use depends on the capabilities of the underlying physical environment and the topology of the virtual network. For example, if the physical switch supports up to four ports in an LACP port channel, you can connect up to four physical NICs per host to a LAG.

You can configure LACP timeout value. LACP functions by sending frames down all the links that have the protocol enabled. If it finds a device on the other end of the link that also has LACP enabled, it sends frames independently along the same links, enabling the two units to detect multiple links between themselves and then combine them into a single logical link. The timeout value is the amount of time that a LAG interface waits for a Protocol Data Unit (PDU) from the remote system before bringing the LACP session down. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate.

Port Channel Configuration on the Physical Switch

For each host on which you want to use LACP, you must create a separate LACP port channel on the physical switch. You must consider the following requirements when configuring LACP on the physical switch:

- The number of ports in the LACP port channel must be equal to the number of physical NICs that you want to group on the host. For example, if you want to aggregate the bandwidth of two physical NICs on a host, you must create an LACP port channel with two ports on the physical switch. The LAG on the distributed switch must be configured with at least two ports.
- The hashing algorithm of the LACP port channel on the physical switch must be the same as the hashing algorithm that is configured to the LAG on the distributed switch.
- All physical NICs that you want to connect to the LACP port channel must be configured with the same speed and duplex.

Configure LACP Teaming and Failover for Distributed Port Groups

To handle the network traffic of distributed port groups by using a LAG, you assign physical NICs to the LAG ports and set the LAG as active in the teaming and failover order of distributed port groups.

Table 185: LACP Teaming and failover configuration of distributed port groups

Failover Order	Uplinks	Description
Active	A single LAG	You can only use one active LAG or multiple standalone uplinks to handle the traffic of distributed port groups. You cannot configure multiple active LAGs or mix active LAGs and standalone uplinks.
Standby	Empty	Having an active LAG and standby uplinks and the reverse is not supported. Having a LAG and another standby LAG is not supported.
Unused	All standalone uplinks and other LAGs if any	Because only one LAG must be active and the Standby list must be empty, you must set all standalone uplinks and other LAGs to unused.

Configure a Link Aggregation Group to Handle the Traffic for Distributed Port Groups

To aggregate the bandwidth of multiple physical NICs on hosts, you can create a link aggregation group (LAG) on the distributed switch and use it to handle the traffic of distributed port groups.

- Verify that for every host where you want to use LACP, a separate LACP port channel exists on the physical switch. See [LACP Support on a vSphere Distributed Switch](#).
- Verify that the vSphere Distributed Switch where you configure the LAG is version 6.5 or later.
- Verify that enhanced LACP is supported on the distributed switch.

Newly created LAGs do not have physical NICs assigned to their ports and are unused in the teaming and failover order of distributed port groups. To handle the network traffic of distributed port groups by using a LAG, you must migrate the traffic from standalone uplinks to the LAG.

Create a Link Aggregation Group

To migrate the network traffic of distributed port groups to a link aggregation group (LAG), you create a LAG on the distributed switch.

1. In the vSphere Client, navigate to the distributed switch.
2. On the **Configure** tab, expand **Settings** and select **LACP**.
3. Click the **New Link Aggregation Group** icon.
4. Name the new LAG.
5. Set the number of ports to the LAG.

Set the same number of ports to the LAG as the number of ports in the LACP port channel on the physical switch. A LAG port has the same function as an uplink on the distributed switch. All LAG ports form a NIC team in the context of the LAG.

6. Select the LACP negotiating mode of the LAG.

Option	Description
Active	All LAG ports are in an Active negotiating mode. The LAG ports initiate negotiations with the LACP port channel on the physical switch by sending LACP packets.
Passive	The LAG ports are in Passive negotiating mode. The LAG ports respond to LACP packets they receive but do not initiate LACP negotiation.

If the LACP-enabled ports on the physical switch are in Active negotiating mode, you can set the LAG ports in Passive mode and the reverse.

7. Select a load balancing mode from the hashing algorithms that LACP defines.

Note: The hashing algorithm must be the same as the hashing algorithm set to the LACP port channel on the physical switch.

8. Select the link aggregation timeout mode.

Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the LACP timeout preference you select. In a fast timeout, PDUs are sent every second, and, in a slow timeout, PDUs are sent every 30 seconds. A slow timeout is the default setting.

Note: A fast timeout is supported only on distributed switch version 7.0.2 or later.

9. Set the VLAN and the NetFlow policies for the LAG.

This option is active when overriding the VLAN and NetFlow policies per individual uplink ports is enabled on the uplink port group. If you set the VLAN and NetFlow policies to the LAG, they override the policies set on the uplink port group level.

10. Click **OK**.

The new LAG is unused in the teaming and failover order of distributed port groups. No physical NICs are assigned to the LAG ports.

As with standalone uplinks, the LAG has a representation on every host that is associated with the distributed switch. For example, if you create LAG1 with two ports on the distributed switch, a LAG1 with two ports is created on every host that is associated with the distributed switch.

Set the LAG as standby in the teaming and failover configuration of distributed port groups. In this way, you create an intermediate configuration that lets you migrate the network traffic to the LAG without losing network connectivity.

Set a LAG as Standby in the Teaming and Failover Order

The new link aggregation group (LAG) by default is unused in the teaming and failover order of distributed port groups. Because only one LAG or only standalone uplinks can be active for distributed port groups, you must create an intermediate teaming and failover configuration, where the LAG is standby. This configuration lets you migrate physical NICs to the LAG ports by keeping the network connectivity up.

1. Navigate to the distributed switch.
2. From the **Actions** menu, select **Distributed Port Group > Manage Distributed Port Groups**.
3. Select **Teaming and failover** and click **Next**.
4. Select the port groups where you want to use the LAG.
5. In Failover order, select the LAG and use the up arrow to move it to the Standby uplinks list.
6. Click **Next**, review the message that informs you about the usage of the intermediate teaming and failover configuration, and click **OK**.
7. On the Ready to complete page, click **Finish**.

Migrate physical NICs from standalone uplinks to the LAG ports.

Assign Physical NICs to the Ports of the Link Aggregation Group

Learn how to set the new link aggregation group (LAG) as standby in the teaming and failover order of distributed port groups. Having the LAG as standby lets you safely migrate the physical NICs from standalone uplinks to the LAG ports without losing network connectivity.

- Verify that either all LAG ports or the corresponding LACP-enabled ports on the physical switch are in active LACP negotiating mode.

- Verify that the physical NICs that you want to assign to the LAG ports have the same speed and are configured at full duplex.
1. In the vSphere Client, navigate to the distributed switch where the LAG resides.
 2. From the **Actions** menu, select **Add and Manage Hosts**.
 3. Select **Manage host networking**.
 4. Select the host whose physical NICs you want to assign to the LAG ports and click **Next**.
 5. On the Select network adapter tasks page, select **Manage physical adapters** and click **Next**.
 6. On the Manage physical adapters page, select a NIC and click **Assign an uplink**.
 7. Select a LAG port and click **OK**.
 8. Repeat [Step 6](#) and [Step 7](#) for all physical NICs that you want to assign to the LAG ports.
 9. Complete the wizard.

Configure Two Physical NICs to a LAG in the Add and Manage Hosts Wizard

For example, if you have a LAG with two ports, you configure a physical NIC to each LAG port in the **Add and Manage Hosts** wizard.

Set the LAG as active and all standalone uplinks to unused in the teaming and failover order of distributed port groups.

Set the Link Aggregation Group as Active in the Teaming and Failover Order of the Distributed Port Group


Learn how to set the LAG as active and move all standalone uplinks as unused in the teaming and failover order of the distributed port groups, after you have migrated physical NICs to the ports of the link aggregation group (LAG).

1. Navigate to the distributed switch.
2. From the **Actions** menu, select **Distributed Port Group > Manage Distributed Port Groups**.
3. Select **Teaming and failover** and click **Next**.
4. Select the port groups where you set the LAG as standby and click **Next**.
5. In Failover order, use the up and down arrows to move the LAG in the Active list, all standalone uplinks in the Unused list, and leave the Standby list empty.
6. Click **Next** and click **Finish**.

You safely migrated network traffic from standalone uplinks to a LAG for distributed port groups and created a valid LACP teaming and failover configuration for the groups.

Topology of a Distributed Switch that Uses a LAG

If you configure a LAG with two ports to handle the traffic of a distributed port group, you can check the topology of the distributed switch to view how it changed as a result of the new configuration.

Figure 83: Distributed Switch Topology with a LAG

Edit a Link Aggregation Group

Learn how to modify the settings of a link aggregation group (LAG) if you need to add more ports to the group or change the LACP negotiating mode, the load balancing algorithm, or the VLAN and NetFlow policies.

1. In the vSphere Client, navigate to the vSphere Distributed Switch.
2. On the **Configure** tab, expand **Settings** and select **LACP**.
3. Click the **New Link Aggregation Group** icon.
4. In the **Name** text box, type a new name for the LAG.
5. Change the number of ports for the LAG if you want to add more physical NICs to it.
The new NICs must be connected to ports that are part of an LACP port channel on the physical switch.
6. Change the LACP negotiating mode of the LAG.
If all ports on the physical LACP port channel are in Active LACP mode, you can change the LACP mode of the LAG to Passive and the reverse.
7. Change the load balancing mode of the LAG.
You can select from the load balancing algorithms that LACP defines.
8. Select the link aggregation timeout mode.
Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the LACP timeout preference you select. In a fast timeout, PDUs are sent every second, and, in a slow timeout, PDUs are sent every 30 seconds. A slow timeout is the default setting.

Note: You can set a fast LACP timeout by running the following ESXCLI command. However, the setting is not available against an NSX enabled DVS or N-VDS (opaque NSX Distributed vSwitch).

```
esxcli network vswitch dvs vmware lacp timeout set --vds DVS-name --lag-id <integer> --timeout 1
```

9. Change the VLAN and the NetFlow policies.

This option is active when the option for overriding the VLAN and NetFlow policies for individual ports is enabled on the uplink port group. If you change the VLAN and NetFlow policies for the LAG, they override the policies set at the uplink port group level.

10. Click **OK**.

Limitations of the LACP Support on a vSphere Distributed Switch

The LACP support on a vSphere Distributed Switch lets network devices negotiate automatic bundling of links by sending LACP packets to a peer. However, the LACP support on a vSphere Distributed Switch has limitations.

- The LACP is not supported with software iSCSI port binding. iSCSI multipathing over LAG is supported, if port binding is not used.
- The LACP support settings are not available in host profiles.
- The LACP support is not possible between nested ESXi hosts.
- The LACP support does not work with the ESXi dump collector.
- The LACP control packets (LACP PDU) do not get mirrored when port mirroring is enabled.
- The teaming and failover health check does not work for LAG ports. LACP checks the connectivity of the LAG ports.
- The enhanced LACP support works correctly when only one LAG handles the traffic per distributed port or port group.

Back Up and Restore Network Configurations

vSphere enables you to backup and restore the configuration of a vSphere Distributed Switch, distributed and uplink port groups in cases of invalid changes or a transfer to another deployment.

Back Up and Restore a vSphere Distributed Switch Configuration

vCenter Server provides the ability to backup and restore the configuration of a vSphere Distributed Switch. Learn how to restore the virtual network configuration in cases of database or upgrade failure. You can also use a saved switch configuration as a template to create a copy of the switch in the same or a new vSphere environment.

You can import or export a configuration of a distributed switch including its port groups. For information about exporting, importing, and restoring a port group configuration, see [Export, Import, and Restore vSphere Distributed Port Group Configurations](#).

Note: You can use a saved configuration file to restore policies and hosts associations on the distributed switch. You cannot restore the connection of physical NICs to uplink ports or ports of link aggregation groups.

Export vSphere Distributed Switch Configurations

Learn how to export vSphere Distributed Switch and distributed port group configurations to a file. The file preserves valid network configurations, enabling transfer of these configurations to other environments.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Settings > Export Configuration**.
3. Choose to export the distributed switch configuration, or export the distributed switch configuration and all port groups.
4. Optional: Enter notes about this configuration in the **Descriptions** field.
5. Click **OK**.
6. Click **Yes** to save the configuration file to your local system.

Use the exported configuration file to do the following tasks:

- Create a copy of the exported distributed switch in a vSphere environment. See [Import a vSphere Distributed Switch Configuration](#).
- Overwrite the settings on an existing distributed switch. See [Restore a vSphere Distributed Switch Configuration](#).

You can also export, import, and restore only port group configurations. See [Export, Import, and Restore vSphere Distributed Port Group Configurations](#).

Import a vSphere Distributed Switch Configuration

Learn how to import a stored configuration file to create a new vSphere Distributed Switch or to restore a switch that has been deleted earlier.

The configuration file contains the networking settings of the switch. By using it you can also replicate the switch in other virtual environments.

Note: You can use a saved configuration file to replicate the switch instance, its host associations, and policies. You cannot replicate the connection of physical NICs to uplink ports or ports on link aggregation groups.

1. In the vSphere Client, navigate to a data center.
2. Right-click the data center and select **Distributed Switch > Import Distributed Switch**.
3. Browse to the location of the configuration file.
4. To assign the keys from the configuration file to the switch and its port groups, select the **Preserve original distributed switch and port group identifiers** check box and click **Next**.

You can use the **Preserve original distributed switch and port group identifiers** option in the following cases:

- Recreate a deleted switch.
- Restore a switch whose upgrade has failed.

All port groups are recreated and the hosts that have been connected to the switch are added again.

5. Review the settings for the switch and click **Finish**.

A new distributed switch is created with settings from the configuration file. If you have included distributed port group information in the configuration file, the port groups are also created.

Restore a vSphere Distributed Switch Configuration

Learn how to use the restore option to reset the configuration of an existing distributed switch to the settings in the configuration file. Restoring a distributed switch changes the settings on the selected switch back to the settings saved in the configuration file.

Note: You can use a saved configuration file to restore policies and hosts associations on the distributed switch. You cannot restore the connection of physical NICs to uplink ports or ports of link aggregation groups.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch in the navigator and select **Settings > Restore Configuration**.
3. Browse for the configuration backup file to use.
4. Select **Restore distributed switch and all port groups** or **Restore distributed switch only** and click **Next**.
5. Review the summary information for the restore.

Restoring a distributed switch will overwrite the current settings of the distributed switch and its port groups. It will not delete existing port groups that are not part of the configuration file.

6. Click **Finish**.

The distributed switch configuration has been restored to the settings in the configuration file.

Export, Import, and Restore vSphere Distributed Port Group Configurations

Learn to export vSphere distributed port group configurations to a file. The configuration file allows you to preserve valid port group configurations, enabling distribution of these configurations to other deployments.

You can export port group information at the same time you export distributed switch configurations. See [Back Up and Restore a vSphere Distributed Switch Configuration](#).

Export vSphere Distributed Port Group Configurations

Learn how to export a distributed port group configurations to a file. The configuration preserves valid network configurations, enabling distribution of these configurations to other deployments.

1. Locate a distributed port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups**.
2. Right-click the distributed port group and select **Export Configuration**.
3. Optional: In the **Descriptions** field, type notes about this configuration.
4. Click **OK**.

Click **Yes** to save the configuration file to your local system.

You now have a configuration file that contains all the settings for the selected distributed port group. You can use this file to create multiple copies of this configuration on an existing deployment, or overwrite settings of existing distributed port groups to conform to the selected settings.

You can use the exported configuration file to do the following tasks:

- To create a copy of the exported distributed port group, see [Import a vSphere Distributed Port Group Configuration](#).
- To overwrite settings on an existing distributed port group, see [Restore a vSphere Distributed Port Group Configuration](#).

Import a vSphere Distributed Port Group Configuration

Learn how to import to create a distributed port group from a configuration file.

If an existing port group has the same name as the imported port group, the new port group name has a number appended in parentheses. The settings from the imported configuration are applied to the new port group and the settings of the original port group remain unchanged.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Distributed Port Group > Import Distributed Port Group**.
3. Browse to the location of your saved configuration file and click **Next**.
4. Review the import settings before completing the import.
5. Click **Finish**.

Restore a vSphere Distributed Port Group Configuration

Learn how to use the restore option to reset the configuration of an existing distributed port group to the settings in a configuration file.

1. Locate a distributed port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups**.
2. Right-click the distributed port group and select **Restore Configuration**.
3. Select one of the following and click **Next**:
 - **Restore to previous configuration** to roll your port group configuration back one step. You cannot restore the port group configuration completely if you have performed more than one step.
 - **Restore configuration from a file** lets you restore the port group configuration from an exported backup file. You can also use a distributed switch backup file as long as it contains configuration information for the port group.
4. Review the summary information for the restore.

The restore operation overwrites the current settings of the distributed port group with the settings from the backup. If you are restoring the port group configuration from a switch backup file, the restore operation does not delete existing port groups that are not a part of the file.
5. Click **Finish**.

ESXi Configuration Manager Integration

You can manage core vSphere networking configurations in the vSphere Lifecycle Manager cluster level configuration manager. The vSphere networking configurations include *vmknics*, *netstacks*, *vdsvswitches*, *pnics*, *netdump.firewall*, *ipsec*, */etc/hosts*, */etc/resolv.conf*.

Starting vSphere 8.0 Update 3, vSphere Distributed Switch is integrated with vSphere Configuration Profiles (VCP). You can manage host distributed switch configuration in the vSphere Configuration Profiles (VCP) cluster.

Note: Global distributed switch configuration such as MTU is out of the scope of VCP.

You can manage vSphere Distributed Switch in addition to all the host configurations at the cluster level by using a single configuration file in JSON format. You can check if all the hosts in the cluster are compliant with the distributed and the host configurations. You can migrate the ESXi host networking from standard switch to a distributed switch using a remediation work flow at the cluster level. VMkernel network adapter on distributed switch is now fully supported by the configuration manager.

Using the configuration manager, you can configure VCPs that are attached to the VCP cluster which includes the host settings such as, the list of distributed switches that the hosts can join, and the PNIC to VDS uplink mapping.

Export Host Configuration

Learn how desired state configuration helps you to seamlessly manage your vSphere Networking configurations across geographies and different domains.

- You must have a vSphere Life Cycle Manager enabled cluster.
- Ensure that you have hosts in the cluster.

In this task, we are exporting a JSON file holding the desired state specification for a vmknics adapter.

1. In the vSphere Client, navigate to the Cluster.
2. Click **Desired State > Host Settings > Extract from reference host**. The **Extract Settings** dialog box appears.
3. Choose the host from the existing cluster.
4. Click **Next**
5. Click **Download Settings** to download the file.

The JSON file is downloaded.

You can then reuse the JSON file to apply the desired state that it contains to another cluster in the same or in a different vCenter Server instance.

Note: Solution configuration is not included. Cached information is not included. Only host specific property is included.

Note: If you have vSphere Life Cycle Manager enabled cluster, you can not join the host in the cluster to a vSphere Distributed Switch.

Import Host Configuration

Learn how to import a desired state as a JSON file.

- You must have a vSphere Life Cycle Manager enabled cluster.
- Ensure you have hosts in the cluster.
- Ensure you have a valid host configuration JSON file.

Learn how to import a JSON file with host configuration. Before applying a desired state on a cluster, you can scan all hosts in the cluster against the desired state and check the cluster compliance against the desired state. You can run compliance checks to detect any drifts instantly.

1. In the vSphere Client, navigate to the Cluster.
2. Click **Desired State > Host Settings > Import Settings**. The **Import Host Settings** dialog box appears.
3. Click **Browse** to select file. The compliance is validated.
If there are hosts that are not compliant, an error message is displayed.
4. Click **Remediate** to remediate the host compliance issue. The **Remediate Settings** dialog box appears. The **Pre-check** is in-progress.
5. The **Pre-Check** is completed. The host level remediation details are displayed under **Review Impact**.
6. Click **Remediate** to remediate the configuration settings.

You can view the imported settings on the host.

Rollback and Recovery of the Management Network

Learn how to prevent and recover from misconfiguration of the management network by using the rollback and recovery support of the vSphere Distributed Switch and vSphere Standard Switch.

Rollback is available for use on both standard and distributed switches. To fix invalid configuration of the management network, you can connect directly to a host to fix the issues through the DCUI.

vSphere Networking Rollback

Learn how vSphere protects hosts from losing connection to vCenter Server as a result of misconfiguration of the management network by rolling back configuration changes.

In vSphere networking rollback is enabled by default. However, you can activate or deactivate rollbacks at the vCenter Server level.

Host Networking Rollbacks

Host networking rollbacks occur when an invalid change is made to the networking configuration for the connection with vCenter Server. Every network change that disconnects a host also triggers a rollback. The following examples of changes to the host networking configuration might trigger a rollback:

- Updating the speed or duplex of a physical NIC.
- Updating DNS and routing settings.
- Updating teaming and failover policies or traffic shaping policies of a standard port group that contains the management VMkernel network adapter.
- Updating the VLAN of a standard port group that contains the management VMkernel network adapter.
- Increasing the MTU of management VMkernel network adapter and its switch to values not supported by the physical infrastructure.
- Changing the IP settings of management VMkernel network adapters.
- Removing the management VMkernel network adapter from a standard or distributed switch.
- Removing a physical NIC of a standard or distributed switch containing the management VMkernel network adapter.
- Migrating the management VMkernel adapter from vSphere standard to distributed switch.

If a network disconnects for any of these reasons, the task fails and the host reverts to the last valid configuration.

vSphere Distributed Switch Rollbacks

Distributed switch rollbacks occur when invalid updates are made to distributed switches, distributed port groups, or distributed ports. The following changes to the distributed switch configuration trigger a rollback:

- Changing the MTU of a distributed switch.
- Changing the following settings in the distributed port group of the management VMkernel network adapter:
 - Teaming and failover
 - VLAN
 - Traffic shaping
- Blocking all ports in the distributed port group containing the management VMkernel network adapter.
- Overriding the policies on at the level of the distributed port for the management VMkernel network adapter.

If a configuration becomes invalid because of any of the changes, one or more hosts might become out of synchronization with the distributed switch.

If you know where the conflicting configuration setting is located, you can manually correct the setting. For example, if you have migrated a management VMkernel network adapter to a new VLAN, the VLAN might not be actually trunked on the physical switch. When you correct the physical switch configuration, the next distributed switch-to-host synchronization will resolve the configuration problem.

If you are not sure where the problem exists, you can restore the state of the distributed switch or distributed port group to an earlier configuration. See [Restore a vSphere Distributed Port Group Configuration](#).

Disable Network Rollback

Rollback is enabled by default in vSphere. You can disable rollback in vCenter Server by using the vSphere Client.

1. In the vSphere Client, navigate to a vCenter Server instance.
2. On the **Configure** tab, expand **Settings** and select **Advanced Settings**.
3. Click **Edit Settings**.
4. Select the `config.vpxd.network.rollback` key, and change the value to false.
If the key is not present, you can add it and set the value to false.
5. Click **OK**.
6. Restart vCenter Server to apply the changes.

Disable Network Rollback by Using the vCenter Server Configuration File

Rollback is enabled by default in vSphere. You can disable rollback by editing the `vpxd.cfg` configuration file of vCenter Server directly.

1. On the host machine of , navigate to the directory `/etc/vmware-vpx`.
2. Open the `vpxd.cfg` file for editing.
3. In the `<network>` element, set the `<rollback>` element to `false`:

```
<config>
  <vpxd>
    <network>
      <rollback>false</rollback>
    </network>
  </vpxd>
</config>
```

4. Save and close the file.
5. Restart the vCenter Server system.

Resolve Errors in the Management Network Configuration on a vSphere Distributed Switch

Learn how to use the Direct Console User Interface (DCUI) to restore the connection between vCenter Server and a host that accesses the management network through a distributed switch.

Verify that the management network is configured on a port group on the distributed switch.

If networking rollback is disabled, misconfiguring the port group for the management network on the distributed switch leads to loss of connection between vCenter Server and the hosts that are added to the switch. You have to use the DCUI to connect each host individually.

If the uplinks that you use to restore the management network are also used by VMkernel adapters that handle other types of traffic (vMotion, Fault Tolerance, and so on), the adapters lose network connectivity after the restore.

For more information about accessing and using the DCUI, see the *vSphere Security* documentation.

Note: Recovery of the management connection on a distributed switch is not supported on stateless ESXi instances.

1. Connect to the DCUI of the host.
2. From the **Network Restore Options** menu, select **Restore vDS**.
3. Configure the uplinks and optionally the VLAN for the management network.
4. Apply the configuration.

The DCUI creates a local ephemeral port and applies the values you provided for the VLAN and uplinks. The DCUI moves the VMkernel adapter for the management network to the new local port to restore connectivity to vCenter Server.

After the connection of the host to vCenter Server is restored, correct the configuration of the distributed port group and re-add the VMkernel adapter to the group.

vSphere Networking Policies

Policies set at the standard switch or distributed port group level apply to all of the port groups on the standard switch or to ports in the distributed port group. The exceptions are the configuration options that are overridden at the standard port group or distributed port level.

Apply Networking Policies on a vSphere Standard or Distributed Switch

You apply networking policies differently on vSphere Standard Switches and vSphere Distributed Switches. Not all policies that are available for a vSphere Distributed Switch are also available for a vSphere Standard Switch.

Table 186: Virtual Switch Objects Where Policies Apply

Virtual Switch	Virtual Switch Object	Description
vSphere Standard Switch	Entire switch	When you apply policies on the entire standard switch, the policies are propagated to all standard port groups on the switch.
	Standard port group	You can apply different policies on individual port groups by overriding the policies that are inherited from the switch.
vSphere Distributed Switch	Distributed port group	When you apply policies on a distributed port group, the policies are propagated to all ports in the group.
	Distributed port	You can apply different policies on individual distributed ports by overriding the policies that are inherited from the distributed port group.
	Uplink port group	You can apply policies at uplink port group level, and the are policies are propagated to all ports in the group.
	Uplink port	You can apply different policies on individual uplink ports by overriding the policies that are inherited from the uplink port group.

Table 187: Policies Available for a vSphere Standard Switch and vSphere Distributed Switch

Policy	Standard Switch	Distributed Switch	Description
Teaming and failover	Yes	Yes	Lets you configure the physical NICs that handle the network traffic for a standard switch, standard port group, distributed port group, or distributed port. You arrange the physical NICs in a failover order and apply different load balancing policies over them.
Security	Yes	Yes	Provides protection of traffic against MAC address impersonation and unwanted port scanning. The networking security policy is implemented in Layer 2 of the networking protocol stack.
Traffic shaping	Yes	Yes	Lets you restrict the network bandwidth that is available to ports, but also to allow bursts of traffic to flow through at higher speeds. ESXi shapes outbound network traffic on standard switches and inbound and outbound traffic on distributed switches.
VLAN	Yes	Yes	Lets you configure the VLAN tagging for a standard or distributed switch. You can configure External Switch Tagging(EST), Virtual Switch Tagging (VST), and Virtual Guest Tagging (VGT).
Monitoring	No	Yes	Enables and disables NetFlow monitoring on a distributed port or port group.
Traffic filtering and marking	No	Yes	Lets you protect the virtual network from unwanted traffic and security attacks or apply a QoS tag to a certain traffic type.
Resources allocation	No	Yes	Lets you associate a distributed port or port group with a user-defined network resource pool. In this way, you can better control the bandwidth that is available to the port or port group. You can use the resource allocation policy with vSphere Network I/O Control version 2 and 3.
Port blocking	No	Yes	Lets you selectively block ports from sending and receiving data.

Configure Overriding Networking Policies on Port Level

Learn how to apply different policies for distributed ports, you configure the per-port overriding of the policies that are set at the port group level. You can also enable the reset of any configuration that is set on per-port level when a distributed port disconnects from a virtual machine.

1. Locate a distributed port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups**.
2. Right-click the distributed port group and select **Edit Settings**.
3. Select the **Advanced** page.

Option	Description
Configure reset at disconnect	From the drop-down menu, enable or disable reset at disconnect. When a distributed port is disconnected from a virtual machine, the configuration of the distributed port is reset to the distributed port group setting. Any per-port overrides are discarded.
Override port policies	Select the distributed port group policies to be overridden on a per-port level.

4. Optional: Use the policy pages to set overrides for each port policy.
5. Click **OK**.

What is Teaming and Failover Policy

NIC teaming lets you increase the network capacity of a virtual switch by including two or more physical NICs in a team. To determine how the traffic is rerouted in case of adapter failure, you include physical NICs in a failover order. To determine how the virtual switch distributes the network traffic between the physical NICs in a team, you select load balancing algorithms depending on the needs and capabilities of your environment.

NIC Teaming Policy

You can use NIC teaming to connect a virtual switch to multiple physical NICs on a host to increase the network bandwidth of the switch and to provide redundancy. A NIC team can distribute the traffic between its members and provide passive failover in case of adapter failure or network outage. You set NIC teaming policies at virtual switch or port group level for a vSphere Standard Switch and at a port group or port level for a vSphere Distributed Switch.

Note: All ports on the physical switch in the same team must be in the same Layer 2 broadcast domain.

Load Balancing Policy

The Load Balancing policy determines how network traffic is distributed between the network adapters in a NIC team. vSphere virtual switches load balance only the outgoing traffic. Incoming traffic is controlled by the load balancing policy on the physical switch.

For more information about each load balancing algorithm, see [Load Balancing Algorithms Available for Virtual Switches](#).

Network Failure Detection Policy

You can specify one of the following methods that a virtual switch uses for failover detection.

Link status only	Relies only on the link status that the network adapter provides. Detects failures, such as removed cables and physical switch power failures. However, link status does not detect the following configuration errors: - Physical switch port that is blocked by spanning tree or is misconfigured to the wrong VLAN . - Pulled cable that connects a physical switch to another networking devices, for example, an upstream switch .
Beacon probing	Sends out and listens for Ethernet broadcast frames, or beacon probes, that physical NICs send to detect link failure in all physical NICs in a team. ESXi hosts send beacon packets every second. Beacon probing is most useful to detect failures in the closest physical switch to the ESXi host, where the failure does not cause a link-down event for the host. Use beacon probing with three or more NICs in a team because ESXi can detect failures of a single adapter. If only two NICs are assigned and one of them loses connectivity, the switch cannot determine which NIC needs to be taken out of service because both do not receive beacons and as a result all packets sent to both uplinks. Using at least three NICs in such a team allows for $n-2$ failures where n is the number of NICs in the team before reaching an ambiguous situation.

Failback Policy

By default, a failback policy is enabled on a NIC team. If a failed physical NIC returns online, the virtual switch sets the NIC back to active by replacing the standby NIC that took over its slot.

If the physical NIC that stands first in the failover order experiences intermittent failures, the failback policy might lead to frequent changes in the NIC that is used. The physical switch sees frequent changes in MAC addresses, and the physical switch port might not accept traffic immediately when an adapter becomes online. To minimize such delays, you might consider changing the following settings on the physical switch:

- Disable Spanning Tree Protocol (STP) on physical NICs that are connected to ESXi hosts .
- For Cisco based networks, enable PortFast mode for access interfaces or PortFast trunk mode for trunk interfaces. This might save about 30 seconds during the initialization of the physical switch port.
- Deactivate the trunking negotiation.

Notify Switches Policy

By using the notify switches policy, you can determine how the ESXi host communicates failover events. When a physical NIC connects to the virtual switch or when traffic is rerouted to a different physical NIC in the team, the virtual switch sends notifications over the network to update the lookup tables on physical switches. Notifying the physical switch offers lowest latency when a failover or a migration with vSphere vMotion occurs.

Load Balancing Algorithms Available for Virtual Switches

Learn to configure various load balancing algorithms on a virtual switch to determine how network traffic is distributed between the physical NICs in a team.

Route Based on Originating Virtual Port

The virtual switch selects uplinks based on the virtual machine port IDs on the vSphere Standard Switch or vSphere Distributed Switch.

Route Based on Originating Virtual Port is the default load balancing method on the vSphere Standard Switch and vSphere Distributed Switch.

Each virtual machine running on an ESXi host has an associated virtual port ID on the virtual switch. To calculate an uplink for a virtual machine, the virtual switch uses the virtual machine port ID and the number of uplinks in the NIC team. After the virtual switch selects an uplink for a virtual machine, it always forwards traffic through the same uplink for this virtual machine as long as the machine runs on the same port. The virtual switch calculates uplinks for virtual machines only once, unless uplinks are added or removed from the NIC team.

The port ID of a virtual machine is fixed while the virtual machine runs on the same host. If you migrate, power off, or delete the virtual machine, its port ID on the virtual switch becomes free. The virtual switch stops sending traffic to this port, which reduces the overall traffic for its associated uplink. If a virtual machine is powered on or migrated, it might appear on a different port and use the uplink, which is associated with the new port.

Table 188: Considerations on Using Route Based on Originating Virtual Port

Considerations	Description
Advantages	• An even distribution of traffic if the number virtual NICs is greater than the number of physical NICs in the team. • Low resource consumption, because in most cases the virtual switch calculates uplinks for virtual machines only once. • No changes on the physical switch are required.

Considerations	Description
Disadvantages	- The virtual switch is not aware of the traffic load on the uplinks and it does not load balance the traffic to uplinks that are less used. - The bandwidth that is available to a virtual machine is limited to the speed of the uplink that is associated with the relevant port ID, unless the virtual machine has more than one virtual NIC.

Route Based on Source MAC Hash

The virtual switch selects an uplink for a virtual machine based on the virtual machine MAC address. To calculate an uplink for a virtual machine, the virtual switch uses the virtual machine MAC address and the number of uplinks in the NIC team.

Table 189: Considerations on Using Route Based on Source MAC Hash

Considerations	Description
Advantages	- A more even distribution of the traffic than Route Based on Originating Virtual Port, because the virtual switch calculates an uplink for every packet. - Virtual machines use the same uplink because the MAC address is static. Powering a virtual machine on or off does not change the uplink that the virtual machine uses. - No changes on the physical switch are required.
Disadvantages	- The bandwidth that is available to a virtual machine is limited to the speed of the uplink that is associated with the relevant port ID, unless the virtual machine uses multiple source MAC addresses. - Higher resource consumption than Route Based on Originating Virtual Port, because the virtual switch calculates an uplink for every packet. - The virtual switch is not aware of the load of the uplinks, so uplinks might become overloaded.

Route Based on IP Hash

The virtual switch selects uplinks for virtual machines based on the source and destination IP address of each packet.

To calculate an uplink for a virtual machine, the virtual switch takes the last octet of both source and destination IP addresses in the packet, puts them through a XOR operation, and then runs the result through another calculation based on the number of uplinks in the NIC team. The result is a number between 0 and the number of uplinks in the team minus one. For example if a NIC team has four uplinks, the result is a number between 0 and 3 as each number is associated with a NIC in the team. For non-IP packets, the virtual switch takes two 32-bit binary values from the frame or packet from where the IP address would be located.

Any virtual machine can use any uplink in the NIC team depending on the source and destination IP address. In this way, each virtual machine can use the bandwidth of any uplink in the team. If a virtual machine runs in an environment with a large number of independent virtual machines, the IP hash algorithm can provide an even spread of the traffic between the NICs in the team. When a virtual machine communicates with multiple destination IP addresses, the virtual switch can generate a different hash for each destination IP. In this way, packets can use different uplinks on the virtual switch that results in higher potential throughput.

However, if your environment has a small number of IP addresses, the virtual switch might consistently pass the traffic through one uplink in the team. For example, if you have a database server that is accessed by one application server, the virtual switch always calculates the same uplink, because only one source-destination pair exists.

Physical Switch Configuration

To ensure that IP hash load balancing works correctly, you must have an Etherchannel configured on the physical switch. An Etherchannel bonds multiple network adapters into a single logical link. When ports are bound into an Etherchannel, every time the physical switch receives a packet from the same virtual machine MAC address on different ports, the switch updates its content addressable memory (CAM) table correctly.

For example, if the physical switch receives packets on ports 01 and 02 from MAC address A, the switch makes a 01-A and a 02-A entry in its CAM table. As a result, the physical switch distributes the incoming traffic to the correct ports. Without an Etherchannel, the physical switch first makes a record that a packet from MAC address A is received on port 01, then updates the same record that a packet from MAC address A is received on port 02. Hence, the physical switch forwards incoming traffic only on port 02, and might result in packets not reaching their destination and overloading the corresponding uplink.

Limitations and Configuration Requirements

- ESXi hosts support IP hash teaming on a single physical switch or stacked switches.
- ESXi hosts support only 802.3ad link aggregation in Static mode . You can only use a static Etherchannel with vSphere Standard Switches. LACP is not supported. If you enable IP hash load balancing without 802.3ad link aggregation and the reverse, you might experience networking disruptions.
- You must use Link Status Only as network failure detection with IP hash load balancing.
- You must set all uplinks from the team in the Active failover list . The Standby and Unused lists must be empty.
- The number of ports in the Etherchannel must be same as the number of uplinks in the team.

Considerations on Using Route Based on IP Hash

Considerations	Description
Advantages	• A more even distribution of the load compared to Route Based on Originating Virtual Port and Route Based on Source MAC Hash, as the virtual switch calculates the uplink for every packet. • A potentially higher throughput for virtual machines that communicate with multiple IP addresses.
Disadvantages	• Highest resource consumption compared to the other load balancing algorithms. • The virtual switch is not aware of the actual load of the uplinks. • Requires changes on the physical network. • Complex to troubleshoot.

Route Based on Physical NIC Load

Route Based on Physical NIC Load is based on Route Based on Originating Virtual Port, where the virtual switch checks the actual load of the uplinks and takes steps to reduce it on overloaded uplinks. Available only for vSphere Distributed Switch.

The distributed switch calculates uplinks for virtual machines by taking their port ID and the number of uplinks in the NIC team. The distributed switch tests the uplinks every 30 seconds, and if their load exceeds 75 percent of usage, the port ID of the virtual machine with the highest I/O is moved to a different uplink.

Table 190: Considerations on Using Route Based on Physical NIC Load

Considerations	Description
Advantages	• Low resource consumption because the distributed switch calculates uplinks for virtual machines only once and checking the of uplinks has minimal impact. • The distributed switch is aware of the load of uplinks and takes care to reduce it if needed. • No changes on the physical switch are required.
Disadvantages	• The bandwidth that is available to virtual machines is limited to the uplinks that are connected to the distributed switch.

Use Explicit Failover Order

No actual load balancing is available with this policy. The virtual switch always uses the uplink that stands first in the list of Active adapters from the failover order and that passes failover detection criteria. If no uplinks in the Active list are available, the virtual switch uses the uplinks from the Standby list.

Configure NIC Teaming, Failover, and Load Balancing on a vSphere Standard Switch or Standard Port Group

Learn how to configure failover order to determine how network traffic is rerouted in case of adapter failure. Include two or more physical NICs in a team to increase the network capacity of a vSphere Standard Switch or standard port group. Select a load balancing algorithm to determine how the standard switch distributes the traffic between the physical NICs in a team.

Configure NIC teaming, failover, and load balancing depending on the network configuration on the physical switch and the topology of the standard switch. See [What is Teaming and Failover Policy](#) and [Load Balancing Algorithms Available for Virtual Switches](#) for more information.

If you configure the teaming and failover policy on a standard switch, the policy is propagated to all port groups in the switch. If you configure the policy on a standard port group, it overrides the policy inherited from the switch.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Navigate to the Teaming and Failover policy for the standard switch, or standard port group.

Option	Action
Standard Switch	1. Select the switch from the list. 2. Click Edit settings and select Teaming and failover.
Standard port group	1. Select the switch where the port group resides. 2. From the switch topology diagram, select the standard port group and click Edit settings. 3. Select Teaming and failover. 4. Select Override next to the policies that you want to override.

4. From the **Load balancing** drop-down menu, specify how the virtual switch load balances the outgoing traffic between the physical NICs in a team.

Option	Description
Route based on the originating virtual port	Select an uplink based on the virtual port IDs on the switch. After the virtual switch selects an uplink for a virtual machine or

Option	Description
	a VMkernel adapter, it always forwards traffic through the same uplink for this virtual machine or VMkernel adapter.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch uses the data at those fields to compute the hash . IP-based teaming requires that the physical switch is configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Use explicit failover order	From the list of active adapters, always use the highest order uplink that passes failover detection criteria. No actual load balancing is performed with this option.

5. From the **Network failure detection** drop-down menu, select the method that the virtual switch uses for failover detection.

Option	Description
Link status only	Relies only on the link status that the network adapter provides. This option detects failures such as removed cables and physical switch power failures.
Beacon probing	Sends out and listens for beacon probes on all NICs in the team, and uses this information, in addition to link status, to determine link failure. sends beacon packets every second. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

6. From the **Notify switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

Note: If the **Notify switches** is set to **Yes**, then all the connected ports, port groups, and distributed switches are connected back to the host when vCenter Server reconnects with ESXi hosts.

Note: Set this option to **No** if a connected virtual machine is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

7.

8. From the **Failback** drop-down menu, select whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes**, the default selection, the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** for a standard port, a failed adapter is left inactive after recovery until another currently active adapter fails and must be replaced.

9. Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move uplinks into different groups.

Option	Description
Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

10. Click **OK**.

Configure NIC Teaming, Failover, and Load Balancing on a Distributed Port Group or Distributed Port

Learn how to configure failover order to determine how network traffic is rerouted in case of adapter failure. Include two or more physical NICs in a team to increase the network capacity of a distributed port group or port. Select a load balancing algorithm to determine how the distributed switch load balances the traffic between the physical NICs in a team.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

Configure NIC teaming, failover, and load balancing according with the network configuration on the physical switch and the topology of the distributed switch. See [What is Teaming and Failover Policy](#) and [Load Balancing Algorithms Available for Virtual Switches](#) for more information.

If you configure the teaming and failover policy for a distributed port group, the policy is propagated to all ports in the group. If you configure the policy for a distributed port, it overrides the policy inherited from the group.

Note: Setting fallback option is not supported with **Route based on physical NIC load** teaming policy.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Navigate the Teaming and Failover policy on the distributed port group or port.

Option	Action
Distributed port group	1. From the Actions menu, select Distributed Port Group > Manage Distributed Port Groups. 2. Select the port group and click Next. 3. Select Teaming and failover.
Distributed port	1. On the Networks tab, click Distributed Port Groups and double-click a distributed port group. 2. On the Ports tab, select a port and click Edit distributed port settings. 3. Select Teaming and failover. 4. Select Override next to the properties that you want to override.

3. From the **Load balancing** drop-down menu, specify how the virtual switch load balances the outgoing traffic between the physical NICs in a team.

Option	Description
Route based on the originating virtual port	Select an uplink based on the virtual port IDs on the switch. After the virtual switch selects an uplink for a virtual machine or a VMkernel adapter, it always forwards traffic through the same uplink for this virtual machine or VMkernel adapter.
Route based on IP hash	Select an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, the switch uses the data at those fields to compute the hash . IP-based teaming requires that the physical switch is configured with EtherChannel.
Route based on source MAC hash	Select an uplink based on a hash of the source Ethernet.
Route based on physical NIC load	Available for distributed port groups or distributed ports. Select an uplink based on the current load of the physical network adapters connected to the port group or port. If an uplink remains busy at 75 percent or higher for 30 seconds, the host

Option	Description
	proxy switch moves a part of the virtual machine traffic to a physical adapter that has free capacity. Note: Selecting Route based on physical NIC load prevents you from setting a failback option for a distributed port group.
Use explicit failover order	From the list of active adapters, always use the highest order uplink that passes failover detection criteria. No actual load balancing is performed with this option.

4. From the **Network failure detection** drop-down menu, select the method that the virtual switch uses for failover detection.

Option	Description
Link status only	Relies only on the link status that the network adapter provides. This option detects failures such as removed cables and physical switch power failures.
Beacon probing	Sends out and listens for beacon probes on all NICs in the team, and uses this information, in addition to link status, to determine link failure. sends beacon packets every second. The NICs must be in an active/active or active/standby configuration because the NICs in an unused state do not participate in beacon probing.

5. From the **Notify switches** drop-down menu, select whether the standard or distributed switch notifies the physical switch in case of a failover.

Note: If the **Notify switches** is set to **Yes**, then all the connected ports, port groups, and distributed switches are connected back to the host when vCenter Server reconnects with ESXi hosts.

Note: Set this option to **No** if a connected virtual machine is using Microsoft Network Load Balancing in unicast mode. No issues exist with Network Load Balancing running in multicast mode.

6. From the **Failback** drop-down menu, select whether a physical adapter is returned to active status after recovering from a failure.

If failback is set to **Yes**, the default selection, the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any.

If failback is set to **No** for a distributed port, a failed adapter is left inactive after recovery only if the associated virtual machine is running. When the **Failback** option is **No** and a virtual machine is powered off, if all active physical adapters fail and then one of them recovers, the virtual NIC is connected to the recovered adapter instead of to a standby one after the virtual machine is powered on. Powering a virtual machine off and then on leads to reconnecting the virtual NIC to a distributed port. The distributed switch considers the port as newly added, and assigns it the default uplink port, that is, the active uplink adapter.

7. Specify how the uplinks in a team are used when a failover occurs by configuring the Failover Order list.

If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, use the up and down arrow keys to move uplinks into different groups.

Option	Description
Active adapters	Continue to use the uplink if the network adapter connectivity is up and active.
Standby adapters	Use this uplink if one of the active physical adapter is down.
Unused adapters	Do not use this uplink.

- Review your settings and apply the configuration.

What is VLAN Policy

VLAN policies determine how VLANs function across your network environment.

A virtual local area network (VLAN) is a group of hosts with a common set of requirements, which communicate as if they were attached to the same broadcast domain, regardless of their physical location. A VLAN has the same attributes as a physical local area network (LAN), but it allows for end stations to be grouped together even if not on the same network switch.

The scope of VLAN policies can be distributed port groups and ports, and uplink port groups and ports.

Configure VLAN Tagging on a Distributed Port Group or Distributed Port

To apply VLAN tagging globally on all distributed ports, you must set the VLAN policy on a distributed port group. To integrate the virtual traffic on the port with physical VLANs in a different way from the parent distributed port group, you must use the VLAN policy on a distributed port.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

- On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
- Navigate to the VLAN policy on the distributed port group or distributed port.

Option	Action
Distributed port group	- From the Actions menu, select Distributed Port Group > Manage Distributed Port Groups. - Select VLAN and click Next. - Select the port group and click Next.
Distributed port	- On the Networks tab, click Distributed Port Groups and double-click a distributed port group. - On the Ports tab, select a port and click the Edit distributed port settings icon. - Select VLAN. - Select Override next to the properties to override.

- From the **VLAN type** drop-down menu, select the type of VLAN traffic filtering and marking, and click **Next**.

Option	Description
None	Do not use VLAN. Use this option in case of External Switch Tagging.
VLAN	Tag traffic with the ID from the VLAN ID field. Type a number between 1 and 4094 for Virtual Switch Tagging.
VLAN Trunking	Pass VLAN traffic with ID within the VLAN trunk range to guest operating system. You can set multiple ranges and individual VLANs by using a comma-separated list. For example: 1702–1705, 1848–1849. Use this option for Virtual Guest Tagging.
Private VLAN	Associate the traffic with a private VLAN created on the distributed switch.

4. Review your settings and apply the configuration.

Configure VLAN Tagging on an Uplink Port Group or Uplink Port

To configure VLAN traffic processing generally for all member uplinks, you must set the VLAN policy on an uplink port. To handle VLAN traffic through the port in a different way than for the parent uplink port group, you must set the VLAN policy on an uplink .

To override the VLAN policy at the port level, enable the port-level overrides. See [Configure Overriding Networking Policies on Port Level](#).

Use the VLAN policy at the uplink port level to propagate a trunk range of VLAN IDs to the physical network adapters for traffic filtering. The physical network adapters drop the packets from the other VLANs if the adapters support filtering by VLAN. Setting a trunk range improves networking performance because physical network adapters filter traffic instead of the uplink ports in the group.

If you have a physical network adapter that does not support VLAN filtering, the VLANs still might not be blocked. In this case, configure VLAN filtering on a distributed port group or on a distributed port.

For information about VLAN filtering support, see the technical documentation from the adapter vendors.

1. In the vSphere Client, navigate to a distributed switch.
2. On the **Networks** tab, click **Uplink Port Groups**.
3. Navigate to the VLAN policy is on the uplink port group or port.

Option	Action
Uplink port group	1. Right-click an uplink port group in the list and select Edit Settings. 2. Click VLAN.
Uplink port	1. Double-click an uplink port group. 2. On the Ports tab, select a port and click the Edit distributed port settings tab. 3. Click VLAN and select Override.

4. Enter a **VLAN trunk range** value to propagate to the physical network adapters.
For trunking of several ranges and individual VLANs, separate the entries with commas.
5. Click **OK**.

What is Security Policy

Networking security policy provides protection of traffic against MAC address impersonation and unwanted port scanning

The security policy of a standard or distributed switch is implemented in Layer 2 (Data Link Layer) of the network protocol stack. The three elements of the security policy are promiscuous mode, MAC address changes, and forged transmits. See the *vSphere Security* documentation for information about potential networking threats.

Configure the Security Policy for a vSphere Standard Switch or Standard Port Group

For a vSphere standard switch, you can configure the security policy to reject MAC address and promiscuous mode changes in the guest operating system of a virtual machine. You can override the security policy that is inherited from the standard switch on individual port groups.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Navigate to the Security policy for the standard switch or port group.

Option	Action
vSphere Standard Switch	1. Select a standard switch from the list. 2. Click Edit settings. 3. Select Security.
Standard port group	1. Select the standard switch where the port group resides. 2. In the topology diagram, select a standard port group. 3. Click Edit settings. 4. Select Security and select Override next to the options to override.

4. Reject or accept promiscuous mode activation or MAC address changes in the guest operating system of the virtual machines attached to the standard switch or port group.

Option	Description
Promiscuous mode	• Reject. The VM network adapter receives only frames that are addressed to the virtual machine. • Accept. The virtual switch forwards all frames to the virtual machine in compliance with the active VLAN policy for the port to which the VM network adapter is connected. Note: Promiscuous mode is insecure mode of operation. Firewalls, port scanners, intrusion detection systems, must run in promiscuous mode.
MAC address changes	• Reject. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter, the switch drops all inbound frames to the adapter. If the guest OS changes the effective MAC address of the virtual machine back to the MAC address of the VM network adapter, the virtual machine receives frames again. • Accept. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter, the switch allows frames to the new address to pass.
Forged transmits	• Reject. The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the <code>.vmx</code> configuration file. • Accept. The switch does not perform filtering, and permits all outbound frames.
Status	Enable or disable the MAC learning feature. The default is disabled.
Allow unicast flooding	When a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic

Option	Description
	to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, if MAC learning is enabled.
MAC Limit	The number of MAC addresses that can be learned is configurable. The maximum value is 4096 per port, which is the default.
MAC Limit Policy	The policy for when the MAC limit is reached. The options are: • Drop - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled. • Allow - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

5. Click **OK**.

Configure the Security Policy for a Distributed Port Group or Distributed Port

Learn how to set a security policy on a distributed port group to allow or reject promiscuous mode and MAC address changes from the guest operating system of the virtual machines associated with the port group. You can override the security policy inherited from the distributed port groups on individual ports.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Navigate to the Security policy for the distributed port group or port.

Option	Action
Distributed port group	1. From the Actions menu, select Distributed Port Group > Manage Distributed Port Groups. 2. Select Security and click Next. 3. Select the port group and click Next.
Distributed port	1. On the Networks tab, click Distributed Port Groups and double-click a distributed port group . 2. On the Ports tab, select a port and click the Edit settings icon. 3. Select Security. 4. Select Override next to the properties to override.

3. Reject or accept promiscuous mode activation or MAC address changes in the guest operating system of the virtual machines attached to the distributed port group or port.

Option	Description
Promiscuous mode	• Reject. The VM network adapter receives only frames that are addressed to the virtual machine. • Accept. The virtual switch forwards all frames to the virtual machine in compliance with the active VLAN policy for the port to which the VM network adapter is connected.

Option	Description
	Note: Promiscuous mode is insecure mode of operation. Firewalls, port scanners, intrusion detection systems, must run in promiscuous mode.
MAC address changes	• Reject. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter, the switch drops all inbound frames to the adapter. If the guest OS changes the effective MAC address of the virtual machine back to the MAC address of the VM network adapter, the virtual machine receives frames again. • Accept. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter, the switch allows frames to the new address to pass.
Forged transmits	• Reject. The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the <code>.vmmx</code> configuration file. • Accept. The switch does not perform filtering, and permits all outbound frames.
Status	Enable or disable the MAC learning feature. The default is disabled.
Allow unicast flooding	When a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, if MAC learning is enabled.
MAC Limit	The number of MAC addresses that can be learned is configurable. The maximum value is 4096 per port, which is the default.
MAC Limit Policy	The policy for when the MAC limit is reached. The options are: • Drop - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled. • Allow - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

4. Review your settings and apply the configuration.

What is Traffic Shaping Policy

A traffic shaping policy is defined by average bandwidth, peak bandwidth, and burst size. You can establish a traffic shaping policy for each port group and each distributed port or distributed port group.

ESXi shapes outbound network traffic on standard switches and inbound and outbound traffic on distributed switches. Traffic shaping restricts the network bandwidth available on a port, but can also be configured to allow bursts of traffic to flow through at higher speeds.

Average Bandwidth

Establishes the number of bits per second to allow across a port, averaged over time. This number is the allowed average load.

Peak Bandwidth

Maximum number of bits per second to allow across a port when it is sending or receiving a burst of traffic. This number limits the bandwidth that a port uses when it is using its burst bonus.

Burst Size

Maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus if it does not use all its allocated bandwidth. When the port needs more bandwidth than specified by the average bandwidth, it might be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter limits the number of bytes that have accumulated in the burst bonus and transfers traffic at a higher speed.

Configure Traffic Shaping for a vSphere Standard Switch or Standard Port Group

ESXi lets you shape outbound traffic on standard switches or port groups. The traffic shaper restricts the network bandwidth available to any port, but you can also configure it to temporarily allow bursts of traffic to flow through a port at higher speeds.

The traffic shaping policies that you set at switch or port group level are applied at each individual port that participates in the switch or port group. For example, if you set an average bandwidth of 100000 Kbps on a standard port group, 100000 Kbps averaged over time can pass through each port that is associated with the standard port group.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Navigate to the traffic shaping policy on the standard switch or port group.

Option	Action
vSphere Standard Switch	1. Select a standard switch from the list. 2. Click Edit settings. 3. Select Traffic shaping.
Standard port group	1. Select the standard switch where the port group resides. 2. In the topology diagram, select a standard port group. 3. Click Edit settings. 4. Select Traffic shaping and select Override next to the options to override.

4. Configure traffic shaping policies.

Option	Description
Status	Enables setting limits on the amount of networking bandwidth allocated for each port that is associated with the standard switch or port group.
Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time (the allowed average load).
Peak Bandwidth	The maximum number of bits per second to allow across a port when it is sending a burst of traffic. This setting tops the bandwidth used by a port whenever it is using its burst bonus. This parameter can never be smaller than the average bandwidth.
Burst Size	The maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than the average bandwidth specifies, the

Option	Description
	port can temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that can accumulate in the burst bonus and can be transferred at a higher speed.

- For each traffic shaping policy (**Average Bandwidth**, **Peak Bandwidth**, and **Burst Size**), enter a bandwidth value.
- Click **OK**.

Edit the Traffic Shaping Policy on a Distributed Port Group or Distributed Port

You can shape both inbound and outbound traffic on vSphere distributed port groups or distributed ports. The traffic shaper restricts the network bandwidth for any port in the group, but might also be configured to temporarily allow “bursts” of traffic to flow through a port at higher speeds.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

The traffic shaping policies that you set at distributed port group level are applied on each individual port that participates in the port group. For example, if you set an average bandwidth of 100000 Kbps on a distributed port group, 100000 Kbps averaged over time can pass through each port that is associated with the distributed port group.

- On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
- Navigate to the Traffic Shaping policy for the distributed port group or port.

Option	Action
Distributed port group	- From the Actions menu, select Distributed Port Group > Manage Distributed Port Groups. - Select Traffic shaping and click Next. - Select the port group and click Next.
Distributed port	- On the Networks tab, click Distributed Port Groups and double-click a distributed port group . - On the Ports tab, select a port and click the Edit distributed port settings icon. - Select Traffic shaping. - Select Override next to the properties to override.

- Configure traffic shaping policies.

Note: The traffic is classified to ingress and egress according to the traffic direction in the switch, not in the host.

Option	Description
Status	Enable either Ingress traffic shaping or Egress traffic shaping by using the Status drop-down menus.
Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time, that is, the allowed average load.
Peak Bandwidth	The maximum number of bits per second to allow across a port when it is sending/sending or receiving a burst of traffic. This parameter tops the bandwidth used by a port whenever it is using its burst bonus.
Burst Size	The maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than the average bandwidth specifies, the port

Option	Description
	can temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that might accumulate in the burst bonus and be transferred at a higher speed.

- Review your settings and apply the configuration.

What is Resource Allocation Policy

The Resource Allocation policy allows you to associate a distributed port or port group with a user-created network resource pool. This policy provides you with greater control over the bandwidth given to the port or port group.

For information about creating and configuring network resource pools, see [vSphere Network I/O Control](#).

Edit the Resource Allocation Policy on a Distributed Port Group

Learn how to associate a distributed port group with a network resource pool to give you greater control over the bandwidth that is given to the distributed port group.

- Enable Network I/O Control on the distributed switch. See [Enable Network I/O Control on a vSphere Distributed Switch](#).
 - Create and configure network resource pools. See [Create a Network Resource Pool](#).
- On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
 - Right-click the distributed switch in the navigator and select **Distributed Port Groups > Manage Distributed Port Groups**.
 - Select the **Resource allocation** check box and click **Next**.
 - Select the distributed port group to configure and click **Next**.
 - Add or remove the distributed port group from the network resource pool and click **Next**.
 - To add the distributed port group, select a user-defined resource pool from the **Network resource pool** drop-down menu.
 - To remove the distributed port group, select **default** from the **Network resource pool** drop-down menu.
 - Review your settings in the **Ready to complete** section and click **Finish**.
Use the **Back** button to change any settings.

What is Monitoring Policy

The monitoring policy enables or disables NetFlow monitoring on a distributed port or port group.

NetFlow settings are configured at the vSphere distributed switch level. See [Configure the NetFlow Settings of a vSphere Distributed Switch](#).

Manage NetFlow Monitoring on a Distributed Port Group or Distributed Port

Learn how to configure NetFlow to monitor IP packets that are passing through the ports of a distributed port group or through individual distributed ports.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

You configure the NetFlow settings on the vSphere Distributed Switch. See [Configure the NetFlow Settings of a vSphere Distributed Switch](#)

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Navigate to the monitoring policy for the distributed port group or distributed port.

Option	Action
Distributed port group	1. From the Actions menu, select Distributed Port Group > Manage Distributed Port Groups. 2. Select Monitoring and click Next. 3. Select the port group and click Next.
Distributed port	1. On the Networks tab, click Distributed Port Groups and double-click a distributed port group. 2. On the Ports tab, select a port and click the Edit distributed port settings icon. 3. Select Monitoring. 4. Select Override next to the properties to override.

3. From the **NetFlow** drop-down menu, select **Enabled** or **Disabled**, and click **Next**.
4. Verify your settings and apply the configuration.

What is Traffic Filtering and Marking Policy

In a vSphere distributed switch, by using the traffic filtering and marking policy, you can protect the virtual network from unwanted traffic and security attacks or apply a QoS tag to a certain type of traffic.

The traffic filtering and marking policy represents an ordered set of network traffic rules for security and for QoS tagging of the data flow through the ports of a distributed switch. In general, a rule consists of a qualifier for traffic, and of an action for restricting or prioritizing the matching traffic.

The vSphere distributed switch applies rules on traffic at different places in the data stream. The distributed switch applies traffic filter rules on the data path between the virtual machine network adapter and distributed port, or between the uplink port and physical network adapter for rules on uplinks. The traffic filtering and marking policy does not support vSphere Distributed Switch configured with Network Offloads Compatibility.

Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group

Set traffic rules at the level of distributed port groups or of uplink port groups to introduce filtering and priority tagging for traffic access over virtual machines, VMkernel adapters, or physical adapters.

Enable Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group

Enable the traffic filtering and marking policy on a port group if you want to configure traffic security and marking on all virtual machine network adapters or uplink adapters that are participating in the group.

Note: You can disable the traffic filtering and marking policy on a particular port to avoid processing the traffic flowing through the port. See [Disable Traffic Filtering and Marking on a Distributed Port or Uplink Port](#)

Note: The traffic filtering and marking policy does not support vSphere Distributed Switch configured with Network Offloads Capability.

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. Click the **Enable and reorder** button.
5. Click **Enable all traffic rules**.
6. Click **OK**.

Set up traffic marking or filtering on the data that is flowing through the ports of the distributed port group or through the uplink port group. See [Mark Traffic on a Distributed Port Group or Uplink Port Group](#) and [Filter Traffic on a Distributed Port Group or Uplink Port Group](#).

Mark Traffic on a Distributed Port Group or Uplink Port Group

Learn to assign priority tags to traffic, such as VoIP and streaming video, that has higher networking requirements for bandwidth, low latency, and so on. You can mark the traffic with a CoS tag in Layer 2 of the network protocol stack or with a DSCP tag in Layer 3.

Priority tagging is a mechanism to mark traffic that has higher QoS demands. In this way, the network can recognize different classes of traffic. The network devices can handle the traffic from each class according to its priority and requirements.

You can also re-tag traffic to either raise or lower the importance of the flow. By using a low QoS tag, you can restrict data tagged in a guest operating system.

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. If traffic filtering and marking is disabled, click **Enable and reorder** > **Enable all traffic rules** > **OK**.
5. Click **Add** to create a new rule, or select a rule and click **Edit** to edit it.
6. In the network traffic rule dialog box, select the **Tag** option from the **Action** drop-down menu.
7. Set the priority tag for the traffic within the scope of the rule.

Option	Description
CoS value	Mark the traffic matching the rule with a CoS priority tag in network Layer 2. Select the check box and type a value from 0 to 7.
DSCP value	Mark the traffic associated with the rule with a DSCP tag in network Layer 3. Select the check box and type a value from 0 to 63.

8. Specify the kind of traffic that the rule is applicable to.

To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.

- a) From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.

The direction also influences how you are going to identify the traffic source and destination.

- b) By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.

A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.

- Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group. For example, you can select NFS for data transfers to network storage.
- Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.
Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
- Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.

9. In the rule dialog box, click **OK** to save the rule.

Voice over IP Traffic Marking

Voice over IP (VoIP) flows have special requirements for QoS in terms of low loss and delay. The traffic related to the Session Initiation Protocol (SIP) for VoIP usually has a DSCP tag equal to 26, which stands for Assured Forwarding Class 3 with Low Drop Probability (AF31).

For example, to mark outgoing SIP UDP packets to a subnet 192.168.2.0/24, you can use the following rule:

Rule Parameter	Parameter Value
Action	Tag
DSCP value	26
Traffic direction	Egress
Traffic qualifiers	IP Qualifier
Protocol	UDP
Destination port	5060
Source address	IP address matches 192.168.2.0 with prefix length 24

Filter Traffic on a Distributed Port Group or Uplink Port Group

Learn how to allow or stop traffic for securing the data that flows through the ports of a distributed port group or uplink port group.

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. If traffic filtering and marking is disabled, click **Enable and reorder** > **Enable all traffic rules** > **OK**.
5. Click **Add** to create a new rule, or select a rule and click **Edit** to edit it.
6. In the network traffic rule dialog box, use the Action options to let traffic pass through the ports of the distributed port group or uplink port group, or to restrict it.
7. Specify the kind of traffic that the rule is applicable to.

To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.

- a) From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.

The direction also influences how you are going to identify the traffic source and destination.

- b) By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.

A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.

- Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group. For example, you can select NFS for data transfers to network storage.
- Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.
Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
- Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.

8. In the rule dialog box, click **OK** to save the rule.

Network Traffic Rules on a Distributed Port Group or Uplink Port Group

Learn how to define traffic rules in a distributed port group or uplink port group to introduce a policy for processing traffic related to virtual machines or to physical adapters. You can filter specific traffic or describe its QoS demands.

Note: You can override the rules of the policy for traffic filtering and marking at port level. See [Network Traffic Rules on a Distributed Port or Uplink Port](#).

View Traffic Rules on a Distributed Port Group or Uplink Group

View the traffic rules that form the traffic filtering and marking policy of a distributed port group or uplink port group.

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. If traffic filtering and marking is disabled, click **Enable and reorder** > **Enable all traffic rules** > **OK**.
5. Examine **Action** to see if the rule filters traffic (Allow or Drop) or marks traffic (Tag) with special QoS demands.
6. From the upper list, select the rule for which you want to view the criteria for locating traffic.
The traffic qualifying parameters of the rule appear in the Traffic Qualifiers list.

Edit a Traffic Rule on a Distributed Port Group or Uplink Port Group

Learn to create or edit traffic rules, and use their parameters to configure a policy for filtering or marking the traffic on a distributed port group or uplink port group.

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. If traffic filtering and marking is disabled, click **Enable and reorder** > **Enable all traffic rules** > **OK**.
5. Click **Add** to create a new rule, or select a rule and click **Edit** to edit it.

Name the network traffic rule, and deny, allow, or tag the target traffic.

Modify Rule Priorities on a Distributed Port Group or Uplink Port Group

Learn how to reorder the rules that form the traffic filtering and marking policy of a distributed port group or uplink port group to change the sequence of actions for processing traffic.

The vSphere distributed switch applies network traffic rules in a strict order. If a packet already satisfies a rule, the packet might not be passed to the next rule in the policy.

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. Click the **Enable and reorder** button.
5. If traffic filtering and marking is disabled, click **Enable all traffic rules** to enable it.
6. Select a rule and use the **Move up** or **Move down** button to change its priority.
7. Click **OK** to apply the changes.

Delete a Traffic Rule on a Distributed Port Group or Uplink Port Group

Learn how to delete a traffic rule on a distributed port group or uplink port group to stop processing packets flowing to virtual machines or physical adapters in a specific way.

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. If traffic filtering and marking is disabled, click **Enable and reorder > Enable all traffic rules > OK**.
5. Select the rule and click the **Delete** button.
6. Click **OK**.

Disable Traffic Filtering and Marking on a Distributed Port Group or Uplink Port Group

Learn how to let traffic flow to virtual machines or physical adapters without additional control related to security or QoS by disabling the traffic filtering and marking policy.

Note: You can enable and set up the traffic filtering and marking policy on a particular port. See [Enable Traffic Filtering and Marking on a Distributed Port or Uplink Port](#).

1. Locate a distributed port group or an uplink port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups** to see the list of distributed port groups, or click **Uplink Port Groups** to see the list of uplink port groups.
2. Click a distributed port group or an uplink port group and select the **Configure** tab.
3. Under Settings, select **Traffic Filtering And Marking**.
4. Click the **Enable and reorder** button.
5. Use the toggle button to disable all traffic rules.
6. Click **OK**.

Traffic Filtering and Marking on a Distributed Port or Uplink Port

Learn how to filter traffic or describe its QoS demands for an individual virtual machine, VMkernel adapter, or physical adapter by configuring the traffic filtering and marking policy on a distributed port or uplink port.

Enable Traffic Filtering and Marking on a Distributed Port or Uplink Port

Enable the traffic filtering and marking policy on a port to configure traffic security and marking on a virtual machine network adapter, VMkernel adapter, or uplink adapter. The traffic filtering and marking policy does not support vSphere Distributed Switch configured with Network Offloads Compatibility.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. Click the **Enable and reorder** button.
5. Use the toggle button to override the default settings.
6. Optional: Click **Enable all traffic rules**.
If traffic rules are enabled on group level, after you override the default settings for the port, the traffic rules are automatically enabled.
7. Click **OK**.

Set up traffic filtering or marking for the data flowing through the distributed port or through the uplink port. See [Mark Traffic on a Distributed Port or Uplink Port](#) and [Filter Traffic on a Distributed Port or Uplink Port](#).

Mark Traffic on a Distributed Port or Uplink Port

Assign priority tags in a rule for traffic that needs special treatment such as VoIP and streaming video. You can mark the traffic for a virtual machine, VMkernel adapter, or physical adapter with a CoS tag in Layer 2 of the network protocol stack or with a DSCP tag in Layer 3.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

Priority tagging is a mechanism to mark traffic that has higher QoS demands. In this way, the network can recognize different classes of traffic. The network devices can handle the traffic from each class according to its priority and requirements.

You can also re-tag traffic to either raise or lower the importance of the flow. By using a low QoS tag, you can restrict data tagged in a guest operating system.

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.

2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. If traffic filtering and marking is not enabled at the port level, click the Enable and reorder button, override the default settings and click **Enable all traffic rules**.
If traffic rules are enabled at group level, after you override the default settings for the port, the traffic rules are automatically enabled.
5. Click **Add** to create a new rule, or select a rule and click **Edit** to edit it.
You can change a rule inherited from the distributed port group or uplink port group. In this way, the rule becomes unique within the scope of the port.
6. In the network traffic rule dialog box, select the **Tag** option from the **Action** drop-down menu.
7. Set the priority tag for the traffic within the scope of the rule.

Option	Description
CoS value	Mark the traffic matching the rule with a CoS priority tag in network Layer 2. Select the check box and type a value from 0 to 7.
DSCP value	Mark the traffic associated with the rule with a DSCP tag in network Layer 3. Select the check box and type a value from 0 to 63.

8. Specify the kind of traffic that the rule is applicable to.
To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.
 - a) From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.
The direction also influences how you are going to identify the traffic source and destination.
 - b) By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.
A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.
 - Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group. For example, you can select NFS for data transfers to network storage.
 - Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.
Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
 - Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.
9. In the rule dialog box, click **OK** to save the rule.

Filter Traffic on a Distributed Port or Uplink Port

By using a rule, permit or stop traffic for securing data flows through a virtual machine, a VMkernel adapter, or a physical adapter.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. If traffic filtering and marking is not enabled at the port level, click the Enable and reorder button, override the default settings and click **Enable all traffic rules**.
 If traffic rules are enabled at group level, after you override the default settings for the port, the traffic rules are automatically enabled.
5. Click **Add** to create a new rule, or select a rule and click **Edit** to edit it.
 You can change a rule inherited from the distributed port group or uplink port group. In this way, the rule becomes unique within the scope of the port.
6. In the network traffic rule dialog box, select the **Allow** action to let traffic pass through the distributed port or uplink port, or the **Drop** action to restrict it.
7. Specify the kind of traffic that the rule is applicable to.
 To determine if a data flow is in the scope of a rule for marking or filtering, the vSphere distributed switch examines the direction of the traffic, and properties like source and destination, VLAN, next level protocol, infrastructure traffic type, and so on.
 - a) From the **Traffic direction** drop-down menu, select whether the traffic must be ingress, egress, or both so that the rule recognizes it as matching.
 The direction also influences how you are going to identify the traffic source and destination.
 - b) By using qualifiers for system data type, Layer 2 packet attributes, and Layer 3 packet attributes set the properties that packets must have to match the rule.
 A qualifier represents a set of matching criteria related to a networking layer. You can match traffic to system data type, Layer 2 traffic properties, and Layer 3 traffic properties. You can use the qualifier for a specific networking layer or can combine qualifiers to match packets more precisely.
 - Use the system traffic qualifier to match packets to the type of virtual infrastructure data that is flowing through the ports of the group. For example, you can select NFS for data transfers to network storage.
 - Use the MAC traffic qualifier to match packets by MAC address, VLAN ID, and next level protocol.
 Locating traffic with a VLAN ID on a distributed port group works with Virtual Guest Tagging (VGT). To match traffic to VLAN ID if Virtual Switch Tagging (VST) is active, use a rule on an uplink port group or uplink port.
 - Use the IP traffic qualifier to match packets by IP version, IP address, and next level protocol and port.
8. In the rule dialog box, click **OK** to save the rule.

Network Traffic Rules on a Distributed Port or Uplink Port

Learn how to define traffic rules in a distributed port or uplink port group to introduce a policy for processing traffic related to a virtual machine or to a physical adapter. You can filter specific traffic or describe its QoS demands.

View Traffic Rules on a Distributed Port or Uplink Port

Review the traffic rules that form the traffic filtering and marking policy of a distributed port or uplink port.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. If traffic filtering and marking is not enabled at the port level, click the Enable and reorder button, override the default settings and click **Enable all traffic rules**.
If traffic rules are enabled at group level, after you override the default settings for the port, the traffic rules are automatically enabled.
5. Examine **Action** to see if the rule filters traffic (Allow or Drop) or marks traffic (Tag) with special QoS demands.
6. From the upper list, select the rule for which you want to view the criteria for locating traffic.
The traffic qualifying parameters of the rule appear in the Traffic Qualifiers list.

Edit a Traffic Rule on a Distributed Port or Uplink Port

Learn how to create or edit traffic rules, and use their parameters to configure a policy for filtering or marking the traffic on a distributed port or uplink port.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. If traffic filtering and marking is not enabled at the port level, click the Enable and reorder button, override the default settings and click **Enable all traffic rules**.
If traffic rules are enabled at group level, after you override the default settings for the port, the traffic rules are automatically enabled.
5. Click **Add** to create a new rule, or select a rule and click **Edit** to edit it.
You can change a rule inherited from the distributed port group or uplink port group. In this way, the rule becomes unique within the scope of the port.

Name the network traffic rule, and deny, allow, or tag the target traffic.

Modify Rule Priorities on a Distributed Port or Uplink Port

Learn how to reorder the rules that form the traffic filtering and marking policy of a distributed port or uplink port to change the sequence of actions for analyzing traffic for security and QoS.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

The vSphere distributed switch applies network traffic rules in a strict order. If a packet already satisfies a rule, the packet might not be passed to the next rule in the policy.

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. Click the **Enable and reorder** button.
5. If traffic filtering and marking is not enabled at the port level, override the default settings and click **Enable all traffic rules**.
If traffic rules are enabled at group level, after you override the default settings for the port, the traffic rules are automatically enabled.
6. Select a rule and use the **Move up** or **Move down** button to change its priority.
7. Click **OK** to apply the changes.

Delete a Traffic Rule on a Distributed Port or Uplink Port

Delete a traffic rule on a distributed port or uplink port to stop filtering or marking certain type of packets that are flowing to a virtual machine or a physical adapter.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. If traffic filtering and marking is not enabled at the port level, click the **Enable and reorder** button, override the default settings and click **Enable all traffic rules**.
If traffic rules are enabled at group level, after you override the default settings for the port, the traffic rules are automatically enabled.
5. Select the rule and click the **Delete** button.
6. Click **OK**.

Disable Traffic Filtering and Marking on a Distributed Port or Uplink Port

Disable the traffic filtering and marking policy on a port to let traffic flow to a virtual machine or a physical adapter without filtering for security or marking for QoS.

To override a policy on distributed port level, enable the port-level override option for this policy. See [Configure Overriding Networking Policies on Port Level](#).

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Select the **Traffic Filtering and Marking** tab.
4. Click the **Enable and reorder** button.
5. Use the toggle buttons to override the default settings at the port level and disable all traffic rules.
6. Click **OK**.

Qualifying Traffic for Filtering and Marking

The traffic that you want to filter or want to mark with QoS tags can be matched to the type of carried infrastructure data, such as data for storage, vCenter Server management, and so on, and to Layer 2 and Layer 3 properties.

To match the traffic in the scope of the rule more precisely, you can combine criteria for system data type, Layer 2 header, and Layer 3 header.

System Traffic Qualifier

By using the system traffic qualifier in a rule for port group or port, you can determine whether certain system data traffic must be marked with a QoS tag, allowed, or dropped.

System Traffic Type

You can select the type of traffic through the ports of the group that carries system data, that is, traffic for management from vCenter Server, storage, VMware vSphere®vMotion®, and vSphere Fault Tolerance. You can mark or filter only a specific traffic type, or for all system data traffic except for a infrastructure feature. For example, you can mark with a QoS value or filter the traffic for management from vCenter Server , storage and vMotion, but not the traffic carrying the Fault Tolerance data.

MAC Traffic Qualifier

By using the MAC traffic qualifier in a rule, you can define matching criteria for the Layer 2 (Data Link Layer) properties of packets such as MAC address, VLAN ID, and next level protocol that consumes the frame payload.

Protocol Type

The **Protocol type** attribute of the MAC traffic qualifier corresponds to the EtherType field in Ethernet frames. EtherType represents the type of next level protocol that is going to consume the payload of the frame.

You can select a protocol from the drop-down menu or type its hexadecimal number. For example, to capture traffic for the Link Layer Discovery Protocol (LLDP) protocol, type 88CC.

VLAN ID

You can use the VLAN ID attribute of the MAC traffic qualifier to mark or filter traffic in a particular VLAN.

Note:

The VLAN ID qualifier on a distributed port group works with Virtual Guest Tagging (VGT).

If a flow is tagged with a VLAN ID through Virtual Switch Tagging (VST), it cannot be located by using this ID in a rule on a distributed port group or distributed port. The reason is that the distributed switch checks the rule conditions, including the VLAN ID, after the switch has already untagged the traffic. In this case, to match traffic by VLAN ID successfully, you must use a rule on an uplink port group or uplink port.

Source Address

By using the Source Address group of attributes, you can match packets by the source MAC address or network.

You can use a comparison operator to mark or filter packets that have or do not have the specified source address or network.

You can match the traffic source in several ways.

Table 191: Patterns for Filtering or Marking Traffic by MAC Source Address

Parameters to Match Traffic Source Address	Comparison Operator	Networking Argument Format
MAC address	is or is not	Type the MAC address for matching. Use colons to separate the octets in it.
MAC network	matches or does not match	Type the lowest address in the network and a mask. Set ones at the positions of the network bits, and zeroes for the host part.

For example, for a MAC network with prefix 05:50:56 that is 23 bits long, set the address as 00:50:56:00:00:00 and mask as ff:ff:fe:00:00:00.

Destination Address

By using the Destination Address group of attributes, you can match packets to their destination address. The MAC destination address options have the same format as those for the source address.

Comparison Operators

To match traffic in a MAC qualifier more closely to your needs, you can use affirmative comparison or negation. You can use operators such that all packets except the ones with certain attributes fall in the scope of a rule.

IP Traffic Qualifier

By using the IP traffic qualifier in a rule, you can define criteria for matching traffic to the Layer 3 (Network Layer) properties such as IP version, IP address, next level protocol, and port.

Protocol

The **Protocol** attribute of the IP traffic qualifier represents the next level protocol consuming the payload of the packet. You can select a protocol from the drop-down menu or type its decimal number according to RFC 1700.

For the TCP and UDP protocols, you can also match traffic by source and destination ports.

Source Port

By using the Source port attribute, you can match TCP or UDP packets by the source port. Consider the traffic direction when matching traffic to a source port.

Destination Port

By using the Destination port attribute, you can match TCP or UDP packets by the destination port. Consider the traffic direction when matching traffic to a destination port.

Source Address

By using the Source Address attribute, you can match packets by source address or subnet. Consider the traffic direction when matching traffic to a source address or network.

You can match traffic source in several ways.

Table 192: Patterns for Filtering or Marking Traffic by IP Source Address

Parameters to Match Traffic Source Address	Comparison Operator	Networking Argument Format
IP version	any	Select the IP version from the drop-down menu.
IP address	is or is not	Type the IP address that you want to match.
IP subnet	matches or does not match	Type the lowest address in the subnet and the bit length of the subnet prefix.

Destination Address

Use the Destination Address to match packets by IP address, subnet, or IP version. The destination address has the same format as the one for the source.

Comparison Operators

To match traffic in an IP qualifier more closely to your needs, you can use affirmative comparison or negation. You can define that all packets fall in the scope of a rule except packets with certain attributes.

Manage Policies for Multiple Port Groups on a vSphere Distributed Switch

Learn how to modify networking policies for multiple port groups on a vSphere Distributed Switch.

Create a vSphere Distributed Switch with one or more port groups.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch in the object navigator and select **Distributed Port Group > Manage Distributed Port Groups**.
3. On the Select port group policies page, select the check box next to the policy categories to modify and click **Next**.

Option	Description
Security	Set MAC address changes, forged transmits, and promiscuous mode for the selected port groups.

Option	Description
Traffic shaping	Set the average bandwidth, peak bandwidth, and burst size for inbound and outbound traffic on the selected port groups.
VLAN	Configure how the selected port groups connect to physical VLANs.
Teaming and failover	Set load balancing, failover detection, switch notification, and failover order for the selected port groups.
Resource allocation	Set network resource pool association for the selected port groups.
Monitoring	Enable or disable NetFlow on the selected port groups.
Miscellaneous	Enable or disable port blocking on the selected port groups.

4. On the Select port groups page, select the distributed port group(s) to edit and click **Next**.
5. Optional: On the Security page, use the drop-down menus to edit the security exceptions and click **Next**.

Option	Description
Promiscuous mode	• Reject. Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter. • Accept. Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSphere Distributed Switch that are allowed under the VLAN policy for the port group that the adapter is connected to.
MAC address changes	• Reject. If set to Reject and the guest operating system changes the MAC address of the adapter to anything other than what is in the <code>.vmx</code> configuration file, all inbound frames are dropped. If the Guest OS changes the MAC address back to match the MAC address in the <code>.vmx</code> configuration file, inbound frames are passed again. • Accept. Changing the MAC address from the Guest OS has the intended effect. Frames to the new MAC address are received.
Forged transmits	• Reject. Any outbound frame with a source MAC address that is different from the one currently set on the adapter are dropped. • Accept. No filtering is performed and all outbound frames are passed.

6. Optional: On the VLAN page, use the drop-down menus to edit the VLAN policy and click **Next**.

Option	Description
None	Do not use VLAN.
VLAN	In the VLAN ID field, enter a number between 1 and 4094.
VLAN trunking	Enter a VLAN trunk range .
Private VLAN	Select an available private VLAN to use.

7. Optional: On the Traffic shaping page, use the drop-down menus to enable or disable Ingress or Egress traffic shaping and click **Next**.

Option	Description
Status	If you enable either Ingress traffic shaping or Egress traffic shaping , you are setting limits on the amount of networking bandwidth allocated for each VMkernel adapter or virtual network adapter associated with this port group. If you disable

Option	Description
	the policy, services have a free, clear connection to the physical network by default.
Average bandwidth	Establishes the number of bits per second to allow across a port, averaged over time, that is, the allowed average load.
Peak bandwidth	The maximum number of bits per second to allow across a port when it is sending or receiving a burst of traffic. This maximum number tops the bandwidth used by a port whenever it is using its burst bonus.
Burst size	The maximum number of bytes to allow in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by Average bandwidth , it might be allowed to transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that can be accumulated in the burst bonus and transferred at a higher speed.

8. Optional: On the Teaming and failover page, use the drop-down menus to edit the settings and click **Next**.

Option	Description
Load balancing	IP-based teaming requires that the physical switch be configured with ether channel. For all other options, ether channel should be disabled. Select how to choose an uplink. • Route based on the originating virtual port. Choose an uplink based on the virtual port where the traffic entered the distributed switch. • Route based on IP hash. Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash. • Route based on source MAC hash. Choose an uplink based on a hash of the source Ethernet. • Route based on physical NIC load. Choose an uplink based on the current loads of physical NICs. • Use explicit failover order. Always use the highest order uplink, from the list of Active adapters, which passes failover detection criteria.
Network failure detection	Select the method to use for failover detection. • Link status only. Relies solely on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch. • Beacon probing. Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. Do not use beacon probing with IP-hash load balancing.
Notify switches	Select Yes or No to notify switches in the case of failover. Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode.

Option	Description
	If you select Yes, whenever a virtual NIC is connected to the distributed switch or whenever that virtual NIC's traffic is routed over a different physical NIC in the team because of a failover event, a notification is sent out over the network to update the lookup tables on physical switches. Use this process for the lowest latency of failover occurrences and migrations with vMotion. If the Notify switches is set to Yes, then all the connected ports, port groups, and distributed switches are connected back to the host when vCenter Server reconnects with ESXi hosts.
Failback	Select Yes or No to disable or enable failback. This option determines how a physical adapter is returned to active duty after recovering from a failure. • Yes (default). The adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. • No. A failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.
Failover order	Select how to distribute the work load for uplinks. To use some uplinks but reserve others in case the uplinks in use fail, set this condition by moving them into different groups. • Active uplinks. Continue to use the uplink when the network adapter connectivity is up and active. • Standby uplinks. Use this uplink if one of the active adapter's connectivity is down. When using IP-hash load balancing, do not configure standby uplinks. • Unused uplinks. Do not use this uplink.

9. Optional: On the Resource allocation page, use the **Network resource pool** drop-down menu to add or remove resource allocations and click **Next**.

10. Optional: On the Monitoring page, use the drop-down menu to enable or disable NetFlow and click **Next**.

Option	Description
Disabled	NetFlow is disabled on the distributed port group.
Enabled	NetFlow is enabled on the distributed port group. You can configure NetFlow settings at the vSphere Distributed Switch level.

11. Optional: On the Miscellaneous page, select **Yes** or **No** from the drop-down menu and click **Next**.

Select **Yes** to shut down all ports in the port group. This shutdown might disrupt the normal network operations of the hosts or virtual machines using the ports.

12. Review your settings on the Ready to complete page and click **Finish**.

Use the **Back** button to change any settings.

Port Blocking Policies

Port blocking policies allow you to selectively block ports from sending or receiving data.

Edit the Port Blocking Policy for a Distributed Port Group

You can block all ports in a distributed port group.

Blocking the ports of a distributed port group might disrupt the normal network operations of the hosts or virtual machines using the ports.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch in the object navigator and select **Distributed Port Group > Manage Distributed Port Groups**.
3. Select the **Miscellaneous** check box and click **Next**.
4. Select one or more distributed port group to configure and click **Next**.
5. From the **Block all ports** drop-down menu, enable or disable port blocking, and click **Next**.
6. Review your settings and click **Finish**.

Edit the Blocking Policy for a Distributed Port or Uplink Port

You can block an individual distributed port or uplink port.

Enable the port-level overrides. See [Configure Overriding Networking Policies on Port Level](#)

Blocking the flow through a port might disrupt the normal network operations on the host or virtual machine using the port.

1. Navigate to a distributed switch and then navigate to a distributed port or an uplink port.
 - To navigate to the distributed ports of the switch, click **Networks > Distributed Port Groups**, click a distributed port group from the list, and click the **Ports** tab.
 - To navigate to the uplink ports of an uplink port group, click **Networks > Uplink Port Groups**, click an uplink port group from the list, and click the **Ports** tab.
2. Select a port from the list.
3. Click **Edit distributed port settings**.
4. In the **Miscellaneous** section, select the **Override** check box, and from the drop-down menu enable or disable port blocking.
5. Click **OK**.

What is MAC Learning Policy

MAC learning provides network connectivity to deployments where multiple MAC addresses are used from one vNIC.

For example, in a nested hypervisor deployment where an ESXi VM runs on an ESXi host and multiple VMs run inside the ESXi VM. Without MAC learning, when the ESXi VM's vNIC connects to a switch port, it only contains a static MAC address. VMs running inside the ESXi VM do not have network connectivity because their packets have different source MAC addresses. With MAC learning, the vSwitch inspects the source MAC address of every packet coming from the vNIC, learns the MAC address in its MAC table, and allows the packet to pass. If a MAC address that is learned is not used for a certain period of time, it is removed.

MAC learning also supports unknown unicast flooding. Normally, when a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is enabled by default, but only if MAC learning is enabled.

The number of MAC addresses that can be learned is configurable. The maximum value is 4096 per port, which is the default. You can also set the policy for when the limit is reached. The options are:

- Drop - Packets from an unknown source MAC address are dropped. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.
- Allow - Packets from an unknown source MAC address are forwarded although the address will not be learned. Packets inbound to this MAC address will be treated as unknown unicast. The port will receive the packets only if it has unknown unicast flooding enabled.

In vSphere 6.7 and later, MAC Learning can be enabled on a Distributed Virtual Port group using the vSphere API. You can configure the MAC Learning Policy on vSphere Distributed Switch, Distributed Virtual Port Group, and Distributed Virtual Port. If a MAC learning policy is not set on Distributed Virtual Port Group, it is inherited from vSphere Distributed Switch and if it is not enabled on the DVport, it is inherited from Distributed Virtual Port Group. See *vSphere Web Services API Reference* for more information.

How to Use VLANs to Isolate Network Traffic

VLANs let you segment a network into multiple logical broadcast domains at Layer 2 of the network protocol stack.

VLAN Configuration

Virtual LANs (VLANs) enable a single physical LAN segment to be further isolated so that groups of ports are isolated from one another as if they were on physically different segments.

Benefits of Using VLANs in vSphere

The VLAN configuration in a vSphere environment provides certain benefits.

- Integrates ESXi hosts into a pre-existing VLAN topology.
- Isolates and secures network traffic.
- Reduces congestion of network traffic.

VLAN Tagging Modes

vSphere supports three modes of VLAN tagging in ESXi: External Switch Tagging (EST), Virtual Switch Tagging (VST), and Virtual Guest Tagging (VGT).

Tagging Mode	VLAN ID on switch port groups	Description
EST	0	The physical switch performs the VLAN tagging. The host network adapters are connected to access ports on the physical switch.
VST	Between 1 and 4094	The virtual switch performs the VLAN tagging before the packets leave the host. The host network adapters must be connected to trunk ports on the physical switch.
VGT	• 4095 for standard switch • Range of and individual VLANs for distributed switch	The virtual machine performs the VLAN tagging. The virtual switch preserves the VLAN tags when it forwards the packets between the virtual machine networking stack and external switch. The host network adapters must be connected to trunk ports on the physical switch.

Tagging Mode	VLAN ID on switch port groups	Description
		The vSphere Distributed Switch supports a modification of VGT. For security reasons, you can configure a distributed switch to pass only packets that belong to particular VLANs. Note: For VGT you must have an 802.1Q VLAN trunking driver installed on the guest operating system of the virtual machine.

Private VLANs

Private VLANs are used to solve VLAN ID limitations by adding a further segmentation of the logical broadcast domain into multiple smaller broadcast subdomains.

A private VLAN is identified by its primary VLAN ID. A primary VLAN ID can have multiple secondary VLAN IDs associated with it. Primary VLANs are **Promiscuous**, so that ports on a private VLAN can communicate with ports configured as the primary VLAN. Ports on a secondary VLAN can be either **Isolated**, communicating only with promiscuous ports, or **Community**, communicating with both promiscuous ports and other ports on the same secondary VLAN.

To use private VLANs between a host and the rest of the physical network, the physical switch connected to the host needs to be private VLAN-capable and configured with the VLAN IDs being used by ESXi for the private VLAN functionality. For physical switches using dynamic MAC+VLAN ID based learning, all corresponding private VLAN IDs must be first entered into the switch's VLAN database.

Create a Private VLAN

Create the necessary private VLANs on the vSphere Distributed Switch to be able to assign distributed ports to participate to a private VLAN.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Settings** and select **Private VLAN**.
3. Click **Edit**.
4. To add a primary VLAN, above Primary VLAN ID click the **plus sign (+)** button.
The primary private VLAN also appears under Secondary Private VLAN ID.
5. To add a secondary VLAN, in the right pane click the **plus sign (+)** button.
6. From the drop-down menu in the **Secondary VLAN type** column, select either **Isolated** or **Community**.
7. Click **OK**.

Configure a distributed port group or port to associate traffic with the private VLAN. See [Configure VLAN Tagging on a Distributed Port Group or Distributed Port](#).

Remove a Primary Private VLAN

Remove unused primary VLANs from the configuration of a vSphere Distributed Switch.

Verify that no port groups are configured to use the primary VLAN and its associated secondary VLANs.

When you remove a primary private VLAN, you also remove the associated secondary private VLANs.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Settings** and select **Private VLAN**.
3. Click **Edit**.
4. Select the primary private VLAN to remove.
5. Click the **times sign (x)** button above the Primary VLAN ID list.
6. Click **OK**.

Remove a Secondary Private VLAN

Remove unused secondary private VLANs from the configuration of a vSphere Distributed Switch.

Verify that no port groups are configured to use the secondary VLAN.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Settings** and select **Private VLAN**.
3. Click **Edit**.
4. Select a primary private VLAN.
The secondary private VLANs associated with it appear on the right.
5. Select the secondary private VLAN to remove.
6. Above the secondary VLAN ID list, click the **times sign (x)** button and click **OK**.

Manage Network Resources

vSphere provides several different methods to help you manage your network resources.

DirectPath I/O

DirectPath I/O allows virtual machine access to physical PCI functions on platforms with an I/O Memory Management Unit.

The following features are unavailable for virtual machines configured with DirectPath:

- Hot adding and removing of virtual devices
- Suspend and resume
- Record and replay
- Fault tolerance
- High availability
- DRS (limited availability. The virtual machine can be part of a cluster, but cannot migrate across hosts)
- Snapshots

Enable Passthrough for a Network Device on a Host

Passthrough devices provide the means to use resources efficiently and improve performance of your environment. You can enable DirectPath I/O passthrough for a network device on a host.

CAUTION: If your ESXi host is configured to boot from a USB device or an SD card attached to a USB channel, make sure that you do not enable DirectPath I/O passthrough for the USB controller. Passing through a USB controller on

an ESXi host that boots from a USB device or SD card might put the host in a state where its configuration cannot be persisted.

1. Browse to a host in the vSphere Client navigator.
2. On the **Configure** tab, expand **Hardware** and click **PCI Devices**.
3. To enable DirectPath I/O passthrough for a PCI network device on the host, click **Edit**.

A list of available passthrough devices appears.

Icon	Description
green icon	A device is active and can be enabled.
orange icon	The state of the device has changed, and you must reboot the host before you can use the device.

4. Select the network device to be used for passthrough and click **OK**.
The selected PCI device appears in the table. Device information is displayed at the bottom of the screen.

Configure a PCI Device on a Virtual Machine

Passthrough devices provide the means to more efficiently use resources and improve performance in your environment. You can configure a passthrough PCI device on a virtual machine in the vSphere Client.

Verify that a passthrough networking device is configured on the host of the virtual machine. See [Enable Passthrough for a Network Device on a Host](#).

When using passthrough devices with a Linux kernel version 2.6.20 or earlier, avoid MSI and MSI-X modes because these modes have significant performance impact.

1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
2. Power off the virtual machine.
3. From the Actions menu, select Edit Settings.
4. Select the **Virtual Hardware** tab in the dialog box displaying the settings.
5. Expand the **Memory** section, and set the **Limit** to **Unlimited**.
6. Click the **Add new device** button and under Other devices, select **PCI Device**.
The **New PCI device** drop-down menu is added to the list in the **Virtual Hardware** tab.
7. From the **New PCI device** drop-down menu select the passthrough device to use, and click **OK**.
8. Power on the virtual machine.

Adding a DirectPath I/O device to a virtual machine sets memory reservation to the memory size of the virtual machine.

Hot-add and Hot-remove support for VMDirectPath I/O Devices

Passthrough devices provide the means to use resources efficiently and improve the performance of your environment.

In vSphere 8.0 and earlier, passthrough devices were added or removed only when the VM was powered off. The current version of vSphere offers hot-add and hot-remove support for VMDirectPath I/O devices. Hot-add refers to the procedure of adding a passthrough device when the VM is powered on.

The following features are available for hot-add and hot-remove support for VMDirectPath I/O devices.

- Hot-add of a passthrough enabled device to a VM by vSphere API. You can perform the operation in one of the following ways:
 - a. You enable the passthrough for an existing device that is discovered at ESXi boot time and hot-add the device to the VM.
 - b. Alternatively, you can physically hot-add a new device to an empty PCIe hot-plug slot of the host, enable passthrough for the newly added device, and then hot-add the device to the VM.

Note: This feature requires PCIe Native hot plug certification for the server.
- Hot-remove of a passthrough device from a VM by vSphere API.
- Surprise hot-remove of passthrough device which is in use by a VM, can be physically pulled out of the device from the host.

Note: This feature requires PCIe Native surprise hot plug certification for the server.

Note: Hot-add and hot-remove of VMDirectPath I/O is supported for only NVMe devices. Support for hot-add and hot-remove of VMdirectPathI/O is not available in the vSphere Client.

The following features are not supported for VMs with hot plug of passthrough enabled"

- Input–output Memory Management Unit (IOMMU) exposed to the guest operating system.
- CPU hot plug.
- Memory hot plug.
- Dynamic Direct path I/O.
- Support for VMs with the number of vCPUs > 128.
- Virtual interrupts also known as interrupt posting might be disabled for VMs with the hot plug of the passthrough enabled.
- Graceful physical hot-removal of passthrough device from the host is not supported.
- UI support for hot-add and hot-remove.
- All features that are unavailable with VM Direct Path I/O.

Requirements for Platforms, Devices, and Guest Operating Systems

Server and Device Requirements

- Ensure that hot-add and hot-remove of VMDirectPath I/O devices feature is supported by your server Original equipment manufacturer (OEM).
- The server platform firmware must be Unified Extensible Firmware Interface (UEFI).
- Ensure that the server platform and device conforms to the requirements mentioned in the [KB 2142307](#).
- The server should be certified for the VM Direct Path I/O.
- For surprise hot-remove of a passthrough device from a VM the server model should be certified for PCIe Native surprise hot plug.

Note: For more information on server certification, see [VMware Compatibility Guide](#).

Guest Operating Systems Requirements

- Ensure that the Guest Operating System has support of NVMe hot plug and NVMe surprise hot-removal.
- Use the latest and stable GOS distribution that contains all the fixes of NVMe driver I/O stack.

The vCenter and ESXi Requirements

The ESXi and the vCenter versions must be 8.0 update 1 or later.

Limits

- Maximum number of passthrough devices that can be hot-added in a single `ReconfigureVM` API call is 1.

- Maximum number of passthrough devices that can be hot-removed in a single `ReconfigureVM` API call is 1.
- Other VM reconfiguration changes cannot be requested in the same `ReconfigureVM` API call when hot-add or hot-remove a passthrough device is requested.
- Maximum number of NVMe passthrough devices that a VM can support is 32.

Enable Hot-Add and Hot-Remove of VMDirectPath I/O Devices

You can enable hot-add and hot-remove of VMDirectPath I/O devices by preparing your vCenter and the ESXi host.

Prepare the vCenter and the ESXi host.

- Deploy the vCenter and stateful install ESXi with the build combination.
- Set the VMKernel boot parameter.

```
set -s maxIntrCookies -v 4096
```

- Restart the ESXi host.
- Deploy the data center and add the host.

1. Browse to a host in the vSphere Client.
2. On the **Configure** tab, expand **Hardware** and click **PCI Devices**.
3. Select the **NVMe PCI** device and press **Toggle Passthrough**.

Prepare and configure the VM.

Prepare and configure VM

Passthrough devices provide the means to use resources efficiently and improve the performance of your environment. You can prepare and configure your virtual machine.

Deploy and configure the VM (UEFI boot)

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and select **Edit Settings**.
3. Expand the **Memory** section and set the **Reserve all guest memory (All locked)**.
4. Click **Save**.
5. Set **fixedPassthruHotPlugEnabled** to **TRUE** for the VM using vSphere API `ReconfigVM_Task()` For more information, see <https://developer.broadcom.com>.

Note: Search for the term *Hotplug* using the search bar in <https://developer.broadcom.com>.

- a) On vCenter, **Right click on VM** → **Compatibility** → **Upgrade VM Compatibility** to **vmx-20** and click save.
- b) Set `motherboardLayout` to `ACPI` using the vSphere API. For more information, see <https://developer.broadcom.com>.

Note: Search for the term *Hotplug* using the search bar in <https://developer.broadcom.com>.

6. (Optional) Add the passthrough enabled PCI devices to the VM.
7. Power on the VM.

Note: Verify if the ongoing hot-plug operation has succeeded before proceeding with the next hot-plug operation.

8. Perform hot-add and hot-remove of passthrough devices using vSphere API. For more information, see <https://developer.broadcom.com>.

Note: Search for the term *Hotplug* using the search bar in <https://developer.vmware.com/samples>.

Verify if the hot plug operation has succeeded.

Verify If the Hot Plug Operation is Successful using vSphere Client

To check if your hot-add, hot-remove, and surprise hot-remove operation has succeeded, you can verify using the vSphere Client.

1. Locate the virtual machine in the vSphere Client.
2. Select a data center, folder, cluster, resource pool, or host and click the VMs tab.
3. Click **Virtual Machines** and click the virtual machine from the list.
4. In the **PCI Devices** table, verify the following:

Option	Description
Verify if the hot-add of a passthrough -enabled device to a VM using vSphere API is successful.	The hot-added passthrough device is displayed.
Verify if the hot-remove of a passthrough device from a VM using vSphere API is successful.	The hot-removed passthrough device is not displayed.
Verify if the surprise hot-remove (physical hot-remove) of a passthrough device of a VM from the ESXi host is successful.	• Surprise hot-removed device is not displayed. • In the Host View, a. Browse to a host in the vSphere Client. b. On the Configure tab, expand Hardware, and click PCI Devices. c. Verify that the surprise hot-removed device is not displayed.

Verify If the Hot Plug Operation is Successful using VM

To check if your hot-add, hot-remove, and surprise hot-remove operation has succeeded, you can verify using the Virtual Machine.

1. In the Virtual Machine, open an SSH connection to the Guest Operating System.
2. Log in as the root user.
3. Run the command `lspci`.
4. Verify the results.

Option	Description
Verify if the hot-add of a passthrough -enabled device to a VM using vSphere API is successful.	The hot-added passthrough device is displayed. Note: The passthrough device might have a different Spotfire Binary Data File (SBDF) address than the one on the host.
Verify if the hot-remove of a passthrough device from a VM by using vSphere API is successful.	The hot-removed passthrough device is not displayed.
Verify if the surprise hot-remove (physical hot-remove) of a passthrough device of a VM from the ESXi host is successful.	The surprise hot-removed device is not displayed.

What is Single Root I/O Virtualization (SR-IOV)

vSphere supports Single Root I/O Virtualization (SR-IOV). You can use SR-IOV for networking of virtual machines that are latency sensitive or require more CPU resources.

Overview of SR-IOV

SR-IOV is a specification that allows a single Peripheral Component Interconnect Express (PCIe) physical device under a single root port to appear as multiple separate physical devices to the hypervisor or the guest operating system.

SR-IOV uses physical functions (PFs) and virtual functions (VFs) to manage global functions for the SR-IOV devices. PFs are full PCIe functions that are capable of configuring and managing the SR-IOV functionality. It is possible to configure or control PCIe devices using PFs, and the PF has full ability to move data in and out of the device. VFs are lightweight PCIe functions that support data flowing but have a restricted set of configuration resources.

The number of virtual functions provided to the hypervisor or the guest operating system depends on the device. SR-IOV enabled PCIe devices require appropriate BIOS and hardware support, as well as SR-IOV support in the guest operating system driver or hypervisor instance. See [SR-IOV Support](#).

Using SR-IOV in vSphere

In vSphere, a virtual machine can use an SR-IOV virtual function for networking. The virtual machine and the physical adapter exchange data directly without using the VMkernel as an intermediary. Bypassing the VMkernel for networking reduces latency and improves CPU efficiency.

In vSphere, though a virtual switch (standard switch or distributed switch) does not handle the network traffic of an SR-IOV enabled virtual machine connected to the switch, you can control the assigned virtual functions by using switch configuration policies at port group or port level.

SR-IOV Support

vSphere supports SR-IOV in an environment with specific configuration only. Some features of vSphere are not functional when SR-IOV is enabled.

Supported Configurations

To use SR-IOV in vSphere, your environment must meet several configuration requirements.

Table 193: Supported Configurations for Using SR-IOV

Component	Requirements
Physical host	• Must be compatible with the ESXi release. • Must have an Intel or AMD processor. • Must support I/O memory management unit (IOMMU), and must have IOMMU enabled in the BIOS. • Must support SR-IOV, and must have SR-IOV enabled in the BIOS. Contact the server vendor to determine whether the host supports SR-IOV.
Physical NIC	• Must be compatible with the ESXi release. • Must be supported for use with the host and SR-IOV according to the technical documentation from the server vendor. • Must have SR-IOV enabled in the firmware. • Must use MSI-X interrupts.

Component	Requirements
PF driver in ESXi for the physical NIC	• Must be certified by VMware. • Must be installed on the ESXi host. The ESXi release provides a default driver for certain NICs, while for others you must download and manually install it.
Guest OS	Must be supported by the NIC on the installed ESXi release according to the technical documentation from the NIC vendor.
VF driver in the guest OS	• Must be compatible with the NIC. • Must be supported on the guest OS release according to the technical documentation from the NIC vendor. • Must be Microsoft WLK or WHCK certified for Windows virtual machines. • Must be installed on the operating system. The operating system release contains a default driver for certain NICs, while for others you must download and install it from a location provided by the vendor of the NIC or the host.

To verify that your physical hosts and NICs are compatible with ESXi releases, see the *VMware Compatibility Guide*.

Availability of Features

The following features are not available for virtual machines configured with SR-IOV:

- vSphere vMotion
- Storage vMotion
- vShield
- NetFlow
- VXLAN Virtual Wire
- vSphere High Availability
- vSphere Fault Tolerance
- vSphere DRS
- vSphere DPM
- Virtual machine suspend and resume
- Virtual machine snapshots
- MAC-based VLAN for passthrough virtual functions
- Hot addition and removal of virtual devices, memory, and vCPU
- Participation in a cluster environment
- Network statistics for a virtual machine NIC using SR-IOV passthrough

Note: Attempts to enable or configure unsupported features with SR-IOV in the result in unexpected behavior in your environment.

Supported NICs

All NICs must have drivers and firmware that support SR-IOV. Some NICs might require SR-IOV to be enabled on the firmware. To learn what NICs are supported for virtual machines configured with SR-IOV, see the [VMware Compatibility Guide](#).


SR-IOV Component Architecture and Interaction

vSphere SR-IOV support relies on the interaction between the virtual functions (VFs) and the physical function (PF) of the NIC port for better performance, and interaction between the driver of the PF and the host switch for traffic control.

In a host that runs virtual machine traffic on top of SR-IOV physical adapters, virtual machine adapters directly contact the virtual functions to communicate data. However, the ability to configure networks is based on the active policies for the port holding the virtual machines.

On an ESXi host without SR-IOV, the virtual switch sends external network traffic through its ports on the host from or to the physical adapter for the relevant port group. The virtual switch also applies the networking policies on managed packets.

Figure 84: Data and Configuration Paths in the SR-IOV Support of vSphere



Data Path in SR-IOV

After the virtual machine network adapter is assigned to a virtual function, the VF driver in the guest operating system uses the I/O memory management unit (IOMMU) technology to access the virtual function that must receive or send the data over the network. The VMkernel, that is, the virtual switch in particular, does not process the data flow, which reduces the overall latency of SR-IOV enabled workloads.

Configuration Path in SR-IOV

If the guest operating system attempts to change the configuration of a virtual machine adapter mapped to a VF, the change takes place if it is allowed by the policy on the port associated with the virtual machine adapter.

The configuration workflow consists of the following operations:

1. The guest operating system requests a configuration change on the VF.
2. The VF forwards the request to the PF through a mailbox mechanism.
3. The PF driver checks the configuration request with the virtual switch (standard switch or host proxy switch of a distributed switch).
4. The virtual switch verifies the configuration request against the policy on the port with which the VF enabled virtual machine adapter is associated.
5. The PF driver configures the VF if the new settings are in compliance with the port policy of the virtual machine adapter.

For example, if the VF driver tries to modify the MAC address, the address remains the same if MAC address change is not allowed in the security policy for the port group or port. The guest operating system might show that the change is successful but a log message indicates that the operation has failed. As a result, the guest operating system and the virtual device save different MAC addresses. The network interface in the guest operating system might not be able to acquire an IP address and communicate. In this case, you have to reset the interface in the guest operating system to get the latest MAC address from the virtual device and acquire an IP address.

vSphere and Virtual Function Interaction

Virtual functions (VFs) are lightweight PCIe functions that contain all the resources necessary for data exchange but have a minimized set of configuration resources. Interaction between vSphere and VFs is limited.

- The physical NIC must use MSI-X interrupts.
- VFs do not implement rate control in vSphere. Every VF can potentially use the entire bandwidth of a physical link.
- When a VF device is configured as a passthrough device on a virtual machine, the standby and hibernate functions for the virtual machine are not supported.
- The maximum number of VFs that you can create and the maximum number of VFs that you can use for passthrough are different. The maximum number of VFs that you can instantiate depends on the NIC capability and on the hardware configuration of the host. However, due to the limited number of interrupt vectors available for passthrough devices, only a limited number of all instantiated VFs can be used on an ESXi host.

The total number of interrupt vectors on each ESXi host can scale up to 4096 in the case of 32 CPUs. When the host boots, devices on the host such as storage controllers, physical network adapters, and USB controllers consume a subset of the 4096 vectors. If these devices require more than 1024 vectors, the maximum number of potentially supported VFs is reduced.

- The number of VFs that is supported on an Intel NIC might be different from the number that is supported on an Emulex NIC. See the technical documentation from the NIC vendor.
- If you have Intel and Emulex NICs present with SR-IOV enabled, the number of VFs available for the Intel NICs depends on how many VFs are configured for the Emulex NIC, and the reverse. You can use the following formula to estimate the maximum number of VFs for use if all 3072 interrupt vectors are available for passthrough:

$$3X + 2Y < 3072$$

where X is the number of Intel VFs, and Y is the number of Emulex VFs.

This number might be smaller if other types of devices on the host use more than 1024 interrupt vectors from the total of 4096 vectors on the host.

- vSphere SR-IOV supports up to 1024 VFs on supported Intel and Emulex NICs.
- vSphere SR-IOV supports up to 64 VFs on a supported Intel or Emulex NIC.
- If a supported Intel NIC loses connection, all VFs from the physical NIC stop communication completely, including that between VFs.

- If a supported Emulex NIC loses connection, all VFs stop communication with the external environment, but communication between VFs still works
- VF drivers offer many different features, such as IPv6 support, TSO, and LRO checksum. See the technical documentation of the NIC vendor for more details.

DirectPath I/O vs SR-IOV

SR-IOV offers performance benefits and tradeoffs similar to those of DirectPath I/O. DirectPath I/O and SR-IOV have similar functionality but you use them to accomplish different things.

SR-IOV is beneficial in workloads with very high packet rates or very low latency requirements. Like DirectPath I/O, SR-IOV is not compatible with certain core virtualization features, such as vMotion. SR-IOV does, however, allow for a single physical device to be shared amongst multiple guests.

With DirectPath I/O you can map only one physical function to one virtual machine. SR-IOV lets you share a single physical device, allowing multiple virtual machines to connect directly to the physical function.

Configure a Virtual Machine to Use SR-IOV

To use the capabilities of SR-IOV, you must enable the SR-IOV virtual functions on the host and connect a virtual machine to the functions.

Verify that the configuration of your environment supports SR-IOV. See [SR-IOV Support](#).

The traffic passes from an SR-IOV passthrough adapter to the physical adapter in compliance with the active policy on the associated port on the standard or distributed switch.

To examine which virtual function is assigned to an SR-IOV passthrough network adapter, on the **Summary** tab for the virtual machine expand the **VM Hardware** panel and check the properties of the adapter.

The topology diagram of the switch marks virtual machine adapters that use virtual functions with the  icon.

Set up the traffic passing through the virtual functions attached to the virtual machine by using the networking policies on the switch, port group, and port. See [Networking Options for the Traffic Related to an SR-IOV Enabled Virtual Machine](#).

Enable SR-IOV on a Host Physical Adapter

Before you can connect virtual machines to virtual functions, use the vSphere Client to enable SR-IOV and set the number of virtual functions on your host.

1. In the **Hosts and Clusters** view, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Physical adapters**.
You can look at the SR-IOV property to see whether a physical adapter supports SR-IOV.
3. Select the physical adapter and click **Edit adapter settings**.
4. Under SR-IOV, select **Enabled** from the **Status** drop-down menu.
5. In the **Number of virtual functions** text box, type the number of virtual functions that you want to configure for the adapter.
A value of 0 means SR-IOV is not enabled for that physical function.
6. Click **OK**.
7. Restart the host.

The virtual functions become active on the NIC port represented by the physical adapter entry. They appear in the PCI Devices list in the **Settings** tab for the host.

You can use the `esxcli network sriovnic vCLI` commands to examine the configuration of virtual functions on the host. Associate a virtual machine with a virtual function through an SR-IOV passthrough network adapter.

Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine

To ensure that a virtual machine and a physical NIC can exchange data, you must associate the virtual machine with one or more virtual functions as SR-IOV passthrough network adapters.

- Verify that the virtual functions exist on the host.
 - Verify that the passthrough networking devices for the virtual functions are active in the PCI Devices list on the **Settings** tab for the host.
 - Verify that the virtual machine compatibility is ESXi 5.5 and later.
 - Verify that Red Hat Enterprise Linux 6 and later or Windows has been selected as the guest operating system when the virtual machine was created.
1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
 2. Power off the virtual machine.
 3. From the Actions menu, select Edit Settings.
 4. Select the **Virtual Hardware** tab in the dialog box displaying the settings.
 5. From the **Add new device** drop-down menu, select **Network Adapter**.
The New Network section is added to the list in the **Virtual Hardware** tab.
 6. Expand the New Network section and connect the virtual machine to a port group.
The virtual NIC does not use this port group for data traffic. The port group is used to extract the networking properties, for example VLAN tagging, to apply on the data traffic.
 7. From the **Adapter type** drop-down menu, select **PCI Device passthrough**.
 8. From the **Physical function** drop-down menu, select the physical adapter to back the passthrough virtual machine adapter.
 9. To allow changes in the MTU of packets from the guest operating system, use the **Guest OS MTU Change** drop-down menu.
 10. Expand the Memory section, select **Reserve all guest memory (All locked)** and click **OK**.
I/O memory management unit (IOMMU) must reach all virtual machine memory so that the passthrough device can access the memory by using direct memory access (DMA).
 11. Power on the virtual machine.

When you power on the virtual machine, the ESXi host selects a free virtual function from the physical adapter and maps it to the SR-IOV passthrough adapter. The host validates all properties of the virtual machine adapter and the underlying virtual function against the settings of the port group to which the virtual machine belongs.

Networking Options for the Traffic Related to an SR-IOV Enabled Virtual Machine

In vSphere you can configure certain networking features on a virtual machine adapter that is associated virtual function (VF). Use settings for the switch, for the port group, or for a port depending on the type of the virtual switch (standard or distributed) that handles the traffic.

Table 194: Networking Options for a Virtual Machine Adapter That Uses a VF

Networking Option	Description
MTU size	Change the size of the MTU, for example, to enable jumbo frames.
Security policy for VF traffic	- If the guest operating system changes the initially set MAC address of a virtual machine network adapter that uses a VF, accept or drop incoming frames for the new address by setting the MAC address changes option. - Enable global promiscuous mode for virtual machine network adapters, including adapters that use VFs.
VLAN tagging mode	Configure VLAN tagging in the standard or distributed switch, and let the tagged traffic reach the virtual machines that are associated with VFs, that is, enable Virtual Guest Tagging (VGT).

Using an SR-IOV Physical Adapter to Handle Virtual Machine Traffic

In vSphere both the physical function (PF) and virtual functions (VFs) of an SR-IOV capable physical adapter can be configured to handle virtual machine traffic.

The PF of an SR-IOV physical adapter controls the VFs that virtual machines use, and can carry the traffic flowing through the standard or distributed switch that handles the networking of these SR-IOV enabled virtual machines.

The SR-IOV physical adapter works in different modes depending on whether it backs the traffic of the switch.

Mixed Mode

The physical adapter provides virtual functions to virtual machines attached to the switch and directly handles traffic from non SR-IOV virtual machines on the switch.

You can check whether an SR-IOV physical adapter is in mixed mode in the topology diagram of the switch. An SR-IOV physical adapter in mixed mode appears with the  icon in the list of physical adapters for a standard switch or in the list of uplink group adapters for a distributed switch.

SR-IOV Only Mode

The physical adapter provides virtual functions to virtual machines connected to a virtual switch, but does not back traffic from non SR-IOV virtual machines on the switch.

To verify whether the physical adapter is in SR-IOV only mode, examine the topology diagram of the switch. In this mode, the physical adapter is in a separate list called External SR-IOV Adapters and appears with the  icon.

Non SR-IOV Mode

The physical adapter is not used for traffic related to VF aware virtual machines. It handles traffic from non SR-IOV virtual machines only.

Enable SR-IOV by Using Host Profiles or an ESXCLI Command

You can configure the virtual functions on an ESXi host by using an ESXCLI command, or by using a host profile to set up multiple hosts simultaneously or to set up stateless hosts.

Enable SR-IOV in a Host Profile

For multiple hosts or a stateless host, you can configure the virtual functions of the physical NIC by using a host profile and apply the profile on a host by using Auto Deploy.

- Verify that the configuration of your environment supports SR-IOV. See [SR-IOV Support](#).
- Create a host profile based on the SR-IOV capable host. See the *vSphere Host Profiles* documentation.

For information about running ESXi by using Auto Deploy with host profiles, see the *vCenter Server Installation and Setup* documentation.

You can also enable SR-IOV virtual functions on the host by using the `esxcli system module parameters set vCLI` command on the NIC driver parameter for virtual functions in accordance with the driver documentation. For more information about using ESXCLI commands, see *ESXCLI Concepts and Examples* documentation.

1. From the Home page, click **Host Profiles**.
 2. Select the host profile from the list and click the **Configure** tab.
 3. Click **Edit Host Profile** and expand the **General System Settings** node.
 4. Expand **Kernel Module Parameter** and select the parameter of the physical function driver for creating virtual functions.
For example, the parameter for the physical function driver of an Intel physical NIC is `max_vfs`.
 5. In the **Value** text box, type a comma-separated list of valid virtual function numbers.
Each list entry indicates the number of virtual functions that you want to configure for each physical function. A value of 0 ensures that SR-IOV is not enabled for that physical function.

For example, if you have a dual port, set the value to `x,y` where `x` or `y` is the number of virtual functions you want to enable for a single port.

If the target number of virtual functions on a single host is 30, you might have two dual port cards set to `0,10,10,10`.
- Note:** The number of virtual functions supported and available for configuration depends on your system configuration.
6. Click **Finish**.
 7. Remediate the host profile to the host as required.

The virtual functions appear in the PCI Devices list on the **Settings** tab for the host.

Associate a virtual function with a virtual machine adapter by using the SR-IOV passthrough network adapter type. See [Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine](#).

Enable SR-IOV on a Host Physical Adapter by Using an ESXCLI Command

Learn how to run a console command on ESXi to create SR-IOV virtual functions on a physical adapter to troubleshoot or to configure hosts directly.

Install the vCLI package, deploy the vSphere Management Assistant (vMA) virtual machine, or use the ESXi Shell. See *Getting Started with ESXCLI*.

You can create SR-IOV virtual functions on the host by manipulating the NIC driver parameter for virtual functions in accordance with the driver documentation.

1. To create virtual functions by setting the parameter for virtual functions of the NIC driver, run the `esxcli system module parameters set` command at the command prompt.

```
esxcli system module parameters set -m driver -p vf_param=w,x,y,z
```


Where *driver* is the name of the NIC driver, and *vf_param* is the driver-specific parameter for creating the virtual function.

You can use a comma-separated list to set values for the *vf_param* parameter, where each entry indicates the number of virtual functions for a port. A value of 0 ensures that SR-IOV is not enabled for that physical function.

If you have two dual port NICs, you can set the value to *w, x, y, z*, where *w, x, y*, and *z* is the number of virtual functions you want to enable for a single port. For example, to create 30 virtual functions distributed on two dual port Intel cards by using the ixgbe driver, run the following command for the ixgbe driver and the *max_vfs* parameter:

```
esxcli system module parameters set -m ixgbe -p max_vfs=0,10,10,10
```

2. Restart the host to create the virtual functions.

Associate a virtual function with a virtual machine adapter by using the SR-IOV passthrough network adapter type. See [Assign a Virtual Function as SR-IOV Passthrough Adapter to a Virtual Machine](#).

What is Remote Direct Memory Access for Virtual Machines

vSphere 6.5 and later releases support remote direct memory access (RDMA) communication between virtual machines that have paravirtualized RDMA (PVRDMA) network adapters .

Overview of RDMA

RDMA allows direct memory access from the memory of one computer to the memory of another computer without involving the operating system or CPU. The transfer of memory is offloaded to the RDMA-capable Host Channel Adapters (HCA) . A PVRDMA network adapter provides remote direct memory access in a virtual environment.

Using RDMA in vSphere

In vSphere, a virtual machine can use a PVRDMA network adapter to communicate with other virtual machines that have PVRDMA devices. The virtual machines must be connected to the same vSphere Distributed Switch.

The PVRDMA device automatically selects the method of communication between the virtual machines . For virtual machines that run on the same ESXi host with or without a physical RDMA device, the data transfer is a memcopy between the two virtual machines. The physical RDMA hardware is not used in this case.

For virtual machines that reside on different ESXi hosts and that have a physical RDMA connection, the physical RDMA devices must be uplinks on the distributed switch. In this case, the communication between the virtual machines by way of PVRDMA uses the underlying physical RDMA devices.

For two virtual machines that run on different ESXi hosts, when at least one of the hosts does not have a physical RDMA device, the communication falls back to a TCP-based channel and the performance is reduced.

PVRDMA Support

vSphere 6.5 and later supports PVRDMA only in environments with a specific configuration.

Supported Configurations

To use PVRDMA in vSphere 6.5 or later, your environment must meet several configuration requirements.

Table 195: Supported Configurations for Using PVRDMA

Component	Requirements
vSphere	- ESXi host 6.5 or later. - vCenter Server 6.5 or later.

Component	Requirements
	vSphere Distributed Switch.
Physical host	Must be compatible with the ESXi release.
Host Channel Adapter (HCA)	Must be compatible with the ESXi release. Note: Virtual machines that reside on different ESXi hosts require HCA to use RDMA. You must assign the HCA as an uplink for the vSphere Distributed Switch. PVRDMA does not support NIC teaming. The HCA must be the only uplink on the vSphere Distributed Switch. For virtual machines on the same ESXi hosts or virtual machines using the TCP-based fallback, the HCA is not required.
Virtual machine	Virtual hardware version 13 or later.
Guest OS	Linux (64-bit)

To verify that your physical hosts and HCAs are compatible with ESXi releases, see the *VMware Compatibility Guide*.

Note: Attempts to enable or configure unsupported features with PVRDMA might result in unexpected behavior in your environment.

Support for PVRDMA Namespaces

In releases previous to vSphere 7.0, PVRDMA virtualized public resource identifiers in the underlying hardware to guarantee that a physical resource can be allocated with the same public identifier when a virtual machine resumed operation following the use of vMotion to move it from one physical host server to another. To do this, PVRDMA distributed virtual to physical resource identifier translations to peers when creating a resource. This resulted in additional overhead that can be significant when creating large numbers of resources.

PVRDMA namespaces prevents these additional overheads by letting multiple virtual machines coexist without coordinating the assignment of identifiers. Each virtual machine is assigned an isolated identifier namespace on the RDMA hardware, such that any virtual machine can select its identifiers within the same range without conflicting with other virtual machines. The physical resource identifier no longer changes even after vMotion, so virtual to physical resource identifier translations are no longer necessary.

PVRDMA namespaces are enabled automatically on vSphere 7.0 and later with virtual machine hardware version 17 or later. The underlying hardware must also support PVRDMA namespaces. To learn how to enable PVRDMA namespaces on your environment's hardware, refer to the RDMA vendor documentation.

Support for PVRDMA Native Endpoints

PVRDMA native endpoints are supported in virtual machine hardware version 18 and later beginning with vSphere 7.0 Update 1 and later releases. PVRDMA native endpoints allow PVRDMA to communicate with non-PVRDMA endpoints. To use PVRDMA native endpoints, you must enable PVRDMA namespaces. To learn how to enable PVRDMA namespaces on your environment's specific hardware, refer to the RDMA vendor documentation.

You must configure the virtual machine to use PVRDMA native endpoints. See [Configure a Virtual Machine to Use PVRDMA Native Endpoints](#).

Configure an ESXi Host for PVRDMA

Configure the VMkernel adapter and firewall rule of an ESXi host for PVRDMA communication.

Verify that your ESXi host meets the requirements for PVRDMA. See [PVRDMA Support](#).

Tag a VMkernel Adapter for PVRDMA

Learn how to select a VMkernel adapter and enable it for PVRDMA communication.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Click the **Edit** button.
5. Locate `Net.PVRDMAVmknicsby` using the filter text field.
6. Click the value field and enter the value of the VMkernel adapter that you want to use, for example `vmk0`.
7. Click **OK**.

Enable the Firewall Rule for PVRDMA

Learn how to enable the firewall rule for PVRDMA in the security profile of the ESXi host.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Firewall**.
4. Click the **Edit** button.
5. Locate the `pvrDMA` rule by using the filter text field.
6. Select the check box next to the `pvrDMA` rule and click **OK**.

Assign a PVRDMA Adapter to a Virtual Machine

To enable a virtual machine to exchange data by using RDMA, you must associate the virtual machine with a PVRDMA network adapter.

- Verify that the host on which the virtual machine is running is configured for RDMA. See [Configure an ESXi Host for PVRDMA](#).
- Verify that the host is connected to a vSphere Distributed Switch.
- Verify that the virtual machine uses virtual hardware version 13 and above.
- Verify that the guest operating system is a Linux 64-bit distribution.

You can add up to ten PVRDMA network adapters to a virtual machine when using vSphere 7.0.2 and later.

1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
2. Power off the virtual machine.
3. From the Actions menu, select Edit Settings.
4. Select the **Virtual Hardware** tab in the dialog box displaying the settings.
5. From the **Add new device** drop-down menu, select **Network Adapter**.
The New Network section is added to the list in the **Virtual Hardware** tab.

6. Expand the **New Network** section and connect the virtual machine to a distributed port group.
7. From the **Adapter type** drop-down menu, select PVRDMA.
8. Expand the **Memory** section, select **Reserve all guest memory (All locked)**, and click **OK**.
9. Power on the virtual machine.

Configure a Virtual Machine to Use PVRDMA Native Endpoints

PVRDMA native endpoints are available as an advanced virtual machine configuration.

Verify that your environment supports PVRDMA. See [PVRDMA Support](#).

PVRDMA native endpoints are supported in virtual machine hardware version 18 and later beginning with vSphere 7.0 Update 1 and later releases. To use PVRDMA native endpoints, you must enable PVRDMA namespaces. To learn how to enable PVRDMA namespaces on your environment's specific hardware, refer to the vendor documentation.

You can use the vSphere Client to configure native endpoints or edit the virtual machine's VMX file. If editing the VMX file directly, add the parameter `vrddmax.nativeEndpointSupport = "TRUE"`, where *x* is the index of the PVRDMA adapter. The following procedure uses the vSphere Client to configure native endpoints.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click **VM Options**.
4. Expand **Advanced**.
5. Under Configuration Parameters, click the **Edit Configuration** button.
6. In the dialog box that appears, click **Add Row** to enter a new parameter and its value.
7. Enter the parameter `vrddmax.nativeEndpointSupport`, where *x* is the index of the PVRDMA adapter, and set the value to `TRUE`.

The index *x* is the number of the PVRDMA adapter minus 1. For example, if the PVRDMA adapter you want to enable native endpoints with is labelled "Network Adapter 2," then the index is 1.

Configure a Virtual Machine to use PVRDMA Asynchronous Mode

Learn how to configure a virtual machine to use PVRDMA asynchronous mode. It is available as an advanced virtual machine configuration.

Verify that your environment supports PVRDMA. See [PVRDMA Support](#).

PVRDMA asynchronous mode is available in virtual machines running on vSphere 8.0 and later. Asynchronous mode might improve the throughput and latency for RDMA workloads running in the virtual machine. When asynchronous mode

is enabled, an increased CPU use might be observed in the host. When asynchronous mode is in use, it is recommended that the virtual machine is configured for high latency sensitivity.

1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click **VM Options**.
4. Expand **Advanced**.
5. Under Configuration Parameters, click the **Edit Configuration** button.
6. In the dialog box that appears, click **Add Row** to enter a new parameter and its value.
7. Enter the parameter `vrDMA.asyncMode`, and set the value to `TRUE`.

Network Requirements for RDMA over Converged Ethernet

RDMA over Converged Ethernet ensures low-latency, light-weight, and high-throughput RDMA communication over an Ethernet network. RoCE requires a network that is configured for lossless traffic of information at layer 2 alone or at both layer 2 and layer 3.

RDMA over Converged Ethernet (RoCE) is a network protocol that uses RDMA to provide faster data transfer for network-intensive applications. RoCE allows direct memory transfer between hosts without involving the hosts' CPUs.

There are two versions of the RoCE protocol. RoCE v1 operates at the link network layer (layer 2). RoCE v2 operates at the Internet network layer (layer 3). Both RoCE v1 and RoCE v2 require a lossless network configuration. RoCE v1 requires a lossless layer 2 network, and RoCE v2 requires that both layer 2 and layer 3 are configured for lossless operation.

Lossless Layer 2 Network

To ensure lossless layer 2 environment, you must be able to control the traffic flows. Flow control is achieved by enabling global pause across the network or by using the Priority Flow Control (PFC) protocol defined by Data Center Bridging group (DCB). PFC is a layer 2 protocol that uses the class of services field of the 802.1Q VLAN tag to set individual traffic priorities. It puts on pause the transfer of packets towards a receiver in accordance with the individual class of service priorities. This way, a single link carries both lossless RoCE traffic and other lossy, best-effort traffic. With traffic flow congestion, important lossy traffic can be affected. To isolate different flows from one another, use RoCE in a PFC priority-enabled VLAN.

Lossless Layer 3 Network

RoCE v2 requires that lossless data transfer is preserved at layer 3 routing devices. To enable the transfer of layer 2 PFC lossless priorities across layer 3 routers, configure the router to map the received priority setting of a packet to the corresponding Differentiated Serviced Code Point (DSCP) QoS setting that operates at layer 3. The transferred RDMA packets are marked with layer 3 DSCP, layer 2 Priority Code Points (PCP) or with both. To extract priority information from the packet routers use either DSCP or PCP. In case PCP is used, the packet must be VLAN-tagged and the router must copy the PCP bits of the tag and forward them to the next network. If the packet is marked with DSCP, the router must keep the DSCP bits unchanged.

Like RoCE v1, RoCE v2 must run on a PFC priority-enabled VLAN.

Note: Do not team RoCE NICs, if you intend to use RDMA on those NICs.

For vendor-specific configuration information, refer to the official documentation of the respective device or the switch vendor.

Configure Remote Direct Memory Access Network Adapters

Learn how to install a remote direct memory access (RDMA) network adapter on your ESXi hosts. Once installed, you can use the vSphere Client to view the RDMA adapter and its corresponding network adapter, and configure its VMkernel binding.

Install an RDMA-capable adapter on your ESXi host. For example, Mellanox Technologies MT27700 Family ConnectX-4.

RDMA provides direct memory access from the memory of one host to the memory of another host without involving the remote operating system and CPU. This boosts network and host performance with lower latency, lower CPU load, and faster bandwidth.

View RDMA Capable Network Adapter

ESXi supports RDMA capable network adapters. After you install such an adapter on your ESXi host, the vSphere Client displays its two components, the RDMA adapter, and a physical network adapter.


Install a RDMA-capable adapter that supports RDMA (RoCE v2) on your ESXi host. For example, Mellanox Technologies MT27700 Family ConnectX-4.

You can use the vSphere Client to view the RDMA adapter and its corresponding network adapter.

1. On your ESXi host, install a RDMA-capable adapter that supports RDMA (RoCE v2).
The host discovers the adapter and the vSphere Client displays its two components, an RDMA adapter, and a physical network adapter.
2. Navigate to the host.
3. Under **Networking**, click **RDMA adapters**.

In this example, the RDMA adapter appears on the list as `vmrdma0`. The **Paired Uplink** column displays the network component as the `vmnic1` physical network adapter.

Figure 85: RDMA adapters installed on an ESXi host in a vSphere environment.




4. To verify the description of the adapter, select the RDMA adapter from the list, and click the **Properties** tab.

Configure Remote Direct Memory Access Network Adapters

You can install a remote direct memory access (RDMA) network adapter and configure its VMkernel binding.

1. Install an RDMA-capable adapter that supports RDMA (RoCE v2) on your ESXi host.
The host discovers the adapter and the vSphere Client displays its two components, an RDMA adapter and a physical network adapter.
2. In the vSphere Client, verify that the RDMA adapter is discovered by your host.
 - a) Navigate to the host.
 - b) Click the **Configure** tab.
 - c) Under **Networking**, click **RDMA adapters**.

In this example, the RDMA adapter appears on the list as `vmrdma0`. The **Paired Uplink** column displays the network component as the `vmnic1` physical network adapter.



- d) To verify the description of the adapter, select the RDMA adapter from the list, and click the **Properties** tab.
3. Configure VMkernel binding for the RDMA adapter.
In the configuration, you can use a vSphere standard switch or a vSphere Distributed Switch. The following steps use the standard switch as an example.
 - a) Create a vSphere standard switch and add the network component to the switch.

Note: Make sure to select the physical network adapter that corresponds to the RDMA adapter. In this example, it is the `vmnic1` adapter.


For information about creating the switch, see [Create a vSphere Standard Switch](#) or [Create a vSphere Distributed Switch](#).

- b) Add a VMkernel adapter to the vSphere standard switch that you created.

Assign an appropriate static IPv4 or IPv6 address to the VMkernel adapter, so that your RDMA adapter can discover the NVMe over RDMA target.

For information about adding the VMkernel adapter, see [How do I Set Up VMkernel Networking](#).

The illustration shows that the physical network adapter and the VMkernel adapter are connected to the vSphere standard switch. Through this connection, the RDMA adapter is bound to the VMkernel adapter.




4. Verify the VMkernel binding configuration for the RDMA adapter.

- a) Navigate to the RDMA adapter.

- b) Click the **VMkernel adapters binding** tab and verify that the associated VMkernel adapter appears on the page.

In this example, the `vmrdma0` RDMA adapter is paired to the `vmnic1` network adapter and is connected to the `vmk1` VMkernel adapter.



You can use the RDMA network component of the adapter for such storage configurations as iSER or NVMe over RDMA. To learn more, see the *vSphere Storage* documentation.

What are Jumbo Frames

Jumbo frames let ESXi hosts send larger frames out onto the physical network. The network must support jumbo frames end-to-end that includes physical network adapters, physical switches, and storage devices.

Before enabling jumbo frames, check with your hardware vendor to ensure that your physical network adapter supports jumbo frames.

You can enable jumbo frames on a vSphere distributed switch or vSphere standard switch by changing the maximum transmission unit (MTU) to a value greater than 1280 bytes. With vCenter Server 7.0 Update 3, you can set the size of the maximum transmission unit (MTU) on a vSphere Distributed Switch to up to 9190 bytes to support switches with larger packet sizes.

Enable Jumbo Frames on a vSphere Distributed Switch

Enable jumbo frames for the entire traffic that passes through a vSphere Distributed Switch.

Important: When you change the MTU size of a vSphere Distributed Switch, the physical NICs that are assigned as uplinks are brought down and up again. This causes a short network outage of between 5 to 10 milliseconds for virtual machines or services that are using the uplinks.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Settings** and select **Properties**.
3. Click **Edit**.
4. Click **Advanced** and set the **MTU** property to a value greater than 1500 bytes.
You cannot set the MTU size to a value greater than 9190 bytes.
5. Click **OK**.

Enable Jumbo Frames on a vSphere Standard Switch

Enable jumbo frames for all traffic through a vSphere Standard Switch on a host.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select a standard switch from the virtual switch table and click **Edit settings**.
4. In the **Properties** section, set the **MTU** property to a value greater than 1500 bytes.
You can increase the MTU size up to 9000 bytes.
5. Click **OK**.

Enable Jumbo Frames for a VMkernel Adapter

Jumbo frames reduce the CPU load caused by transferring data. Enable jumbo frames on a VMkernel adapter by changing the maximum transmission units (MTU) of the adapter.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **VMkernel adapters**.
3. Select a VMkernel adapter from the adapter table.
The properties of the adapter appear.
4. Click **Edit**.
5. On the Port properties page, set the **MTU** property to a value greater than 1500.
You can increase the MTU size up to 9000 bytes.

6. Click **OK**.

Enable Jumbo Frame Support on a Virtual Machine

Enabling jumbo frame support on a virtual machine requires an enhanced VMXNET adapter for that virtual machine.

1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
2. From the Actions menu, select Edit Settings.
3. Select the **Virtual Hardware** tab in the dialog box displaying the settings.
4. Expand the **Network adapter** section. Record the network settings and MAC address that the network adapter is using.
5. Click the times-circle button to remove the network adapter from the virtual machine.
6. From the **Add new device** drop-down menu, select **Network Adapter**.
The New Network section is added to the list in the Virtual Hardware tab.
7. Expand the New Network section.
8. From the **Adapter Type** drop-down menu, select **VMXNET 2 (Enhanced)** or **VMXNET 3**.
9. Set the network settings to the ones recorded for the old network adapter.
10. Set the **MAC Address** to **Manual**, and type the MAC address that the old network adapter was using.
11. Click **OK**.
 - Check that the enhanced VMXNET adapter is connected to a standard switch or to a distributed switch with jumbo frames enabled.
 - Inside the guest operating system, configure the network adapter to allow jumbo frames. See the documentation of your guest operating system.
 - Configure all physical switches and any physical or virtual machines to which this virtual machine connects to support jumbo frames.

What is TCP Segmentation Offload

Learn how to use TCP Segmentation Offload (TSO) in VMkernel network adapters and virtual machines to improve the network performance in workloads that have severe latency requirements.

TSO on the transmission path of physical network adapters, and VMkernel and virtual machine network adapters improves the performance of ESXi hosts by reducing the overhead of the CPU for TCP/IP network operations. When TSO is enabled, the network adapter divides larger data chunks into TCP segments instead of the CPU. The VMkernel and the guest operating system can use more CPU cycles to run applications.

To benefit from the performance improvement that TSO provides, enable TSO along the data path on an ESXi host including physical network adapters, VMkernel and guest operating system. By default, TSO is enabled in the VMkernel of the ESXi host , and in the VMXNET 2 and VMXNET 3 virtual machine adapters.

For information about the location of TCP packet segmentation in the data path, see VMware Knowledge Base article [Understanding TCP Segmentation Offload \(TSO\) and Large Receive Offload \(LRO\) in a VMware environment](#).

Manage Software TSO in the VMkernel

If a physical network adapter experiences problems with TSO, you can temporarily enable the software simulation of TSO in the VMkernel until you troubleshoot the problems.

Run these `esxcli network nic software set` console commands to activate or deactivate the software simulation of TSO in the VMkernel.

- Activate the software simulation of TSO in the VMkernel.

```
esxcli network nic software set --ipv4tso=1 -n vmnicX
esxcli network nic software set --ipv6tso=1 -n vmnicX
```

- Deactivate the software simulation of TSO in the VMkernel.

```
esxcli network nic software set --ipv4tso=0 -n vmnicX
esxcli network nic software set --ipv6tso=0 -n vmnicX
```

where *X* in `vmnicX` represents the number of the NIC ports on the host.

The configuration change persists across host reboots.

How to determine whether TSO is Supported on the Physical Network Adapters on an ESXi Host

Examine whether a physical network adapter offloads TCP/IP packet segmentation when you estimate the networking performance on a host that runs latency-sensitive workloads. If a physical network adapter supports TSO, then TSO is enabled by default.

Run the following console command to determine whether TSO is enabled on the physical network adapters on a host.

```
esxcli network nic tso get
```

Manage TSO on an ESXi Host

Activate TCP Segmentation Offload (TSO) on the transmission path to have the NIC divide larger data chunks into TCP segments. Deactivate TSO to have CPU perform TCP segmentation.

By default, a host uses hardware TSO if its physical adapters support it.

1. In the `ESXi Shell`, navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Edit the value of the `Net.UseHwTSO` parameter for IPv4 and of `Net.UseHwTSO6` for IPv6.
 - To activate TSO, set `Net.UseHwTSO` and `Net.UseHwTSO6` to 1.
 - To deactivate TSO, set `Net.UseHwTSO` and `Net.UseHwTSO6` to 0.
5. Click **OK** to apply the changes.
6. To reload the driver module of the physical adapter, run the `esxcli system module set` console command in the ESXi Shell on the host.
 - a) To deactivate the driver, run the `esxcli system module set` command with the `--enabled false` option.

```
esxcli system module set
--enabled false
--module
nic_driver_module
```

- b) To activate the driver, run the `esxcli system module set` command with the `--enabled true` option.

```
esxcli system module set
--enabled true
--module
nic_driver_module
```

If a physical adapter does not support hardware TSO, the VMkernel segments large TCP packets coming from the guest operating system and sends them to the adapter.

How to find if TSO Is Enabled on an ESXi Host

Examine whether hardware TSO is enabled in the VMkernel when you estimate the networking performance on a host that runs latency-sensitive workloads. By default, hardware TSO is enabled on an ESXi host.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Examine the value of the `Net.UseHwTSO` and `Net.UseHwTSO6` parameters.

`Net.UseHwTSO` shows the TSO state for IPv4, and `Net.UseHwTSO6` for IPv6. TSO is enabled if the property is set to 1.

Manage TSO on a Linux Virtual Machine

Activate TSO support on the network adapter of a Linux virtual machine so that the guest operating system redirects TCP packets that need segmentation to the VMkernel.

- Verify that ESXi supports the Linux guest operating system.
See the *VMware Compatibility Guide* documentation.
- Verify that the network adapter on the Linux virtual machine is VMXNET2 or VMXNET3.

In a terminal window on the Linux guest operating system, to activate or deactivate TSO, run the `ethtool` command with the `-K` and `tso` options.

- To activate TSO, run the following command:

```
ethtool-K ethYtsoon
```

- To deactivate TSO, run the following command:

```
ethtool-K ethYtsooff
```

where `Y` in `ethY` is the sequence number of the NIC in the virtual machine.

Manage TSO on a Windows Virtual Machine

By default, TSO is activated on a Windows virtual machine on VMXNET2 and VMXNET3 network adapters. For performance reasons, you might want to deactivate TSO.

- Verify that ESXi supports the Windows guest operating system. See the *VMware Compatibility Guide* documentation.
- Verify that the network adapter on the Windows virtual machine is VMXNET2 or VMXNET3.

1. In the Network and Sharing Center on the Windows control panel, click the name of the network adapter.
2. Click its name.
A dialog box displays the status of the adapter.

3. Click **Properties**, and beneath the network adapter type, click **Configure**.
4. On the **Advanced** tab, set the **Large Send Offload V2 (IPv4)** and **Large Send Offload V2 (IPv6)** properties to **Enabled** or **Disabled**.
5. Click **OK**.
6. Restart the virtual machine.

What is Large Receive Offload

Learn how to use Large Receive Offload (LRO) to reduce the CPU overhead for processing packets that arrive from the network at a high rate.

LRO reassembles incoming network packets into larger buffers and transfers the resulting larger but fewer packets to the network stack of the host or virtual machine. The CPU has to process fewer packets than when LRO is deactivated, which reduces its utilization for networking especially in the case of connections that have high bandwidth.

To benefit from the performance improvement of LRO, enable LRO along the data path on an ESXi host including VMkernel and guest operating system. By default, LRO is activated in the VMkernel and in the VMXNET3 virtual machine adapters.

For information about the location of TCP packet aggregation in the data path, see VMware Knowledge Base article [Understanding TCP Segmentation Offload \(TSO\) and Large Receive Offload \(LRO\) in a VMware environment](#).

Manage Hardware LRO for All VMXNET3 Adapters on an ESXi Host

Activate the hardware capabilities of host physical adapters to aggregate incoming TCP packets for VMXNET3 VM adapters by using the LRO technology instead of consuming resources for assembling in the guest operating system.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Edit the value of the `Net.Vmxnet3HwLRO` parameter.
 - To activate hardware LRO, set `Net.Vmxnet3HwLRO` to 1.
 - To deactivate hardware LRO, set `Net.Vmxnet3HwLRO` to 0.
5. Click **OK** to apply the changes.

Manage Software LRO for All VMXNET3 Adapters on an ESXi Host

Use software LRO in the VMkernel backend of VMXNET3 adapters to improve networking performance of virtual machines if the host physical adapters do not support hardware LRO.

vSphere supports software LRO for both IPv4 and IPv6 packets.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Edit the value of the `Net.Vmxnet3SwLRO` parameter for VMXNET3 adapters.
 - To activate software LRO, set `Net.Vmxnet3SwLRO` to 1.
 - To deactivate software LRO, set `Net.Vmxnet3SwLRO` to 0.

5. Click **OK** to apply the changes.

Determine Whether LRO Is Enabled for VMXNET3 Adapters on an ESXi Host

Examine the status of LRO on an ESXi when you estimate the networking performance on a host that runs latency-sensitive workloads.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Examine the value of the LRO parameters for VMXNET2 and VMXNET3.
 - For hardware LRO, examine the `Net.Vmxnet3HwLRO` parameter. If it is equal to 1, hardware LRO is enabled.
 - For software LRO, examine the `Net.Vmxnet3SwLRO` parameter. If it is equal to 1, hardware LRO is enabled.

Change the Size of the LRO Buffer for VMXNET 3 Adapters

You can change the size of the buffer for packet aggregation for virtual machine connections through VMXNET 3 network adapters. Increase the buffer size to reduce the number of TCP acknowledgments and improve efficiency in workloads.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Enter a value between 1 and 65535 for the `Net.VmxnetLROMaxLength` parameter to set the LRO buffer size in bytes.

By default the size of the LRO buffer is equal to 32000 bytes.

Activate or Deactivate LRO for All VMkernel Adapters on an ESXi Host

Use LRO in the VMkernel network adapters on an ESXi host to improve the networking performance for incoming infrastructure traffic.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Edit the value of the `Net.TcpipDefLROEnabled` parameter.
 - To activate LRO for the VMkernel network adapters on the host, set `Net.TcpipDefLROEnabled` to 1.
 - To deactivate software LRO for the VMkernel network adapters on the host, set `Net.TcpipDefLROEnabled` to 0.

5. Click **OK** to apply the changes.

Change the Size of the LRO Buffer for VMkernel Adapters

You can change the size of the buffer for packet aggregation for VMkernel connections. Increase the buffer size to reduce the number of TCP acknowledgments and improve efficiency in the VMkernel.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System**.
3. Click **Advanced System Settings**.
4. Enter a value between 1 and 65535 for the `Net.TcpipDefLROMaxLength` parameter to set the LRO buffer size in bytes.

By default the size of the LRO buffer is equal to 32768 bytes.

Manage LRO on a VMXNET3 Adapter on a Linux Virtual Machine

If LRO is enabled for VMXNET3 adapters on the host, activate LRO support on a network adapter on a Linux virtual machine to ensure that the guest operating system does not spend resources to aggregate incoming packets into larger buffers.

Verify that the Linux kernel is 2.6.24 and later.

In a terminal window on the Linux guest operating system, run the `ethtool` command with the `-K` and `lro` options.

- To activate LRO, run the following command:

```
ethtool -K ethY lro on
```

where `Y` in `ethY` is the sequence number of the NIC in the virtual machine.

- To deactivate LRO, run the following command:

```
ethtool -K ethY lro off
```

where `Y` in `ethY` is the sequence number of the NIC in the virtual machine.

Manage LRO on a VMXNET3 Adapter on a Windows Virtual Machine

If LRO is enabled for VMXNET3 adapters on the host, activate LRO support on a network adapter on a Windows virtual machine to ensure that the guest operating system does not spend resources to aggregate incoming packets into larger buffers.

- Verify that the virtual machine runs Windows Server 2012 and later or Windows 8 and later.
- Verify that the virtual machine compatibility is ESXi 6.0 and later.
- Verify that the version of the VMXNET3 driver installed on the guest operating system is 1.6.6.0 and later.
- Verify that LRO is enabled globally on a virtual machine that runs Windows Server 2012 and later or Windows 8 and later. See [Manage LRO Globally on a Windows Virtual Machine](#).

On Windows, the LRO technology is also referred to as Receive Side Coalescing (RSC).

1. In the **Network and Sharing Center** of the guest operating system's Control Panel, click the name of the network adapter.

A dialog box displays the status of the adapter.

2. Click **Properties**, and under the VMXNET3 network adapter type, click **Configure**.
3. On the **Advanced** tab, set both **Recv Segment Coalescing (IPv4)** and **Recv Segment Coalescing (IPv6)** to **Enabled** or **Disabled**.
4. Click **OK**.

Manage LRO Globally on a Windows Virtual Machine

To use Large Receive Offload (LRO) on a VMXNET3 adapter on a virtual machine that runs Windows 8 and later or Windows Server 2012 and later, you must enable LRO globally on the guest operating system. On Windows, the LRO technology is also referred to as Receive Side Coalescing (RSC).

1. To verify whether LRO is deactivated globally on a Windows 8 and later or Windows Server 2012 guest OS, run the `netsh int tcp show global` command at the command prompt.

```
netsh int tcp show global
```

The command displays the status of the global Transmission Control Protocol (TCP) parameters that are set on the Windows 8.x OS.

TCP Global Parameters

```
-----
Receive-Side Scaling State           : enabled
Chimney Offload State                : disabled
NetDMA State                        : disabled
Direct Cache Access (DCA)           : disabled
Receive Window Auto-Tuning Level    : normal
Add-On Congestion Control Provider  : none
ECN Capability                      : disabled
RFC 1323 Timestamps                 : disabled
Initial RTO                         : 3000
Receive Segment Coalescing State    : disabled
```

If LRO is globally deactivated on the Windows 8 and later or Windows Server 2012 machine, the Receive Segment Coalescing State property appears as disabled.

2. To enable LRO globally on the Windows OS, run the `netsh int tcp set global rsc=enabled` command at the command prompt:

```
netsh int tcp set global rsc=enabled
```

Enable LRO for the VMXNET3 adapter on the Windows 8 and later or Windows Server 2012 virtual machine. See [Manage LRO on a VMXNET3 Adapter on a Windows Virtual Machine](#).

NetQueue and Networking Performance

NetQueue takes advantage of the ability of some network adapters to deliver network traffic to the system in multiple receive queues that can be processed separately, allowing processing to be scaled to multiple CPUs, improving receive-side networking performance.

To use Rx queues effectively in the physical NICs by managing the vNIC and the VMkernel adapter filters, the NetQueue balancer in ESXi uses load balancing algorithms.

You can activate or deactivate different types of Rx queues. For more information, see the `esxcli network nic queue loadbalancer set` command in the *ESXCLI Reference* documentation.

Activate NetQueue on a Host

NetQueue is enabled by default. To use NetQueue after it has been deactivated, you must activate it again.

1. In the ESXi Shell on the host, use the following command:

```
esxcli system settings kernel set --setting="netNetqueueEnabled" --value="TRUE"
```

2. To configure the NIC driver to use NetQueue, use the `esxcli module parameters set` command .
To configure the driver with eight receive queues on a dual-port Emulex NIC, run this ESXCLI command.

```
esxcli system module parameters set -m tg3 -p force_netq=8,8
```

3. Reboot the host.

Deactivate NetQueue on a Host

NetQueue is activated by default.

Familiarize yourself with the information on configuring NIC drivers in *Getting Started with ESXCLI*.

1. In the ESXCLI, use the following command depending on the host version:

```
esxcli system settings kernel set --setting="netNetqueueEnabled" --value="FALSE"
```

2. To deactivate NetQueue on the NIC driver, use the `esxcli module parameters set` command.
To configure the driver with one receive queues, run the ESXCLI commands on a dual-port Emulex NIC.

```
esxcli system module parameters set -m tg3 -p force_netq=1,1
```

3. Reboot the host.

vSphere Network I/O Control

Use vSphere Network I/O Control to allocate network bandwidth to business-critical applications and to resolve situations where several types of traffic compete for common resources.

What is vSphere Network I/O Control

vSphere Network I/O Control version 3 introduces a mechanism to reserve bandwidth for system traffic based on the capacity of the physical adapters on a host. It enables fine-grained resource control at the VM network adapter level similar to the model that you use for allocating CPU and memory resources..

Version 3 of the Network I/O Control feature offers improved network resource reservation and allocation across the entire switch.

Models for Bandwidth Resource Reservation

Network I/O Control version 3 supports separate models for resource management of system traffic related to infrastructure services, such as vSphere Fault Tolerance, and of virtual machines.

The two traffic categories have different nature. System traffic is strictly associated with an ESXi host. The network traffic routes change when you migrate a virtual machine across the environment. To provide network resources to a virtual machine regardless of its host, in Network I/O Control you can configure resource allocation for virtual machines that is valid in the scope of the entire distributed switch.

Bandwidth Guarantee to Virtual Machines

Network I/O Control version 3 provisions bandwidth to the network adapters of virtual machines by using constructs of shares, reservation and limit. Based on these constructs, to receive sufficient bandwidth, virtualized workloads can rely on admission control in vSphere Distributed Switch, vSphere DRS and vSphere HA. See [Admission Control for Virtual Machine Bandwidth](#).

Availability of Features

SR-IOV is not available for virtual machines configured to use Network I/O Control version 3.

Enable Network I/O Control on a vSphere Distributed Switch

Enable network resource management on a vSphere Distributed Switch to guarantee minimum bandwidth to system traffic for vSphere features and to virtual machine traffic.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. From the **Actions** menu, select **Settings > Edit Settings**.
3. From the **Network I/O Control** drop-down menu, select **Enable**.

Note: When Network Offloads compatibility is enabled, the **Network I/O Control** is disabled. When **Network Offloads** is set to **None**, **Network I/O Control** is supported.

4. Click **OK**.

When enabled, the model that Network I/O Control uses to handle bandwidth allocation for system traffic and virtual machine traffic is based on the Network I/O Control version that is active on the distributed switch. See [What is vSphere Network I/O Control](#).

Allocate Bandwidth for System Traffic

You can configure Network I/O Control to allocate certain amount of bandwidth for traffic generated by vSphere Fault Tolerance, vSphere vMotion, and so on.

You can use Network I/O Control on a distributed switch to configure bandwidth allocation for the traffic that is related to the main vSphere features:

- Management
- Fault Tolerance
- NFS
- vSAN
- vMotion
- vSphere Replication
- vSphere Data Protection Backup
- Virtual machine
- NVMe over TCP

vCenter Server propagates the allocation from the distributed switch to each physical adapter on the hosts that are connected to the switch.

Bandwidth Allocation Parameters for System Traffic

By using several configuration parameters Network I/O Control allocates bandwidth to traffic from basic vSphere system features.

Table 196: Allocation Parameters for System Traffic

Parameter for Bandwidth Allocation	Description
Shares	Shares, from 1 to 100, reflect the relative priority of a system traffic type against the other system traffic types that are active on the same physical adapter.

Parameter for Bandwidth Allocation	Description
	The amount of bandwidth available to a system traffic type is determined by its relative shares and by the amount of data that the other system features are transmitting.
Reservation	The minimum bandwidth, in Mbps, that must be guaranteed on a single physical adapter. The total bandwidth reserved among all system traffic types cannot exceed 75 percent of the bandwidth that the physical network adapter with the lowest capacity can provide. Reserved bandwidth that is unused becomes available to other types of system traffic. However, Network I/O Control does not redistribute the capacity that system traffic does not use to virtual machine placement.
Limit	The maximum bandwidth, in Mbps or Gbps, that a system traffic type can consume on a single physical adapter.

Example of Bandwidth Reservation for System Traffic

The capacity of the physical adapters determines the bandwidth that you guarantee. According to this capacity, you can guarantee minimum bandwidth to a system feature for its optimal operation.

For example, on a distributed switch that is connected to ESXi hosts with 10 GbE network adapters, you might configure reservation to guarantee 1 Gbps for management through vCenter Server, 1 Gbps for vSphere Fault Tolerance, 1 Gbps for vSphere vMotion traffic, and 0.5 Gbps for virtual machine traffic. Network I/O Control allocates the requested bandwidth on each physical network adapter. You can reserve no more than 75 percent of the bandwidth of a physical network adapter, that is, no more than 7.5 Gbps.

You might leave more capacity unreserved to let the host allocate bandwidth dynamically according to shares, limits, and use, and to reserve only bandwidth that is enough for the operation of a system feature.

How to Allocate Bandwidth for System Traffic

Assign bandwidth for host management, virtual machines, NFS storage, vSphere vMotion, vSphere Fault Tolerance, vSAN, and vSphere Replication on the physical adapters that are connected to a vSphere Distributed Switch.

- Verify that vSphere Distributed Switch is version 6.5.0 and later.
- Verify that Network I/O Control on the switch is version 3.

To enable bandwidth allocation for virtual machines by using Network I/O Control, configure the virtual machine system traffic. The bandwidth reservation for virtual machine traffic is also used in admission control. When you power on a virtual machine, admission control verifies that enough bandwidth is available.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Resource Allocation**.
3. Click **System Traffic**.
You see the bandwidth allocation for the types of system traffic.
4. Select the traffic type according to the vSphere feature that you want to provision and click **Edit**.
The network resource settings for the traffic type appear.

5. From the **Shares** drop-down menu, edit the share of the traffic in the overall flow through a physical adapter.
Network I/O Control applies the configured shares when a physical adapter is saturated.
You can select an option to set a pre-defined value, or select **Custom** and type a number from 1 to 100 to set another share.
6. In the **Reservation** text box, enter a value for the minimum bandwidth that must be available for the traffic type.
The total reservation for system traffic must not exceed 75% of the bandwidth supported by the physical adapter with the lowest capacity of all adapters connected to the distributed switch.
7. In the **Limit** text box, set the maximum bandwidth that system traffic of the selected type can use.
8. Click **OK** to apply the allocation settings.

vCenter Server propagates the allocation from the distributed switch to the host physical adapters that are connected to the switch.

Allocate Bandwidth for Virtual Machine Traffic

Version 3 of Network I/O Control lets you configure bandwidth requirements for individual virtual machines. You can also use network resource pools where you can assign a bandwidth quota from the aggregated reservation for the virtual machine traffic and then allocate bandwidth from the pool to individual virtual machines.

How to Allocate Bandwidth for Virtual Machines

Network I/O Control allocates bandwidth for virtual machines by using two models: allocation across the entire vSphere Distributed Switch based on network resource pools and allocation on the physical adapter that carries the traffic of a virtual machine.

Network Resource Pools


A network resource pool represents a part of the aggregated bandwidth that is reserved for the virtual machine system traffic on all physical adapters connected to the distributed switch.

For example, if the virtual machine system traffic has 0.5 Gbps reserved on each 10 GbE uplink on a distributed switch that has 10 uplinks, then the total aggregated bandwidth available for VM reservation on this switch is 5 Gbps. Each network resource pool can reserve a quota of this 5 Gbps capacity.

The bandwidth quota that is dedicated to a network resource pool is shared among the distributed port groups associated with the pool. A virtual machine receives bandwidth from the pool through the distributed port group the VM is connected to.

By default, distributed port groups on the switch are assigned to a network resource pool, called default, whose quota is not configured.

Figure 86: Bandwidth Aggregation for Network Resource Pools Across the Uplinks of a vSphere Distributed Switch



Defining Bandwidth Requirements for a Virtual Machine


You allocate bandwidth to an individual virtual machine similarly to allocating CPU and memory resources. Network I/O Control version 3 provisions bandwidth to a virtual machine according to shares, reservation, and limits that are defined for a network adapter in the VM hardware settings. The reservation represents a guarantee that the traffic from the virtual machine can consume at least the specified bandwidth. If a physical adapter has more capacity, the virtual machine may use additional bandwidth according to the specified shares and limit.

Bandwidth Provisioning to a Virtual Machine on the Host

To guarantee bandwidth, Network I/O Control implements a traffic placement engine that becomes active if a virtual machine has bandwidth reservation configured. The distributed switch attempts to place the traffic from a VM network adapter to the physical adapter that can supply the required bandwidth and is in the scope of the active teaming policy.

The total bandwidth reservation of the virtual machines on a host cannot exceed the reserved bandwidth that is configured for the virtual machine system traffic.

The actual limit and reservation also depends on the traffic shaping policy for the distributed port group the adapter is connected to. For example, if a VM network adapter requires a limit of 200 Mbps and the average bandwidth configured in the traffic shaping policy is 100 Mbps, then the effective limit becomes 100Mbps.

Figure 87: Configuration for Bandwidth Allocation for Individual Virtual Machines

The reservation of bandwidth for a VM in Network I/O Control is guaranteed against the reservation of VM system traffic on the physical adapter that sends the traffic from the VM.

Bandwidth Allocation Parameters for Virtual Machine Traffic

Network I/O Control version 3 allocates bandwidth to individual virtual machines based on configured shares, reservation, and limit for the network adapters in the VM hardware settings.

Table 197: Bandwidth Allocation Parameters for a VM Network Adapter

Parameter for Bandwidth Allocation	Description
Shares	The relative priority, from 1 to 100, of the traffic through this VM network adapter against the capacity of the physical adapter that is carrying the VM traffic to the network.
Reservation	The minimum bandwidth, in Mbps, that the VM network adapter must receive on the physical adapter.
Limit	The maximum bandwidth on the VM network adapter for traffic to other virtual machines on the same or on another host.

Admission Control for Virtual Machine Bandwidth

To guarantee that sufficient bandwidth is available to a virtual machine, vSphere implements admission control at host and cluster levels based on bandwidth reservation and teaming policy.

Bandwidth Admission Control in vSphere Distributed Switch

When you power on a virtual machine, the Network I/O Control feature on a distributed switch verifies that these conditions are satisfied on the host.

- A physical adapter on the host can supply the minimum bandwidth to the VM network adapters in accordance with the teaming policy and reservation.
- The reservation for a VM network adapter is less than the free quota in the network resource pool.

If you change the reservation for a network adapter of a running virtual machine, Network I/O Control verifies again whether the associated network resource pool can accommodate the new reservation. If the pool does not have enough unclaimed quota, the change is not applied.

To use admission control in vSphere Distributed Switch, perform the following tasks:

- Configure bandwidth allocation for the virtual machine system traffic on the distributed switch.
- Configure a network resource pool with a reservation quota from the bandwidth configured for virtual machine system traffic.
- Associate the network resource pool with the distributed port group that connects the virtual machines to the switch.
- Configure the bandwidth requirements of a virtual machine connected to the port group.

Bandwidth Admission Control in vSphere DRS

If you power on a virtual machine that is in a cluster, vSphere DRS places the virtual machine on a host that has the capacity to guarantee the bandwidth reserved for the virtual machine according to the active teaming policy.

vSphere DRS migrates a virtual machine to another host to satisfy the bandwidth reservation of the virtual machine in these situations:

- The reservation is changed to a value that the initial host can no longer satisfy.
- A physical adapter that carries traffic from the virtual machine is offline.

To use admission control in vSphere DRS, perform the following tasks:

- Configure bandwidth allocation for the virtual machine system traffic on the distributed switch.
- Configure the bandwidth requirements of a virtual machine that is connected to the distributed switch.

For more information about resource management according to bandwidth demands of virtual machines, see the *vSphere Resource Management* documentation.

Bandwidth Admission Control in vSphere HA

When a host fails or is isolated, vSphere HA powers on a virtual machine on another host in the cluster according to the bandwidth reservation and teaming policy.

To use admission control in vSphere HA, perform the following tasks:

- Allocate bandwidth for the virtual machine system traffic.
- Configure the bandwidth requirements of a virtual machine that is connected to the distributed switch.

For more information about vSphere HA provides failover based on the bandwidth demands of virtual machines, see the *vSphere Availability* documentation.

Create a Network Resource Pool

Create network resource pools on a vSphere Distributed Switch to reserve bandwidth for a set of virtual machines.

- Verify that vSphere Distributed Switch is version 6.5.0 and later.
- Verify that Network I/O Control on the switch is version 3.

A network resource pool provides a reservation quota to virtual machines. The quota represents a portion of the bandwidth that is reserved for virtual machine system traffic on the physical adapters connected to the distributed switch. You can set aside bandwidth from the quota for the virtual machines that are associated with the pool. The reservation from the network adapters of powered on VMs that are associated with the pool must not exceed the quota of the pool. See [How to Allocate Bandwidth for Virtual Machines](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Resource Allocation**.
3. Click **Network resource pools**.
4. Click the **Add** icon.
5. Optional: Type a name and a description for the network resource pool.
6. Enter a value for **Reservation quota**, in Mbps, from the free bandwidth that is reserved for the virtual machine system traffic.

The maximum quota that you can assign to the pool is determined according to the following formula:

$$\text{max reservation quota} = \text{aggregated reservation for vm system traffic} - \text{quotas of the other resource pools}$$

where

- $\text{aggregated reservation for vm system traffic} = \text{configured bandwidth reservation for the virtual machine system traffic on each pNIC} * \text{number of pNICs connected to the distributed switch}$
- $\text{quotas of the other pools} = \text{the sum of the reservation quotas of the other network resource pools}$

7. Click **OK**.

Add one or more distributed port groups to the network resource pool so that you can allocate bandwidth to individual virtual machines from the quota of the pool. See [Add a Distributed Port Group to a Network Resource Pool](#).

Add a Distributed Port Group to a Network Resource Pool

Add a distributed port group to a network resource pool so that you can allocate bandwidth to the virtual machines that are connected to the port group.

Verify that Network I/O Control is enabled. See [Enable Network I/O Control on a vSphere Distributed Switch](#).

To assign a network resource pool to several distributed port groups at once, you can use the Resource allocation policy in the **Manage Distributed Port Groups** wizard. See [Manage Policies for Multiple Port Groups on a vSphere Distributed Switch](#).

Network I/O Control allocates bandwidth to the virtual machines associated with the distributed port group according to the model implemented in the Network I/O Control version that is active on the distributed switch. See [What is vSphere Network I/O Control](#).

1. Locate a distributed port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups**.
2. Right-click the distributed port group and select **Edit settings**.
3. In the Edit Settings dialog box, click **General**.
4. From the **Network resource pool** drop-down menu, select the network resource pool and click **OK**.
If the distributed switch does not contain network resource pools, you see only the **(default)** option in the drop-down menu.

Configure Bandwidth Allocation for a Virtual Machine

You can configure bandwidth allocation to individual virtual machines that are connected to a distributed port group. You can use shares, reservation, and limit settings for bandwidth.

- Verify that vSphere Distributed Switch is version 6.5.0 and later.
 - Verify that Network I/O Control on the switch is version 3.
1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
 2. From the Actions menu, select Edit Settings.
 3. Expand the Network adapter section of the VM network adapter.
 4. If you want to configure bandwidth allocation for a new VM network adapter, from the **Add new device** drop-down menu select **Network adapter**.
A New Network section displays options for bandwidth allocation and other network adapter settings.
 5. If the VM network adapter is not connected to the distributed port group, select the port group from the drop-down menu next to the Network adapter or New Network label.
 6. From the **Shares** drop-down menu, set the relative priority of the traffic from this virtual machine as shares from the capacity of the connected physical adapter.
Network I/O Control applies the configured shares when a physical adapter is saturated.
You can select an option to set a pre-defined value, or select **Custom** and type a number from 1 to 100 to set another share.
 7. In the **Reservation** text box, reserve a minimum bandwidth that must be available to the VM network adapter when the virtual machine is powered on.
If you provision bandwidth by using a network resource pool, the reservation from the network adapters of powered on VMs that are associated with the pool must not exceed the quota of the pool.
If vSphere DRS is enabled, to power on the virtual machine, make sure that the reservation from all VM network adapters on the host does not exceed the bandwidth reserved for virtual machine system traffic on the host physical adapters.

8. In the **Limit** text box, set a limit on the bandwidth that the VM network adapter can consume.
9. Click **OK**.

Network

I/O Control allocates the bandwidth that you reserved for the network adapter of the virtual machine out of the reservation quota of the network resource pool.

Configure Bandwidth Allocation on Multiple Virtual Machines

With a single operation, configure bandwidth allocation on multiple virtual machines that are connected to a specific network resource pool, for example, after you upgrade Network I/O Control to version 3.

- Verify that vSphere Distributed Switch is version 6.5.0 and later.
- Verify that Network I/O Control on the switch is version 3.
- Verify that the virtual machines are associated with a specific network resource pool through the connected distributed port groups. See [Add a Distributed Port Group to a Network Resource Pool](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Resource Allocation**.
3. Click **Network resource pools**.
4. Select a network resource pool.
5. Click **Virtual Machines**.

A list of the VM network adapters that are connected to the selected network resource pool appears.

6. Select the VM network adapters whose settings you want to configure and click **Edit**.
7. From the **Shares** drop-down menu, set the relative priority of traffic from these virtual machines in the scope of the physical adapters that carry the traffic.

Network I/O Control applies the configured shares when a physical adapter is saturated.

8. In the **Reservation** text box, reserve a minimum bandwidth that must be available to each VM network adapter when the virtual machines are powered on.

If you provision bandwidth by using a network resource pool, the reservation from the network adapters of powered on VMs that are associated with the pool must not exceed the quota of the pool.

9. In the **Limit** text box, set a limit on the bandwidth that each VM network adapter can consume.
10. Click **OK**.

Modify the Quota of a Network Resource Pool

Learn how to modify the quota of bandwidth that you can reserve for virtual machines connected to a set of distributed port groups.

- Verify that vSphere Distributed Switch is version 6.5.0 and later.

- Verify that Network I/O Control on the switch is version 3.
1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
 2. On the **Configure** tab, expand **Resource Allocation**.
 3. Click **Network resource pools**.
 4. Select a network resource pool from the list and click **Edit**.
 5. In the **Reservation quota** text box, enter the bandwidth quota for virtual machines from the aggregation of free bandwidth that is reserved for virtual machine system traffic on all physical adapters on the switch.
 6. Click **OK**.

Remove a Distributed Port Group from a Network Resource Pool

To stop allocating bandwidth to the virtual machines from the reservation quota of a network resource pool, remove the association between the port group to which the virtual machines are connected and the pool.

1. Locate a distributed port group in the vSphere Client.
 - a) Select a distributed switch and click the **Networks** tab.
 - b) Click **Distributed Port Groups**.
2. Right-click the distributed port group and select **Edit Settings**.
3. In the Edit Settings dialog box for the port group, click **General**.
4. From the **Network resource pool** drop-down menu, select **(default)** and click **OK**.

The distributed port group becomes associated with the default VM network resource pool.

Delete a Network Resource Pool

Delete a network resource pool that is no longer in use.

Uncouple the network resource pool from all associated distributed port groups. See [Remove a Distributed Port Group from a Network Resource Pool](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Configure** tab, expand **Resource Allocation**.
3. Click **Network resource pools**.
4. Select a network resource pool and click **Remove**.
5. Click **OK** to delete the resource pool.

Move a Physical Adapter Out of Network I/O Control

Under certain conditions you might need to exclude physical adapters with low capacity from the bandwidth allocation model of Network I/O Control version 3.

- Verify that the host is running ESXi 6.5 and later.

- Verify that vSphere Distributed Switch is version 6.5.0 and later.
- Verify that Network I/O Control on the switch is version 3.

For example, if the bandwidth allocation on a vSphere Distributed Switch is tailored on top of 10 GbE NICs, you might not be able to add a 1GbE NIC to the switch because it cannot meet the higher allocation requirements configured on the 10GbE NICs.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **System** and select **Advanced System Settings** .
3. Set the physical adapters that you need to function outside the scope of Network I/O Control as a comma-separated list to the `Net.IOControlPnicOptOut` parameter.
For example: `vmnic0,vmnic3`
4. Click **OK** to apply the changes.

MAC Address Management

MAC addresses are used in the Layer 2 (Data Link Layer) of the network protocol stack to transmit frames to a recipient. In vSphere, vCenter Server generates MAC addresses for virtual machine adapters and VMkernel adapters, or you can assign addresses manually.

Each network adapter manufacturer is assigned a unique three-byte prefix called an Organizationally Unique Identifier (OUI), which it can use to generate unique MAC addresses.

VMware supports several address allocation mechanisms, each of them with a separate OUI:

- Generated MAC addresses
 - Assigned by vCenter Server
 - Assigned by the ESXi host
- Manually set MAC addresses
- Generated for legacy virtual machines, but no longer used with ESXi

If you reconfigure the network adapter of a powered off virtual machine, for example by changing the automatic MAC address allocation type, or setting a static MAC address, vCenter Server resolves any MAC address conflict before the adapter reconfiguration takes effect.

MAC Address Assignment from vCenter Server

vSphere provides several schemes for automatic allocation of MAC addresses in vCenter Server. You can select the scheme that best suits your requirements for MAC address duplication, OUI requirements for locally administered or universally administered addresses, and so on.

The following schemes of MAC address generation are available in vCenter Server:

- VMware OUI allocation, default allocation
- Prefix-based allocation
- Range-based allocation

After the MAC address is generated, it does not change unless the virtual machine's MAC address conflicts with that of another registered virtual machine. The MAC address is saved in the configuration file of the virtual machine.

Note: If you use invalid prefix- or range-based allocation values, an error is logged in the `vpd.log` file. vCenter Server does not allocate MAC addresses when provisioning a virtual machine.

Preventing MAC Address Conflicts

The MAC address of a powered off virtual machine is not checked against the addresses of running or suspended virtual machines.

When a virtual machine is powered on again, it might acquire a different MAC address. The change might be caused by an address conflict with another virtual machine. While this virtual machine has been powered off, its MAC address has been assigned to another virtual machine that has been powered on.

If you reconfigure the network adapter of a powered off virtual machine, for example, by changing the automatic MAC address allocation type or setting a static MAC address, vCenter Server resolves MAC address conflicts before the adapter reconfiguration takes effect.

For information about resolving MAC address conflicts, see the *vSphere Troubleshooting* documentation.

What is VMware OUI Allocation

VMware Organizationally Unique Identifier (OUI) allocation assigns MAC addresses based on the default VMware OUI 00:50:56 and the vCenter Server ID.

VMware OUI allocation is the default MAC address assignment model for virtual machines. The allocation works with up to 64 vCenter Server instances, and each vCenter Server can assign up to 64000 unique MAC addresses. The VMware OUI allocation scheme is suitable for small scale deployments.

MAC Address Format

According to the VMware OUI allocation scheme, a MAC address has the format 00:50:56:XX:YY:ZZ where 00:50:56 represents the VMware OUI, XX is calculated as (128+ vCenter Server ID), and YY and ZZ are random two-digit hexadecimal numbers.

The addresses created through the VMware OUI allocation are in the range 00:50:56:80:YY:ZZ - 00:50:56:BF:YY:ZZ.

What is Prefix-Based MAC Address Allocation

You can use prefix-based allocation to specify an OUI other than the default one 00:50:56 by VMware, or to introduce Locally Administered MAC Addresses (LAA) for a larger address space.

Prefix-based MAC address allocation overcomes the limits of the default VMware allocation to provide unique addresses in larger scale deployments. Introducing an LAA prefix leads to a very large MAC address space (2 to the power of 46) instead of an universally unique address OUI which can give only 16 million MAC addresses.

Verify that the prefixes that you provide for different vCenter Server instances in the same network are unique. vCenter Server relies on the prefixes to avoid MAC address duplication issues, see [Troubleshooting MAC Address Allocation](#).

What is Range-Based MAC Address Allocation

You can use range-based allocation to include or exclude ranges of Locally Administered Addresses (LAA).

You specify one or more ranges using a starting and ending MAC addresses, for example, (02:50:68:00:00:02 , 02:50:68:00:00:FF). MAC addresses are generated only from within the specified range.

You can specify multiple ranges of LAA, and vCenter Server tracks the number of used addresses for each range. vCenter Server allocates MAC addresses from the first range that still has addresses available. vCenter Server checks for MAC address conflicts within its ranges.

When using range-based allocation, you must provide different instances of vCenter Server with ranges that do not overlap. vCenter Server does not detect ranges that might be in conflict with other vCenter Server instances. See the [Troubleshoot MAC Address Allocation](#) for more information about resolving issues with duplicate MAC addresses.

Note: Range-based MAC address allocation settings are lost when you upgrade to a new version of vCenter Server. You must manually recreate the range-based MAC address allocation settings after the upgrade.

Assigning a MAC Address

Use the vSphere Client to enable prefix-based or range-based MAC address allocation and to adjust the allocation parameters.

If you are changing from one type of allocation to another, for example changing from the VMware OUI allocation to a range-based allocation, use the vSphere Client. However, when a schema is prefix-based or range-based and you want to change to a different allocation schema, you must edit the `vpzd.cfg` file manually and restart vCenter Server.

Change to or Adjust Range- or Prefixed-Based Allocations

By switching from the default VMware OUI to range- or prefixed-based MAC address allocation through the vSphere Client, you can avoid and resolve MAC address duplication conflicts in vSphere deployments.

Change the allocation scheme from the default VMware OUI to range- or to prefixed-based allocation by using the **Advanced Settings** available for the vCenter Server instance in the vSphere Client.

To switch from range- or prefixed-based allocation back to VMware OUI allocation, or between range- and prefixed-based allocation, edit the `vpzd.cfg` file manually. See [Set or Change Allocation Type](#).

1. In the vSphere Client, navigate to vCenter Server instance.
2. On the **Configure** tab, expand **Settings** and select **Advanced Settings**.
3. Click **Edit Settings**.
4. Add or edit parameters for the target allocation type.
Use only one allocation type.
 - Change to prefix-based allocation.

Key	Example Value
<code>config.vpzd.macAllocScheme.prefixScheme.prefix</code>	005026
<code>config.vpzd.macAllocScheme.prefixScheme.prefixLength</code>	23

`prefix` and `prefixLength` determine the range of MAC address prefixes that newly added vNICs have. `prefix` is the starting OUI of MAC addresses related to the instance, and `prefixLength` determines the length of the prefix in bits.

For example, the settings from the table result in VM NIC MAC addresses starting with 00:50:26 or 00:50:27.

- Change to range-based allocation.

Key	Example Value
<code>config.vpzd.macAllocScheme.rangeScheme.range[X].begin</code>	005067000000

Key	Example Value
<code>config.vpxd.macAllocScheme.rangeScheme.range[X].end</code>	005067ffff

`X` in `range[X]` stands for the range sequence number. For example, 0 in `range[0]` represents the allocation settings of the first range for MAC address allocation.

5. Click **Save**.

Set or Change Allocation Type

If you are changing from range- or prefixed-based allocation to the VMware OUI allocation, you must set the allocation type in the `vpxd.cfg` file and restart the vCenter Server.

Decide on an allocation type before changing the `vpxd.cfg` file. For information on allocation types, see [MAC Address Assignment from vCenter Server](#)

1. On the host machine of , navigate to the directory `/etc/vmware-vpx`.
2. Open the `vpxd.cfg` file.
3. Decide on an allocation type to use and enter the corresponding XML code in the file to configure the allocation type. The following are examples of XML code to use.

Note: Use only one allocation type.

- VMware OUI allocation

```
<vpxd>
  <macAllocScheme>
    <VMwareOUI>true</VMwareOUI>
  </macAllocScheme>
</vpxd>
```

- Prefix-based allocation

```
<vpxd>
  <macAllocScheme>
    <prefixScheme>
      <prefix>005026</prefix>
      <prefixLength>23</prefixLength>
    </prefixScheme>
  </macAllocScheme>
</vpxd>
```

- Range-based allocation

```
<vpxd>
  <macAllocScheme>
    <rangeScheme>
      <range id="0">
        <begin>005067000001</begin>
        <end>005067000001</end>
      </range>
    </rangeScheme>
  </macAllocScheme>
</vpxd>
```

4. Save the `vpzd.cfg`.
5. Restart the vCenter Server host.

MAC Address Generation on ESXi Hosts

An ESXi host generates the MAC address for a virtual machine adapter when the host is not connected to vCenter Server. Such addresses have a separate VMware OUI to avoid conflicts.

The ESXi host generates the MAC address for a virtual machine adapter in one of the following cases:

- The host is not connected to vCenter Server.
- The virtual machine configuration file does not contain the MAC address and information about the MAC address allocation type.

MAC Address Format

The host generates MAC addresses that consists of the VMware OUI `00:0C:29` and the last three octets in hexadecimal format of the virtual machine UUID. The virtual machine UUID is based on a hash calculated by using the UUID of the ESXi physical machine and the path to the configuration file (`.vmtx`) of the virtual machine.

Preventing MAC Address Conflicts

All MAC addresses that have been assigned to network adapters of running and suspended virtual machines on a given physical machine are tracked for conflicts.

If you import a virtual machine with a host-generated MAC address from one vCenter Server to another, select the **Imported** option when you power on the virtual machine to regenerate the address and avoid potential conflicts in the target vCenter Server or between the vCenter Server systems.

How to set a Static MAC Address to a Virtual Machine

In most network deployments, generated MAC addresses are a good approach. However, you might need to set a static MAC address for a virtual machine adapter with unique value.

The following cases show when you might set a static MAC address:

- Virtual machine adapters on different physical hosts share the same subnet and are assigned the same MAC address, causing a conflict.
- Ensure that a virtual machine adapter always has the same MAC address.

By default, VMware uses the Organizationally Unique Identifier (OUI) `00:50:56` for manually generated addresses, but all unique manually generated addresses are supported.

Note: Make sure that no other non-VMware devices use addresses assigned to VMware components. For example, you might have physical servers in the same subnet, which use `11:11:11:11:11:11`, `22:22:22:22:22:22` as static MAC addresses. The physical servers do not belong to the vCenter Server inventory, and vCenter Server is not able to check for address collision.

VMware OUI in Static MAC Addresses

By default, static MAC addresses have the VMware Organizationally Unique Identifier (OUI) as the prefix. However, the range of free address provided by the VMware OUI is restricted.

If you decide to use the VMware OUI, part of the range is reserved for use by vCenter Server, host physical NICs, virtual NICs, and for future use.

You can set a static MAC address that contains the VMware OUI prefix in compliance with the following format:

00:50:56:XX:YY:ZZ

where *XX* is a valid hexadecimal number between 00 and 3F, and *YY* and *ZZ* are valid hexadecimal numbers between 00 and FF. To avoid conflict with MAC addresses that are generated by vCenter Server or are assigned to VMkernel adapters for infrastructure traffic, the value for *XX* must not be greater than 3F.

The maximum value for a manually generated MAC address is as follows.

00:50:56:3F:FF:FF

To avoid conflicts between the generated MAC addresses and the manually assigned ones, select a unique value for *XX:YY:ZZ* from your hard-coded addresses.

Assign a Static MAC Address

You can assign static MAC addresses to the virtual NIC of a powered off virtual machine by using the vSphere Client.

1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
2. Power off the virtual machine.
3. From the Actions menu, select Edit Settings.
4. Select the **Virtual Hardware** tab in the dialog box displaying the settings.
5. Expand the network adapter section.
6. Under MAC Address, select **Manual** from the drop-down menu.
7. Type the static MAC address, and click **OK**.
8. Power on the virtual machine.

Assign a Static MAC Address in the Virtual Machine Configuration File

To set a static MAC address for a virtual machine, you can edit the configuration file of the virtual machine by using the vSphere Client.

1. Locate the virtual machine in the .
 - a) Select a data center, folder, cluster, resource pool, or host and click the **VMs** tab.
 - b) Click **Virtual Machines** and click the virtual machine from the list.
2. Power off the virtual machine.
3. From the **Actions** drop-down menu, select **Edit Settings**.
4. Select the **VM Options** tab and expand **Advanced**.
5. Click **Edit Configuration**.
6. To assign a static MAC address, add or edit parameters as required.

Parameter	Value
ethernet X.addressType	static

Parameter	Value
ethernet X .address	MAC_address_of_the_virtual_NIC

X next to `ethernet` stands for the sequence number of the virtual NIC in the virtual machine.

For example, 0 in `ethernet0` represents the settings of the first virtual NIC device added to the virtual machine.

7. Click **OK**.

8. Power on the virtual machine.

Configure vSphere for IPv6

Configure ESXi hosts and vCenter Server for operation in a pure IPv6 environment for larger address space and improved address assignment.

IPv6 is designated by the Internet Engineering Task Force (IETF) as the successor to IPv4 providing the following benefits:

- Increased address length. The increased address space resolves the problem of address exhaustion and eliminates the need for network address translation. IPv6 uses 128-bit addresses compared with the 32-bit addresses used by IPv4.
- Ability for improved address autoconfiguration of nodes.

vSphere IPv6 Connectivity

In an environment that is based on vSphere 6.0 and later, nodes and features can communicate over IPv6 transparently supporting static and automatic address configuration.

IPv6 in the Communication between vSphere Nodes

The nodes in a vSphere deployment can communicate using IPv6 and accept assigned addresses according to the network configuration.

Table 198: IPv6 Support of the Nodes in a vSphere Environment

Connection Type	IPv6 Support	Address Configuration of vSphere nodes
ESXi to ESXi	Yes	• Static • Automatic: AUTOCONF/DHCPv6
vCenter Server machine to ESXi	Yes	• Static • Automatic: AUTOCONF/DHCPv6
vCenter Server machine to machine	Yes	• Static • Automatic: AUTOCONF/DHCPv6
ESXi to vSphere Client machine	Yes	• Static • Automatic: AUTOCONF/DHCPv6
Virtual machine to virtual machine	Yes	• Static • Automatic: AUTOCONF/DHCPv6
ESXi to iSCSI Storage	Yes	• Static • Automatic: AUTOCONF/DHCPv6
ESXi to NFS Storage	Yes	• Static • Automatic: AUTOCONF/DHCPv6

Connection Type	IPv6 Support	Address Configuration of vSphere nodes
ESXi to Active Directory	No Use LDAP through vCenter Server to connect ESXi to the Active Directory database	-
vCenter Server to Active Directory	No Use LDAP to connect vCenter Server to the Active Directory database	-

IPv6 Connectivity of vSphere Features

Certain vSphere features do not support IPv6:

- vSphere DPM over Intelligent Platform Management Interface (IPMI) and Hewlett-Packard Integrated Lights-Out (iLO). vSphere 6.5 and later supports only Wake-On-LAN (WOL) to bring a host out of standby mode.
- Authentication Proxy
- vSphere Management Assistant and ESXCLI connected to Active Directory.
Use LDAP to connect the vSphere Management Assistant or the ESXCLI to the Active Directory database.

IPv6 Connectivity of Virtual Machines

Virtual machines can exchange data in the network over IPv6. vSphere supports both static and automatic assignment of IPv6 addresses for virtual machines.

Configuring one or more IPv6 addresses is also possible when you customize the guest operating system of a virtual machine.

vSAN Connectivity

vSAN supports IPv6. Use NFS 4.1 with AUTH_SYS.

FQDNs and IPv6 Addresses

In vSphere, you should use fully qualified domain names (FQDNs) that are mapped to IPv6 addresses on the DNS server. You can use IPv6 addresses if they have a valid FQDN on the DNS server for reverse lookup.

To deploy vCenter Server in a pure IPv6 environment, you must use FQDNs only.

Deploy vSphere on IPv6

Run vSphere in a pure IPv6 environment to use an extended address space and flexible address assignment.

If you plan to deploy vCenter Server and ESXi hosts in an IPv6 network, you must perform additional steps.

Enable IPv6 on a vSphere Installation

If you have a greenfield deployment of vSphere 6.5 in an IPv6 network, configure ESXi and vCenter Server for pure IPv6 management connection by configuring IPv6 on the deployment nodes and connecting them.

- Verify that the IPv6 addresses for vCenter Server, the ESXi hosts and an external database, if used, are mapped to fully qualified domain names (FQDNs) on the DNS server.

- Verify that the network infrastructure provides IPv6 connectivity for the ESXi hosts, vCenter Server and external database if used.
 - Verify that you have version 6.5 of vCenter Server installed with FQDN that is mapped to an IPv6 address. See the *vCenter Server Installation and Setup* documentation.
 - Verify that the hosts have ESXi 6.5 installed. See the *vCenter Server Installation and Setup* documentation.
1. In the Direct Console User Interface (DCUI), configure each ESXi host as a pure IPv6 node.
 - a) In the DCUI, press F2 and log in to the host.
 - b) From the **Configure Management Network** menu, select **IPv6 Configuration** and press Enter.
 - c) Assign an IPv6 address to the host.

Address Assignment Option	Description
Automatic address assignment using DHCPv6	1. Select the Use dynamic IPv6 address and network configuration option and select Use DHCPv6. 2. Press Enter to save the changes.
Static address assignment	1. Select the Set static IPv6 address and network configuration option and enter the IPv6 address of the host and the default gateway. 2. Press Enter to save the changes.

- d) From the **Configure Management Network** menu, select **IPv4 Configuration** and press Enter.
 - e) Select **Disable IPv4 configuration for management network** and press Enter.
2. In the vSphere Client, add the hosts to the inventory.

Enable IPv6 on an Upgraded vSphere Environment

On an IPv4 deployment of vSphere 6.5 that consists of an installed or upgraded vCenter Server and upgraded ESXi, configure ESXi and vCenter Server for pure IPv6 management connection by enabling IPv6 on the deployed nodes and reconnecting them.

- Verify that the network infrastructure provides IPv6 connectivity for the ESXi hosts, vCenter Server and external database if used.
- Verify that the IPv6 addresses for vCenter Server, the ESXi hosts and an external database, if used, are mapped to fully qualified domain names (FQDNs) on the DNS server.
- Verify that you have version 6.x of vCenter Server installed or upgraded. See the *vCenter Server Installation and Setup* and *vCenter Server Upgrade* documentation.
- Verify that all ESXi hosts are upgraded to version 6.x. See the *VMware ESXi Upgrade* documentation.

1. In the vSphere Client, disconnect the hosts from vCenter Server.
2. Configure each ESXi host as a pure IPv6 node.
 - a) Open an SSH connection and log in to the ESXi host.
 - b) Run the following command:


```
esxcli network ip interface ipv6 set -i vmk0 -e true
```
 - c) Assign an IPv6 address to the management network.

Address Assignment Option	Description
Static address assignment	1. Open an SSH connection and log in to the ESXi host.

Address Assignment Option	Description
	- Set a static IPv6 address for the management network vmk0 by running the following command: <pre>esxcli network ip interface ipv6 address add -I IPv6_address -i vmk0</pre> - Set the default gateway for the management network vmk0 by running the following command: <pre>esxcli network ip interface ipv6 set -i vmk0 -g default_gateway_IPv6_address</pre> - Add a DNS server by running the following command: <pre>esxcli network ip dns server add -s DNS_server_IPv6_address</pre>
Automatic address assignment using DHCPv6	- Open an SSH connection and log in to the ESXi host. - Enable DHCPv6 for the management network vmk0 by running the following command: <pre>esxcli network ip interface ipv6 -i vmk0 -enable-dhcpv6 = true</pre> - Enable IPv6 Router Advertised for the management network vmk0 by running the following command: <pre>esxcli network ip interface ipv6 set -i vmk0 -enable-router-adv =true</pre> - Add a DNS server or use the DNS setting published by DHCPv6 by running one the following commands: <pre>esxcli network ip dns server add -s DNS_server_IPv6_address</pre> <pre>esxcli network ip interface ipv6 set -i vmk0 --peer-dns=true</pre>

3. Deactivate the IPv4 configuration for the management network.

- Open an SSH connection and log in to the ESXi host.
- Run the following command:

```
esxcli network ip interface ipv4 set -i vmk0 --type=none
```

4. If vCenter Server uses an external database, configure the database as an IPv6 node.
5. Configure vCenter Server as a pure IPv6 node and restart it.
6. Deactivate the IPv4 on the database server.
7. In the vSphere Client, add the hosts to the inventory.
8. Deactivate the IPv4 in the network infrastructure.


Activate or Deactivate IPv6 Support on a Host

The IPv6 support in vSphere lets hosts work in an IPv6 network that has a large address space, enhanced multicasting, simplified routing, and so on.

In ESXi 6.0 and later releases, IPv6 is enabled by default.

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **TCP/IP Configuration**.
3. Click **Edit**.

Figure 88: TCP/IP Configuration



4. Use the toggle button to activate or deactivate IPv6 support.
5. Click **OK**.
6. Reboot the host to apply the changes in the IPv6 support.

Configure the IPv6 settings of VMkernel adapters on the host, for example, of the management network. See [Set Up IPv6 on an ESXi Host](#).

Set Up IPv6 on an ESXi Host

To connect an ESXi host over IPv6 to the management network, vSphere vMotion, shared storage, vSphere Fault Tolerance, and so on, edit the IPv6 settings of the VMkernel adapters on the host.

Verify that IPv6 is enabled on the ESXi host. See [Activate or Deactivate IPv6 Support on a Host](#) .

1. In the , navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **VMkernel adapters**.
3. Select the VMkernel adapter on the target distributed or standard switch and click **Edit**.
4. In the Edit Settings dialog box, click **IPv6 settings**.
5. Configure the address assignment of the VMkernel adapter.

IPv6 Address Option	Description
Obtain IPv6 address automatically through DHCP	Receive an IPv6 address for the VMkernel adapter from a DHCPv6 server.
Obtain IPv6 address automatically through Router Advertisement	Receive an IPv6 address for the VMkernel adapter from a router through Router Advertisement.
Static IPv6 addresses	Set one or more addresses. For each address entry, enter the IPv6 address of the adapter, subnet prefix length and IPv6 address of the default gateway.

You can select several assignment options according the configuration of your network.

6. Click **OK** to apply the changes on the VMkernel adapter.

Set Up IPv6 on vCenter Server

Use the vSphere Client to configure vCenter Server for communication with ESXi hosts in an IPv6 network.

1. On the vSphere Client main page, click **Administration** > **Deployment** > **System Configuration**.
2. Under **System Configuration**, select a node from the list.
3. Click **Login**.
4. Use the vSphere Client to log in as `administrator@your_domain_name` to the vCenter Server instance.
The address is of the type `http://appliance-IP-address-or-FQDN/ui`.
5. Under Common, select **Networking** and click **Edit**.
6. Expand the network interface name to edit the IP address settings.
7. Edit the IPv6 address settings.

Option	Description
Enable or Disable IPv6 settings	Enables or disables the IPv6 address based on the toggle switch option.
Obtain IPv6 settings automatically through DHCP	Assigns IPv6 addresses to the appliance automatically from the network by using DHCP.
Obtain IPv6 settings automatically through router advertisement	Assigns IPv6 addresses to the appliance automatically from the network by using router advertisement.
Use static IPv6 addresses	Uses static IPv6 addresses that you set up manually. 1. Click the checkbox. 2. Enter the IPv6 address and the subnet prefix length. 3. Click Add to enter additional IPv6 addresses. 4. Click Save.

Option	Description
	Note: For static IPv4 or IPv6 addresses, DNS server must be set manually.

You can configure the appliance to obtain the IPv6 settings automatically through both DHCP and router advertisement. You can assign static a IPv6 address at the same time.

Note: Second and third party solutions need to re-register, when there is a change in the IPv4 and IPv6 IP addresses.

- Optional: To remove IPv6 addresses that are assigned automatically through Router Advertisement, click **Remove Addresses** and delete the addresses.

You might want to delete certain IPv6 addresses that vCenter Server obtained through Router Advertisement to stop the communication on these addresses and to enforce the configured static addresses.

Connect the ESXi hosts to vCenter Server over IPv6 by using their FQDNs.

Monitor Network Connection and Traffic

Monitor network connection and packets that pass through the ports of a vSphere Standard Switch or a vSphere Distributed Switch to analyze the traffic between virtual machines and hosts.

How to Capture Network Packets

Learn how to use the PacketCapture utility to diagnose networking problems such as slow connection, lost packets, and connectivity problems

PacketCapture is a lightweight tcpdump utility that captures and stores only the minimum amount of data that is needed to diagnose the network problem. PacketCapture is integrated in the rhttpproxy service of ESXi and vCenter Server. You start and stop PacketCapture by editing the rhttpproxy service XML configuration file.

Note: For ESXi 8.0 and earlier, the process for enabling network packet captures was by editing the rhttpproxy service XML configuration file. Starting ESXi 8.0 Update 1 and later, reverse proxy configuration has been moved from XML files to the configstore database. To enable, Capture Network Packets with configstorecli, see <https://kb.vmware.com/s/article/89489>.

- Start capturing packets.
 - Open an SSH connection and log in to the ESXi host or vCenter Server.
 - Open the `config.xml` file for editing.

vSphere Component	File Location
ESXi	<code>/etc/vmware/rhttpproxy/config.xml</code>
vCenter Server	<code>/etc/vmware-rhttpproxy/config.xml</code>

- Make the following changes.

```
<config>
  <packetCapture>
    <enabled>true</enabled>
```

- Optional: Configure PacketCapture options.

Option and Default Value	Description
<code><validity>72</validity></code>	On startup delete all <code>pcap</code> and <code>pcap.gz</code> files that were last modified before the specified period of hours and are not part of the current process.

Option and Default Value	Description
<code><directory>/directory_path</directory></code>	The directory in which <code>pcap</code> and <code>pcap.gz</code> files are stored. The directory must exist and be accessible.
<code><maxDataInPcapFile>52428800</maxDataInPcapFile></code>	The amount of captured data in bytes that each <code>pcap</code> and <code>pcap.gz</code> file can store before rolling over to the next file. The minimum size is 5 MB on vCenter Server and 2.5MB on ESXi. Note: Storing 50 MB of captured data in a <code>pcap</code> file requires a <code>pcap</code> file of about 67.5 MB.
<code><maxPcapFilesCount>5</maxPcapFilesCount></code>	The number of <code>pcap</code> or <code>pcap.gz</code> files to rotate. The minimum number is 2.

- e) Save and close the `config.xml` file.
- f) Reload the `config.xml` file by running the following command.

```
kill -SIGHUP `pidof rhttpproxy`
```

2. Stop capturing packets.

- a) Open an SSH connection and log in to the ESXi host or vCenter Server.
- b) Open the `config.xml` file for editing.
- c) Make the following changes.

```
<config>
  <packetCapture>
    <enabled>false</enabled>
```

- d) Save and close the `config.xml` file.
- e) Reload the `config.xml` file by running the following command.

```
kill -SIGHUP `pidof rhttpproxy`
```

3. Collect the captured data.

The `pcap` or `pcap.gz` files are stored in the following default directories.

vSphere Component	File Location
ESXi	<code>/var/run/log</code>
vCenter Server	<code>/var/log/vmware/rhttpproxy</code>

Copy the `pcap` and `pcap.gz` files to a system that runs a network analyzer tool, such as Wireshark, and examine the packet details.

Before you analyze the `pcap` and `pcap.gz` captured from an ESXi host, use the TraceWrangler utility to fix the frame size metadata. For more information, see <https://kb.vmware.com/kb/52843>

Capture and Trace Network Packets by Using the `pktcap-uw` Utility

Learn how to monitor the traffic that flows through physical network adapters, VMkernel adapters, and virtual machines adapters, and analyze packet information by using the graphical user interface of network analysis tools such as Wireshark.

In vSphere you can monitor packets on a host by using the `pktcap-uw` console utility. You can use the utility without additional installation on an ESXi host. `pktcap-uw` provides many points in the host network stack at which you can monitor traffic.

For detailed analysis of captured packets, you can save packet content from the `pktcap-uw` utility to files in PCAP or PCAPNG format and open them in Wireshark. You can also troubleshoot dropped packets and trace a packet's path in the network stack.

Note: The `pktcap-uw` utility is not fully supported for backward compatibility across vSphere releases. The options of the utility might change in the future.

pktcap-uw Command Syntax for Capturing Packets

Use the `pktcap-uw` utility to inspect the contents of packets while they traverse the network stack on an ESXi host.

pktcap-uw Syntax for Capturing Packets

The `pktcap-uw` command has the following syntax for capturing packets at a certain place in the network stack:

```
pktcap-uw
  switch_port_arguments
  capture_point_options
  filter_options
  output_control_options
```

Note: Certain options of the `pktcap-uw` utility are designed for VMware internal use only and you should use them only under the supervision of VMware Technical Support. These options are not described in the *vSphere Networking* guide.

Table 199: pktcap-uw Arguments for Capturing Packets

Argument Group	Argument	Description
<code>switch_port_arguments</code>	<code>--uplink vmnicX</code>	Capture packets that are related to a physical adapter. You can combine the <code>--uplink</code> and <code>--capture</code> options for monitoring packets at a certain place in the path between the physical adapter and the virtual switch. See Capture Packets That Arrive at a Physical Adapter .
	<code>--vmk vmkX</code>	Capture packets that are related to a VMKernel adapter. You can combine the <code>vmk</code> and <code>--capture</code> options for monitoring packets at a certain place in the path between the VMkernel adapter and the virtual switch. See How to Capture Packets for a VMkernel Adapter .
	<code>--switchport {vmxnet3_port_ID vmkernel_adapter_port_ID}</code>	Capture packets that are related to a VMXNET3 virtual machine adapter or to a VMkernel adapter that is connected to a particular virtual switch port. You can view the ID of the port in the network panel of the <code>esxtop</code> utility.

Argument Group	Argument	Description
		You can combine the switchport and capture options for monitoring packets at a certain place in the path between the VMXNET3 adapter or VMkernel adapter and the virtual switch. See How to Capture Packets for a VMXNET3 Virtual Machine Adapter .
	--lifID <i>lif_ID</i>	Capture packets that are related to the logical interface of a distributed router. See the <i>VMware NSX</i> documentation.
<i>capture_point_options</i>	--capture <i>capture_point</i>	Capture packets at a particular place in the network stack. For example, you can monitor packets right after they arrive from a physical adapter.
	--dir {0 1 2}	Capture packets according to the direction of the flow with regard to the virtual switch. 0 stands for incoming traffic, 1 for outgoing traffic, and 2 for bidirectional traffic. By default, the <code>pktcap-uw</code> utility captures ingress traffic. Use the --dir option together with the --uplink, --vmk, or --switchport option.
	--stage {0 1}	Capture the packet closer to its source or to its destination. Use this option to examine how a package changes while it traverses the points in the stack. 0 stands for traffic closer to source and 1 for traffic closer to destination. Use the --stage option together with the --uplink, --vmk, --switchport, or --dvfilter option.
	--dvfilter <i>filter_name</i> --capture PreDVFilter PostDVFilter	Capture packets before or after a vSphere Network Appliance (DVFilter) intercepts them. See How to Capture Packets at DVFilter Level .
	-A --availpoints	View all capture points that the <code>pktcap-uw</code> utility supports.
	For details about the capture points of the <code>pktcap-uw</code> utility, see Capture Points of the pktcap-uw Utility .	
<i>filter_options</i>	Filter captured packets according to source or destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port. See pktcap-uw Options for Filtering Packets .	
<i>output_control_options</i>	Save the contents of a packet to a file, capture only a number of packets, and capture a number of bytes at the beginning of packets, and so on. See pktcap-uw Options for Output Control .	

The vertical bars | represent alternative values, and the curly brackets {} used with vertical bars specify a list of choices for an argument or option.

pktcap-uw Command Syntax for Tracing Packets

Use the `pktcap-uw` utility to view the path of a packet in the network stack on an ESXi host for latency analysis.

pktcap-uw Syntax for Tracing Packets

The command of the `pktcap-uw` utility has the following syntax for tracing packets in the network stack:

```
pktcap-uw --trace filter_optionsoutput_control_options
```

Options to the pktcap-uw Utility for Tracing Packets

The `pktcap-uw` utility supports the following options when you use it to trace packets:

Table 200: pktcap-uw Options for Tracing Packets

Argument	Description
<i>filter_options</i>	Filter traced packets according to source or destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port. See pktcap-uw Options for Filtering Packets .
<i>output_control_options</i>	Save the contents of a packet to a file and trace only a number of packets. See pktcap-uw Options for Output Control .

pktcap-uw Options for Output Control

Use the options for output control of the `pktcap-uw` utility to save packet contents to a file, capture up to a certain number of bytes from each packet, and limit the number of captured packets.

pktcap-uw Options for Output Control

The options of the `pktcap-uw` utility for output control are valid when you capture and trace packets. For information about the command syntax of the `pktcap-uw` utility, see [pktcap-uw Command Syntax for Capturing Packets](#) and [pktcap-uw Command Syntax for Tracing Packets](#).

Table 201: Options for Output Control That Are Supported by the pktcap-uw Utility

Option	Description
{-o --outfile} <i>pcap_file</i>	Save captured or traced packets to a file in packet capture (PCAP) format. Use this option to examine packets in a visual analyzer tool such as Wireshark.
-P --ng	Save packet content in the PCAPNG file format. Use this option together with the -o or --outfile option.
--console	Print packet details and content to the console output. By default, the <code>pktcap-uw</code> utility shows packet information in the console output.
{-c --count} <i>number_of_packets</i>	Capture the first <i>number_of_packets</i> packets.

Option	Description
{-s --snaplen} <i>snapshot_length</i>	Capture only the first <i>snapshot_length</i> bytes from each packet. If traffic on the host is intensive, use this option to reduce the load on the CPU and storage. To limit the size of captured contents, set a value greater than 24. To capture the complete packet, set this option to 0.
-h	View help about the <code>pktcap-uw</code> utility.

The vertical bars | represent alternative values, and the curly brackets {} used with vertical bars specify a list of choices for an argument or option.

pktcap-uw Options for Filtering Packets

Narrow the range of packets that you monitor by using the `pktcap-uw` utility to apply filtering options for source and destination address, VLAN, VXLAN, and next level protocol consuming the packet payload.

Filter Options

The filter options for `pktcap-uw` are valid when you capture and trace packets. For information about the command syntax of the `pktcap-uw` utility, see [pktcap-uw Command Syntax for Capturing Packets](#) and [pktcap-uw Command Syntax for Tracing Packets](#).

Table 202: Filter Options of the `pktcap-uw` Utility

Option	Description
--srcmac <i>mac_address</i>	Capture or trace packets that have a specific source MAC address. Use colons to separate the octets in it.
--dstmac <i>mac_address</i>	Capture or trace packets that have a specific destination MAC address. Use colons to separate the octets in it.
--mac <i>mac_address</i>	Capture or trace packets that have a specific source or destination MAC address. Use colons to separate the octets in it.
--ethtype <i>0xEthertype</i>	Capture or trace packets at Layer 2 according to the next level protocol that consumes the packet payload. <i>EtherType</i> corresponds to the EtherType field in Ethernet frames. It represents the type of next level protocol that consumes the payload of the frame. For example, to monitor traffic for the Link Layer Discovery Protocol (LLDP) protocol, type --ethtype 0x88CC.
--vlan <i>VLAN_ID</i>	Capture or trace packets that belong to a VLAN.
--srcip <i>IP_address IP_address/subnet_range</i>	Capture or trace packets that have a specific source IPv4 address or subnet.
--dstip <i>IP_address IP_address/subnet_range</i>	Capture or trace packets that have a specific destination IPv4 address or subnet.
--ip <i>IP_address</i>	Capture or trace packets that have a specific source or destination IPv4 address.
--proto <i>0xIP_protocol_number</i>	Capture or trace packets at Layer 3 according to the next level protocol that consumes the payload. For example, to monitor traffic for the UDP protocol, type --proto 0x11.

Option	Description
<code>--srcport source_port</code>	Capture or trace packets according to their source TCP port.
<code>--dstport destination_port</code>	Capture or trace packets according to their destination TCP port.
<code>--tcpport TCP_port</code>	Capture or trace packets according to their source or destination TCP port.
<code>--vxlان ID</code>	Capture or trace packets that belong to a VXLAN.
<code>--rcf pcap_filter_expression</code>	Capture or trace packets using the rich common filter expression. For example, to capture all ingress and egress packets whose IP content length is larger than 1000 bytes, use the filter expression <code>--rcf "ip[2:2]>1000"</code>. To select a specific source host address and port number, use the filter expression <code>--rcf "src host 12.0.0.1 and port 5000"</code>. This example filters traffic for the host address 12.0.0.1 using port 5000. To learn more about filtering network traffic with the <code>--rcf</code> option, refer to the documentation for pcap filter expressions using command-line packet analyzers such as <code>tcpdump</code>. See pcap-filter - packet filter syntax. Note: When using the <code>--rcf</code> option, obey the following limitations. - Do not filter VLAN packets using the <code>--rcf</code> option. To trace VLAN or VXLAN use the <code>pktcap-uw--vlan</code> or <code>--vxlان</code> options. - Do not filter an IP broadcast address. - Do not use <code>--rcf</code> on ENS ports.
<code>--rcf-tcp-data tcp_packet_data_filter</code>	Capture or trace TCP data packets using the rich common filter expression. For example, to capture all HTTP/1.0 response packets with 200 OK, use the filter expression <code>--rcf-tcp-data "HTTP/1.0 200 OK"</code>. To filter on HTTP GET requests returning an index.html file, use the filter expression <code>--rcf-tcp-data "GET / index.html"</code>.

The vertical bars | represent alternative values.

Capturing Packets by Using the pktcap-uw Utility

Capture packets through the `pktcap-uw` utility in the path between a virtual switch and the physical adapters, VMkernel adapters and virtual machine adapters to troubleshoot data transfer in the network stack on an ESXi host.

Capture Packets That Arrive at a Physical Adapter

Monitor host traffic related to the external network by capturing packets at certain points in the path between a vSphere Standard Switch or vSphere Distributed Switch and a physical adapter.

You can specify a certain capture point in the data path between a virtual switch and a physical adapter, or determine a capture point by traffic direction with regard to the switch and proximity to the packet source or destination. For information about supported capture points, see [Capture Points of the pktcap-uw Utility](#).

1. Optional: Find the name of the physical adapter that you want to monitor in the host adapter list.

- In the vSphere Client, on the **Configure** tab for the host, expand **Networking** and select **Physical adapters**.

- In the ESXi Shell to the host, to view a list of the physical adapters and examine their state, run the following ESXCLI command:

```
esxcli network nic list
```

Each physical adapter is represented as `vmnicX`. `X` is the number that ESXi assigned to the physical adapter port.

2. In the ESXi Shell to the host, run the `pktcap-uw` command with the `--uplink vmnicX` argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

```
pktcap-uw
--uplink vmnicX [--capturecapture_point|--dir 0|1] [filter_options] [--outfilepcap_file_path
[--ng]] [--countnumber_of_packets]
```

where the square brackets `[]` enclose the options of the `pktcap-uw--uplink vmnicX` command and the vertical bars `|` represent alternative values.

If you run the `pktcap-uw--uplink vmnicX` command without options, you obtain the content of packets that are incoming to the standard or distributed switch in the console output at the point where they are switched.

- a) Use the `--capture` option to check packets at another capture point or the `--dir` option at another traffic direction.

pktcap-uw Command Option	Goal
<code>--capture UplinkSnd</code>	Monitor packets immediately before they enter the physical adapter device.
<code>--capture UplinkRcv</code>	Monitor packets immediately after they are received in the network stack from the physical adapter.
<code>--dir 1</code>	Monitor packets that leave the virtual switch.
<code>--dir 0</code>	Monitor packets that enter the virtual switch.

- b) Use a `filter_options` to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.
For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

- c) Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the file system.

- d) Use the `--count` option to monitor only a number of packets.

3. If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

Capture Packets That Are Received at `vmnic0` from an IP Address 192.168.25.113

To capture the first 60 packets from a source system that is assigned the IP address 192.168.25.113 at `vmnic0` and save them to a file called `vmnic0_rcv_srcip.pcap`, run the following `pktcap-uw` command:

```
pktcap-uw --uplink vmnic0 --capture UplinkRcv --srcip 192.168.25.113 --outfile vmnic0_rcv_srcip.p-
cap --count 60
```

If the contents of the packet are saved to a file, copy the file from the host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

How to Capture Packets for a VMXNET3 Virtual Machine Adapter

Monitor traffic that flows between a virtual switch and a VMXNET3 virtual machine adapter by using the `pktcap-uw` utility.

Verify that the virtual machine adapter is of type VMXNET3.

You can specify a certain capture point in the data path between a virtual switch and a virtual machine adapter. You can also determine a capture point by traffic direction with regard to the switch and proximity to the packet source or destination. For information about supported capture points, see [Capture Points of the pktcap-uw Utility](#).

1. On the host, learn the port ID of the virtual machine adapter by using the `esxstop` utility.
 - a) In the ESXi Shell to the host, to start the utility, run `esxstop`.
 - b) To switch to the network panel of the utility, press `n`.
 - c) In the USED-BY column, locate the virtual machine adapter, and write down the PORT-ID value for it.
The USED-BY field contains the name of the virtual machine and the port to which the virtual machine adapter is connected.
 - d) Press `Q` to exit `esxstop`.
2. In the ESXi Shell, run `pktcap-uw--switchportport_ID`.
`port_ID` is the ID that the `esxstop` utility displays for the virtual machine adapter in the PORT-ID column.
3. In the ESXi Shell, run the `pktcap-uw` command with the `--switchport port_ID` argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

```
pktcap-uw
--switchport
port_ID [--capturecapture_point|--dir 0|1 --stage 0|1] [filter_options] [--outfile p-
cap_file_path [--ng]] [--countnumber_of_packets]
```

where the square brackets `[]` enclose the options of the `pktcap-uw--switchportport_ID` command and the vertical bars `|` represent alternative values.

If you run the `pktcap-uw--switchportport_ID` command without options, you obtain the content of packets that are incoming to the standard or distributed switch in the console output at the point when they are switched.

- a) To check packets at another capture point or direction in the path between the guest operating system and the virtual switch, use the `--capture` option or combine the values of the `--dir` and `--stage` options.

pktcap-uw Command Options	Goal
<code>--capture VnicTx</code>	Monitor packets when they pass from the virtual machine to the switch.
<code>--capture VnicRx</code>	Monitor packets when they arrive to the virtual machine.
<code>--dir 1 --stage 0</code>	Monitor packets immediately after they leave the virtual switch.
<code>--dir 1</code>	Monitor packets immediately before they enter the virtual machine.

pktcap-uw Command Options	Goal
--dir 0 --stage 1	Monitor packets immediately after they enter the virtual switch.

- b) Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.
For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.
- c) Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.
 - To save packets to a `.pcap` file, use the `--outfile` option.
 - To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the file system.

- d) Use the `--count` option to monitor only a number of packets.
4. If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

Capture Packets That Are Received at a Virtual Machine from an IP Address 192.168.25.113

To capture the first 60 packets from a source that is assigned the IP address 192.168.25.113 when they arrive at a virtual machine adapter with port ID 33554481 and save them to a file called `vmxnet3_rcv_srcip.pcap`, run the following `pktcap-uw` command:

```
pktcap-uw --switchport 33554481 --capture VnicRx --srcip 192.168.25.113 --outfile vmxnet3_rcv_srcip.pcap --count 60
```

If the contents of the packet are saved to a file, copy the file from the host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

How to Capture Packets for a VMkernel Adapter

Monitor packets that are exchanged between a VMkernel adapter and a virtual switch by using the `pktcap-uw` utility.

You can capture packets at a certain capture point in the flow between a virtual switch and a VMkernel adapter. You can also determine a capture point by traffic direction with regard to the switch and proximity to the packet source or destination. For information about supported capture points, see [Capture Points of the pktcap-uw Utility](#).

1. Optional: Find the name of the VMkernel adapter that you want to monitor in the VMkernel adapter list.
 - In the , expand **Networking** on the **Configure** tab for the host and select **VMkernel adapters**.
 - In the ESXi Shell to the host, to view a list of the physical adapters, run the following console command:

```
esxcli network ip interface list
```

Each VMkernel adapter is represented as `vmkX`, where *X* is the sequence number that ESXi assigned to the adapter.

2. In the ESXi Shell to the host, run the `pktcap-uw` command with the `--vmk vmkX` argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

```
pktcap-uw
```



```
--vmk vmkX [--capturecapture_point|--dir 0|1 --stage 0|1] [filter_options] [--outfilep-
cap_file_path [--ng]] [--countnumber_of_packets]
```

where the square brackets [] enclose the options of the `pktcap-uw--vmk vmkX` command and the vertical bars | represent alternative values.

You can replace the `--vmk vmkX` option with `--switchportvmkernel_adapter_port_ID`, where `vmkernel_adapter_port_ID` is the PORT-ID value that the network panel of the `esxtop` utility displays for the adapter.

If you run the `pktcap-uw--vmk vmkX` command without options, you obtain the content of packets that are leaving the VMkernel adapter.

- a) To check transmitted or received packets at a specific place and direction, use the `--capture` option, or combine the values of the `--dir` and `--stage` options.

pktcap-uw Command Options	Goal
<code>--dir 1 --stage 0</code>	Monitor packets immediately after they leave the virtual switch.
<code>--dir 1</code>	Monitor packets immediately before they enter the VMkernel adapter.
<code>--dir 0 --stage 1</code>	Monitor packets immediately before they enter the virtual switch.

- b) Use a `filter_options` to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

- c) Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the file system.

- d) Use the `--count` option to monitor only a number of packets.

3. If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

If the contents of the packet are saved to a file, copy the file from the host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

How to Capture Dropped Packets

Troubleshoot lost connectivity by capturing dropped packets through the `pktcap-uw` utility.

A packet might be dropped at a point in the network stream for many reasons, for example, a firewall rule, filtering in an IOChain and DVfilter, VLAN mismatch, physical adapter malfunction, checksum failure, and so on. You can use the `pktcap-uw` utility to examine where packets are dropped and the reason for the drop.

1. In the ESXi Shell to the host, run the `pktcap-uw --capture Drop` command with options to monitor packets at a particular point, filter captured packets and save the result to a file.

```
pktcap-uw --capture Drop [filter_options] [--outfile pcap_file_path [--ng]] [--count number_of_packets]
```

where the square brackets `[]` enclose the options of the `pktcap-uw--capture Drop` command and the vertical bars `|` represent alternative values.

- a) Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

- b) Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

•

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the file system.

Note: You can see the reason and the place where a packet is dropped only when you capture packets to the console output. The `pktcap-uw` utility saves only the content of packets to a `.pcap` or `.pcapng` file.

- c) Use the `--count` option to monitor only a number of packets.

2. If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

Besides the contents of dropped packets, the output of the `pktcap-uw` utility displays the reason for the drop and the function in the network stack that handled the packet last.

If the contents of the packet are saved to a file, copy the file from the host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

How to Capture Packets at DVFilter Level

Examine how packets change when they pass through a vSphere Network Appliance (DVFilter).

DVFilters are agents that reside in the stream between a virtual machine adapter and a virtual switch. They intercept packets to protect virtual machines from security attacks and unwanted traffic.

1. Optional: To find the name of the DVFilter that you want to monitor, in the ESXi Shell, run the `summarize-dvfilter` command.

The output of the command contains the fast-path and slow-path agents of the DVFilters that are deployed on the host.

2. Run the `pktcap-uw` utility with the `--dvfilterdvfilter_name` argument and with options to monitor packets at a particular point, filter captured packets and save the result to a file.

```
pktcap-uw
--dvFilter
dvfilter_name
--capture PreDVFilter|PostDVFilter [filter_options] [--outfilepcap_file_path [--ng]] [--count-
number_of_packets]
```

where the square brackets `[]` enclose optional items of the `pktcap-uw--dvFilter vmnicX` command and the vertical bars `|` represent alternative values.

- a) Use the `--capture` option to monitor packets before or after the DVFilter intercepts them.

pktcap-uw Command Option	Goal
<code>--capture PreDVFilter</code>	Capture packets before they enter the DVFilter.

pktcap-uw Command Option	Goal
<code>--capture PostDVFilter</code>	Capture packets after they leave the DVFilter.

- b) Use a *filter_options* to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.

For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.

- c) Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

- To save packets to a `.pcap` file, use the `--outfile` option.
- To save packets to a `.pcapng` file, use the `--ng` and `--outfile` options.

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the file system.

- d) Use the `--count` option to monitor only a number of packets.

3. If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

If the contents of the packet are saved to a file, copy the file from the host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

How to use the Capture Points of the `pktcap-uw` Utility

Learn how to use the capture points of the `pktcap-uw` utility to monitor packets when a function handles them at a specific place in the network stack on a host.

Overview of Capture Points

A capture point in the `pktcap-uw` utility represents a place in the path between a virtual switch on one side and a physical adapter, VMkernel adapter or a virtual machine adapter on the other.

You can use certain capture points in combination with an adapter option. For example, you use the `UplinkRcv` point when you capture uplink traffic. You can address other points standalone. For example, use the `Drop` point to inspect all dropped packets.

Note: Certain capture points of the `pktcap-uw` utility are designed for VMware internal use only and you should use them only under the supervision of VMware Technical Support. These capture points are not described in the *vSphere Networking* guide.

Option for Using Capture Points in the `pktcap-uw` Utility

To examine a packet state or content at a capture point, add the `--capture capture_point` option to the `pktcap-uw` utility.

Auto-Selecting a Capture Point

For traffic that is related to a physical, VMkernel or VMXNET3 adapter, by combining the `--dir` and `--stage` options you can auto-select and switch between capture points to examine how a packet changes before and after a point.

Capture Points of the `pktcap-uw` Utility

The `pktcap-uw` utility supports capture points that can be used only when you monitor uplink, VMkernel or virtual machine traffic, and capture points that represent special places in the stack that are not related to the adapter type.

Capture Points That Are Relevant to Physical Adapter Traffic

The `pktcap-uw--uplink vmnicX` command supports capture points for functions that handle traffic at a specific place and direction in the path between the physical adapter and the virtual switch.

Capture Point	Description
UplinkRcv	The function that receives packets from the physical adapter.
UplinkSnd	The function that sends packets to the physical adapter.
PortInput	The function that passes a list of packets from UplinkRcv to a port on the virtual switch.
PortOutput	The function that passes a list of packets from a port on the virtual switch to the UplinkSnd point.

Capture Points That Are Relevant to Virtual Machine Traffic

The `pktcap-uw--switchportvmxnet3_port_ID` command supports capture points for functions that handle traffic packets at a specific place and direction in the path between a VMXNET3 adapter and a virtual switch.

Capture Point	Description
VnicRx	The function in the virtual machine NIC backend that receives packets from the virtual switch.
VnicTx	The function in the virtual machine NIC backend that sends packets from the virtual machine to the virtual switch.
PortOutput	The function that passes a list of packets from a port on the virtual switch to Vmxnet3Rx.
PortInput	The function that passes a list of packets from Vmxnet3Tx to a port on the virtual switch. Default capture point for traffic related to a VMXNET3 adapter.

Capture Points That Are Relevant to VMkernel Adapter Traffic

The `pktcap-uw--vmk vmkX` and `pktcap-uw--switchportvmkernel_adapter_port_ID` commands support capture points that represent functions at a specific place and direction in the path between a VMkernel adapter and a virtual switch.

Capture Point	Description
PortOutput	The function that passes a list of packets from a port on the virtual switch to the VMkernel adapter.
PortInput	The function that passes a list of packets from the VMkernel adapter to a port on the virtual switch. Default capture point for traffic related to a VMkernel adapter.

Capture Points That Are Relevant to Distributed Virtual Filters

The `pktcap-uw--dvfilterdivfilter_name` command requires a capture point that indicates whether to capture packets when they enter the DVFilter or when they leave it.

Capture Point	Description
PreDVFilter	The point before a DVFilter intercepts a packet.
PostDVFilter	The point after a DVFilter intercepts a packet.

Standalone Capture Points

Certain capture points are mapped directly to the network stack rather than to a physical, VMkernel or VMXNET3 adapter.

Capture Point	Description
Drop	Captures dropped packets and shows the place where drops occur.
TcpipDispatch	Capture packets at the function that dispatches traffic to the TCP/IP stack of the VMkernel from the virtual switch, and the reverse.
PktFree	Capture packets right before they are released.
VdrRxLeaf	Capture packets at the receive leaf I/O chain of a dynamic router in VMware NSX. Use this capture point together with the --lifID option.
VdrRxTerminal	Capture packets at the receive terminal I/O chain of a dynamic router in VMware NSX. Use this capture point together with the --lifID option.
VdrTxLeaf	Capture packets at the transmit leaf I/O chain of a dynamic router in VMware NSX. Use this capture point together with the --lifID option.
VdrTxTerminal	Capture packets at the transmit terminal I/O chain of a dynamic router in VMware NSX. Use this capture point together with the --lifID option.

For information about dynamic routers, see the *VMware NSX* documentation.

How to List the Capture Points of the pktcap-uw Utility

View all capture points of the `pktcap-uw` utility to find the name of capture point for monitoring traffic at a certain place in the network stack on the ESXi host.

For information about the capture points of the `pktcap-uw` utility, see [Capture Points of the pktcap-uw Utility](#).

In the ESXi Shell to the host, run the `pktcap-uw -A` command to view all capture points that the `pktcap-uw` utility supports.

How to Trace Packets using the pktcap-uw Utility

Use the `pktcap-uw` utility to trace the path that packets traverse in the network stack for latency analysis and for locating the point where a packet is corrupted or dropped.

The `pktcap-uw` utility shows the path of packets together with timestamps that note the time when a packet is handled by a networking function on ESXi. The utility reports the path of a packet immediately before it is released from the stack.

To view the full path information for a packet, you must print the result from the `pktcap-uw` utility in the console output or save it to a PCAPNG file.

1. In the ESXi Shell to the host, run the `pktcap-uw--trace` command with options to filter traced packets, save the result to a file and limit the number of traced packets.

```
pktcap-uw
--trace [filter_options] [--outfilepcap_file_path [--ng]] [--countnumber_of_packets]
```

where the square brackets `[]` enclose optional items of the `pktcap-uw --trace` command and the vertical bars `|` represent alternative values.

- a) Use a `filter_options` to filter packets according to source and destination address, VLAN ID, VXLAN ID, Layer 3 protocol, and TCP port.
For example, to monitor packets from a source system that has IP address 192.168.25.113, use the `--srcip 192.168.25.113` filter option.
- b) Use options to save the contents of each packet or the contents of a limited number of packets to a `.pcap` or `.pcapng` file.

-

You can open the file in a network analyzer tool such as Wireshark.

By default, the `pktcap-uw` utility saves the packet files to the root folder of the file system.

Note: A `.pcap` file contains only the contents of traced packets. To collect packet paths besides packet content, save the output to a `.pcapng` file.

- c) Use the `--count` option to monitor only a number of packets.
2. If you have not limited the number of packets by using the `--count` option, press Ctrl+C to stop capturing or tracing packets.

If the contents of the packet are saved to a file, copy the file from the host to the system that runs a graphical analyzer tool, such as Wireshark, and open it in the tool to examine the packet details.

Configure the NetFlow Settings of a vSphere Distributed Switch

Analyze virtual machine IP traffic that flows through a vSphere Distributed Switch by sending reports to a NetFlow collector.

vSphere Distributed Switch supports IPFIX (NetFlow version 10).

Note: To configure IPFIX on a vSphere Distributed Switch backed by ESXi on DPU, you must create vmknics on ops TCP/IP stack. Else, the flow information is not exported to collector.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. From the **Actions** menu, select **Settings > Edit Netflow**.
3. Type the **Collector IP address** and **Collector port** of the NetFlow collector.

You can contact the NetFlow collector by IPv4 or IPv6 address.

4. Set an **Observation Domain ID** that identifies the information related to the switch.
5. To see the information from the distributed switch in the NetFlow collector under a single network device instead of under a separate device for each host on the switch, type an IPv4 address in the **Switch IP address** text box.
6. Optional: In the **Active flow export timeout** and **Idle flow export timeout** text boxes, set the time, in seconds, to wait before sending information after the flow is initiated.
7. Optional: To change the portion of data that the switch collects, configure **Sampling Rate**.
The sampling rate represents the number of packets that NetFlow drops after every collected packet. A sampling rate of x instructs NetFlow to drop packets in a *collected packets:dropped packets* ratio 1: x . If the rate is 0, NetFlow samples every packet, that is, collect one packet and drop none. If the rate is 1, NetFlow samples a packet and drops the next one, and so on.
8. Optional: To collect data on network activity between virtual machines on the same host, enable **Process internal flows only**.
Collect internal flows only if NetFlow is enabled on the physical network device to avoid sending duplicate information from the distributed switch and the physical network device.
9. Click **OK**.

Enable NetFlow reporting for traffic from virtual machines connected to a distributed port group or a port. See [Manage NetFlow Monitoring on a Distributed Port Group or Distributed Port](#).

What is Port Mirroring

Port mirroring allows you to mirror a distributed port's traffic to other distributed ports or specific physical switch ports.

Port mirroring is used on a switch to send a copy of packets seen on one switch port (or an entire VLAN) to a monitoring connection on another switch port. Port mirroring is used to analyze and debug data or diagnose errors on a network.

What is Port Mirroring Interoperability

There are some interoperability issues to consider when using vSphere port mirroring with other features of vSphere.

vMotion

vMotion functions differently depending on which vSphere port mirroring session type you select. During vMotion, a mirroring path could be temporarily invalid, but it is restored when vMotion completes.

Table 203: vMotion Interoperability with port mirroring

Port mirroring session type	Source and destination	Interoperable with vMotion	Functionality
Distributed Port Mirroring	Non-uplink distributed port source and destination	Yes	Port mirroring between distributed ports can only be local. If the source and destination are on different hosts due to vMotion, mirroring between them will not work. However, if the source and destination move to the same host, port mirroring works.
Remote Mirroring Source	Non-uplink distributed port source	Yes	When a source distributed port is moved from host A to host B, the original mirroring path from the source port to A's uplink is removed on A, and a new

Port mirroring session type	Source and destination	Interoperable with vMotion	Functionality
			mirroring path from the source port to B's uplink is created on B. Which uplink is used is determined by the uplink name specified in session.
	Uplink port destinations	No	Uplinks can not be moved by vMotion.
Remote Mirroring Destination	VLAN source	No	
	Non-uplink distributed port destination	Yes	When a destination distributed port is moved from host A to host B, all original mirroring paths from source VLANs to the destination port are moved from A to B.
Encapsulated Remote Mirroring (L3) Source	Non-uplink distributed port source	Yes	When a source distributed port is moved from host A to host B, all original mirroring paths from the source port to destination IPs are moved from A to B.
	IP destination	No	
Distributed Port Mirroring (legacy)	IP source	No	
	Non-uplink distributed port destination	No	When a destination distributed port is moved from host A to host B, all original mirroring paths from source IPs to the destination port are invalid because the port mirroring session source still sees the destination on A.

TSO and LRO

TCP Segmentation Offload (TSO) and large receive offload (LRO) might cause the number of mirroring packets to not equal to the number of mirrored packets.

When TSO is enabled on a vNIC, the vNIC might send a large packet to a distributed switch. When LRO is enabled on a vNIC, small packets sent to it might be merged into a large packet.

Source	Destination	Description
TSO	LRO	Packets from the source vNIC might be large packets, and whether they are split is determined by whether their sizes are larger than the destination vNIC LRO limitation.
TSO	Any destination	Packets from the source vNIC might be large packets, and they are split to standard packets at the destination vNIC.
Any source	LRO	Packets from the source vNIC are standard packets, and they might be merged into larger packets at the destination vNIC.

Create a Port Mirroring Session

Create a port mirroring session by using the vSphere Client to mirror vSphere Distributed Switch traffic to ports, uplinks, and remote IP addresses.

Verify that the vSphere Distributed Switch is version 5.0.0 and later.

Types of Port Mirroring Session

To begin a port mirroring session, you must specify the type of port mirroring session.

1. Browse to a distributed switch in the vSphere Client navigator.
2. Click the **Configure** tab and expand **Settings**.
3. Select the **Port mirroring** option and click **New**.
4. Select the session type for the port mirroring session.

Note:

When the session type for the port mirroring session is set to **Distributed Port Mirroring**, **Remote Mirror Source**, or **Remote Mirror Destination**, the switching VLAN tag will not be in the mirrored packets.

Option	Description
Distributed Port Mirroring	Mirror packets from a number of distributed ports to other distributed ports on the same host. If the source and the destination are on different hosts, this session type does not function.
Remote Mirroring Source	Mirror packets from a number of distributed ports to specific uplink ports on the corresponding host.
Remote Mirroring Destination	Mirror packets from a number of VLANs to distributed ports.
Encapsulated Remote Mirroring (L3) Source	Mirror packets from a number of distributed ports to the IP addresses of a remote agent. The virtual machine's traffic is mirrored to a physical or virtual destination through an IP tunnel.

5. Click **Next**.

Port Mirroring Name and Session Details

To continue creating a port mirroring session, specify the name, description, and session details for the new port mirroring session.

1. Set the session properties. Different options are available for configuration depending on which session type you selected.

Option	Description
Name	You can enter a unique name for the port mirroring session, or accept the automatically generated session name.
Status	Use the drop down menu to enable or disable the session.
Session type	Displays the type of session you selected.
Encapsulation Type	Select GRE , ERSPAN TWO , or ERSPAN THREE . Note: This option is enabled when session type is set to Encapsulated Remote Mirroring (L3) Source .
Session ID	Specify an ERSPAN ID if encapsulation type is set to ERSPAN TWO or ERSPAN THREE .

Option	Description
	Note: This option is enabled when session type is set to Encapsulated Remote Mirroring (L3) Source.
Encapsulation VLAN ID	VLAN ID that encapsulates all frames at the destination ports. Note: If the original frames have a VLAN and Preserve original VLAN is not selected, the encapsulation VLAN replaces the original VLAN. This option is enabled when session type is set to Remote Mirroring Source.
Preserve original VLAN	Select Preserve original VLAN to keep the original VLAN in an inner tag so mirrored frames are double encapsulated. This option is enabled when session type is set to Remote Mirroring Source.
Normal I/O on destination ports	Use the drop-down menu to allow or disallow normal I/O on destination ports. This property is only available for uplink and distributed port destinations. If you disallow this option, mirrored traffic will be allowed on outgoing destination ports, but incoming traffic will not be allowed.
TCP/IP Stack	Use the drop down menu to select the type of TCP/IP stack. • Default: The default TCP/IP stack. • Mirror: Using mirror stack instead of default TCP/IP netstack can separate mirror traffic from management traffic. Without mirror stack, mirror traffic is bound to default TCP/IP stack. The management traffic also uses the default TCP/IP stack. If the mirror traffic is large, then it affects the management traffic. If you want to separate the mirror traffic from the default TCP/IP stack, then you can have a dedicated mirror netstack on the ESXi. You can enable this dedicated netstack while configuring Encapsulated Remote Mirroring sessions. Note: If you want to configure ERSPAN on a vSphere Distributed backed by ESXi on DPU, create a vmknic on mirror TCP/IP stack. Note: This option is enabled when session type is set to Encapsulated Remote Mirroring (L3) Source.
Mirrored packet length (Bytes)	Use the check box to enable mirrored packet length in bytes. This puts a limit on the size of mirrored frames. If this option is selected, all mirrored frames are truncated to the specified length.
Sampling rate	Select the rate at which packets are sampled. This is enabled by default for all port mirroring sessions except legacy sessions. Note: Sampling rate is not enabled if NSX transport node and Enhanced Network Stack (ENS) is enabled.
Description	You have the option to enter a description of the port mirroring session configuration.

2. Click **Next**.

Port Mirroring Sources

To continue creating a port mirroring session, select sources and traffic direction for the new port mirroring session.

You can create a port mirroring session without setting the source and destinations. When a source and destination are not set, a port mirroring session is created without the mirroring path. This allows you to create a port mirroring session with the correct properties set. Once the properties are set, you can edit the port mirroring session to add the source and destination information.

Note: Consider the following limitations when selecting port mirroring sources.

- A source mirror port cannot be used in more than one mirror session.
- A port cannot be used as a mirror source and mirror destination in the same or different mirror sessions at the same time.

1. Select the source of the traffic to be mirrored and the traffic direction.

Depending on the type of port mirroring session you selected, different options are available for configuration.

Option	Description
Add existing ports from a list	Click Select distributed ports . A dialog box displays a list of existing ports. Select the check box next to the distributed port and click OK . You can choose more than one distributed port.
Add existing ports by port number	Click Add distributed ports , enter the port number and click OK .
Set the traffic direction	After adding ports, select the port in the list and click the ingress, egress, or ingress/egress button. Your choice appears in the Traffic Direction column.
Specify the source VLAN	If you selected a Remote Mirroring Destination sessions type, you must specify the source VLAN. Click Add to add a VLAN ID. Edit the ID by using the up and down arrows, or by clicking in the field and entering the VLAN ID manually.

2. Click **Next**.

Port Mirroring Destinations

To complete the creation of a port mirroring session, select ports or uplinks as destinations for the port mirroring session.

You can create a port mirroring session without setting the source and destinations. When a source and destination are not set, a port mirroring session is created without the mirroring path. This allows you to create a port mirroring session with the correct properties set. Once the properties are set, you can edit the port mirroring session to add the source and destination information.

Port mirroring is checked against the VLAN forwarding policy. If the VLAN of the original frames is not equal to or trunked by the destination port, the frames are not mirrored.

1. Select the destination for the port mirroring session.

Depending on which type of session you chose, different options are available.

Option	Description
Select a destination distributed port	Click Select distributed ports to select ports from a list, or click Add distributed ports to add ports by port number. You can add more than one distributed port.

Option	Description
Select an uplink	Select an available uplink from the list and click Add to add the uplink to the port mirroring session. You can select more than one uplink.
Select ports or uplinks	Click Select distributed ports to select ports from a list, or click Add distributed ports to add ports by port number. You can add more than one distributed port. Click Add uplinks to add uplinks as the destination. Select uplinks from the list and click OK .
Specify IP address	Click Add . A new list entry is created. Select the entry and either click Edit to enter the IP address, or click directly in the IP Address field and type the IP address. A warning appears if the IP address is invalid.

- Click **Next**.
- Review the information that you entered for the port mirroring session on the **Ready to complete** page.
- Optional: Use the **Back** button to edit the information.
- Click **Finish**.

The new port mirroring session appears in the Port Mirroring section of the **Settings** tab.

View Port Mirroring Session Details

View port mirroring session details, including status, sources, and destinations.

- On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
- On the **Configure** tab, expand **Settings** and click **Port mirroring**.
- Select a port mirroring session from the list to display more detailed information at the bottom of the screen. Use the tabs to review configuration details.
- Optional: Click **New** to add a new port mirroring session.
- Optional: Click **Edit** to edit the details for the selected port mirroring session.
- Optional: Click **Remove** to delete the selected port mirroring session.

Edit Port Mirroring Session Details, Sources, and Destinations

Edit the details of a port mirroring session, including name, description, status, sources, and destinations.

- On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
- On the **Configure** tab, expand **Settings** and click **Port mirroring**.
- Select a port mirroring session from the list and click **Edit**.
- On the **Properties** page, edit the session properties.

Depending on the type of port mirroring session being edited, different options are available for configuration.

Option	Description
Name	You can enter a unique name for the port mirroring session, or accept the automatically generated session name.
Status	Use the drop-down menu to enable or disable the session.

Option	Description
Normal I/O on destination ports	Use the drop-down menu to allow or disallow normal I/O on destination ports. This property is only available for uplink and distributed port destinations. If you do not select this option, mirrored traffic will be allowed out on destination ports, but no traffic will be allowed in.
Sampling rate	Select the rate at which packets are sampled. This is enabled by default for all port mirroring sessions except legacy sessions. Note: Sampling rate is not enabled if NSX transport node and enhanced network stack is enabled.
Mirrored packet length (Bytes)	Use the check box to enable mirrored packet length in bytes. This puts a limit on the size of mirrored frames. If this option is selected, all mirrored frames are truncated to the specified length.
Description	You have the option to enter a description of the port mirroring session configuration.

5. On the **Sources** page, edit sources for the port mirroring session.

Depending on the type of port mirroring session being edited, different options are available for configuration.

Option	Description
Add existing ports from a list	Click the Select distributed ports to add to this port mirroring session. button. A dialog opens with a list of existing ports. Select the check box next to the distributed port and click OK . You can choose more than one distributed port.
Set the traffic direction	After adding ports, select the port in the list and click the ingress, egress, or ingress/egress button. Your choice is displayed in the Traffic Direction column.

6. In the **Destinations** section, edit the destinations for the port mirroring session.

Depending on the type of port mirroring session being edited, different options are available for configuration.

Option	Description
Select a destination distributed port	Click the Select distributed portsto add to this port mirroring session. button to select ports from a list. You can add more than one distributed port.

7. Click **OK**.

vSphere Distributed Switch Health Check

The health check support helps you identify and troubleshoot configuration errors in a vSphere Distributed Switch.

Use vSphere Distributed Switch health check to examine certain settings on the distributed and physical switches to identify common errors in your environment's networking configuration. The default interval between two health checks is 1 minute.

Important: Use health check to troubleshoot network problems, and then deactivate it after you identify and resolve the problem. After you deactivate vSphere Distributed Switch health check, the generated MAC addresses age out of your physical network environment according to your network policy. For more information, see Knowledge Base article [KB 2034795](#).

Configuration Error	Health Check	Required Configuration on the Distributed Switch
The VLAN trunk ranges configured on the distributed switch do not match the trunk ranges on the physical switch.	Checks whether the VLAN settings on the distributed switch match the trunk port configuration on the connected physical switch ports.	At least two active physical NICs
The MTU settings on the physical network adapters, distributed switch, and physical switch ports do not match.	Checks whether the physical access switch port MTU jumbo frame setting based on per VLAN matches the vSphere distributed switch MTU setting.	At least two active physical NICs
The teaming policy configured on the port groups does not match the policy on the physical switch port-channel.	Checks whether the connected access ports of the physical switch that participate in an EtherChannel are paired with distributed ports whose teaming policy is set to IP hash.	At least two active physical NICs and two hosts

Health check is limited to only the access switch port to which the distributed switch uplink connects.

Manage vSphere Distributed Switch Health Check

Health check monitors for changes in vSphere Distributed Switch configurations. You must enable vSphere Distributed Switch health check to perform checks on distributed switch configurations.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Select the **Configure** tab and expand Setting.
3. Select **Health Check** and click the **Edit** button.
4. Use the drop-down menus to enable or disable health check options.

Option	Description
VLAN and MTU	Reports the status of distributed uplink ports and VLAN ranges.
Teaming and Failover	Checks for any configuration mismatch between the ESXi host and the physical switch used in the teaming policy.

5. Click **OK**.

When you change the configuration of a vSphere Distributed Switch, you can view information about the change in the **Monitor** tab in the vSphere Client. See [View vSphere Distributed Switch Health Status](#).

View vSphere Distributed Switch Health Status

Once you have enabled health check on a vSphere Distributed Switch, you can view the network health status of the hosts connected in the vSphere Client .

Verify that health check for VLAN and MTU, and for teaming policy is enabled on the vSphere Distributed Switch. See [Manage vSphere Distributed Switch Health Check](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. On the **Monitor** tab, click **Health**.
3. In the Host member health status section, examine the overall, VLAN, MTU and teaming health of the hosts connected to the switch.

Switch Discovery Protocol

Switch discovery protocols help vSphere administrators to determine which port of the physical switch is connected to a vSphere standard switch or vSphere distributed switch.

vSphere 5.0 and later supports Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP). CDP is available for vSphere standard switches and vSphere distributed switches connected to Cisco physical switches. LLDP is available for vSphere distributed switches version 5.0.0 and later.

When CDP or LLDP is enabled for a particular vSphere distributed switch or vSphere standard switch, you can view properties of the peer physical switch such as device ID, software version, and timeout from the vSphere Client.

Enable Cisco Discovery Protocol on a vSphere Distributed Switch

Cisco Discovery Protocol (CDP) enables vSphere administrators to determine which port of a physical Cisco switch connects to a vSphere Standard Switch or vSphere Distributed Switch. When CDP is enabled for a vSphere Distributed Switch, you can view the properties of the Cisco switch such as device ID, software version, and timeout.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. From the **Actions** menu, select **Settings > Edit Settings**.
3. In the Edit Settings dialog box, click **Advanced**.
4. In the Discovery Protocol section, select **Cisco Discovery Protocol** from the **Type** drop-down menu.
5. From the **Operation** drop-down menu, select the operational mode of the ESXi hosts connected to the switch.

Option	Description
Listen	ESXi detects and displays information about the associated Cisco switch port, but information about the vSphere Distributed Switch is not available to the Cisco switch administrator.
Advertise	ESXi makes information about the vSphere Distributed Switch available to the Cisco switch administrator, but does not detect and display information about the Cisco switch.
Both	ESXi detects and displays information about the associated Cisco switch and makes information about the vSphere Distributed Switch available to the Cisco switch administrator.

- Click **OK**.

Enable Link Layer Discovery Protocol on a vSphere Distributed Switch

With Link Layer Discovery Protocol (LLDP), vSphere administrators can determine which physical switch port connects to a given vSphere Distributed Switch. When LLDP is enabled for a particular distributed switch, you can view properties of the physical switch (such as chassis ID, system name and description, and device capabilities) from the .

- On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
- From the **Actions** menu, select **Settings > Edit Settings**.
- In the Edit Settings dialog box, click **Advanced**.
- In the Discovery Protocol section, select **Link Layer Discovery Protocol** from the **Type** drop-down menu.
- From the **Operation** drop-down menu, select the operational mode of the ESXi hosts connected to the switch.

Operation	Description
Listen	ESXi detects and displays information about the associated physical switch port, but information about the vSphere Distributed Switch is not available to the switch administrator.
Advertise	ESXi makes information about the vSphere Distributed Switch available to the switch administrator, but does not detect and display information about the physical switch.
Both	ESXi detects and displays information about the associated physical switch and makes information about the vSphere Distributed Switch available to the switch administrator.

- Click **OK**.

View Switch Information

When Cisco Discovery Protocol (CDP) or Link Layer Discovery Protocol (LLDP) is enabled on the distributed switch and the hosts connected to the switch are in Listen or Both operational mode, you can view physical switch information from the vSphere Client.

- In the , navigate to the host.
- On the **Configure** tab, expand **Networking** and click **Physical adapters**.
- Select a physical adapter from the list to view its detailed information.

According to the enabled switch discovery protocol, the properties of the switch appear under the **CDP** or **LLDP** tab. If the information is available in the network, under Peer device capability you can examine the system capabilities of the switch.

View the Topology of an NSX Virtual Distributed Switch

You can examine the structure and components of an NSX Virtual Distributed Switch (N-VDS) by viewing its topology diagram.

The topology diagram of an N-VDS provides a visual representation of the adapters and port groups connected to the switch.

From the diagram you can view the settings of a selected port group and of a selected adapter.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **Networking** and select **Virtual Switches**.
3. Select the N-VDS from the list.

The diagram appears under the list of virtual switches on the host.

You can use the topology diagram to examine whether a virtual machine or VMkernel adapter is connected to the external network and to identify the physical adapter that carries the data.

Configure Protocol Profiles for Virtual Machine Networking

A network protocol profile contains a pool of IPv4 and IPv6 addresses that vCenter Server assigns to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

Network protocol profiles also contain settings for the IP subnet, DNS, and HTTP proxy server.

To configure the networking settings of virtual machines by using from network protocol profiles, perform the following operations:

- Create network profiles at the level of a data center or a vSphere distributed switch.
- Associate a protocol profile with the port group of a vApp virtual machine.
- Enable the transient or static IP allocation policy from the settings of the vApp or from the vApp options of a virtual machine.

Note: If you move a vApp or a virtual machine that retrieves its network settings from a protocol profile to another data center, to power it on you must assign a protocol profile to the connected port group on the destination data center.

Add a Network Protocol Profile

A network protocol profile contains a pool of IPv4 and IPv6 addresses that vCenter Server assigns to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

You can configure network protocol profile ranges for IPv4, IPv6, or both. vCenter Server uses these ranges to dynamically allocate IP addresses to the virtual machines within a vApp, when the vApp uses transient IP allocation policy.

Network protocol profiles also contain settings for the IP subnet, the DNS, and HTTP proxy servers.

Note: If you move a vApp or a virtual machine that retrieves its network settings from a protocol profile to another data center, to power on the vApp or virtual machine you must assign a protocol profile to the connected port group on the destination data center.

1. Navigate to a data center that is associated with a vApp.
2. On the **Configure** tab, select **More > Network Protocol Profiles**.
Existing network protocol profiles are listed.
3. Click the **Add** button.
The **Add Network Protocol Profile** wizard opens.
4. On the **Name and network** page, enter the name of the network protocol profile and select the networks that use this profile. Click **Next**.
A network can be associated with one network protocol profile at a time.

5. On the **IPv4** page, configure the relevant IPv4 settings.
 - a) In the **Subnet** and the **Gateway** text boxes, enter the IP subnet and gateway.
 - b) To indicate that the DHCP server is available on the network, select the **DHCP Present** radio button.
 - c) In the **DNS server addresses** text box, enter the DNS server information.
 - d) To specify an IP pool range, enable the **IP Pool** option.
 - e) If you enable IP pools, enter a comma-separated list of host address ranges in the **IP pool range** text box.
 A range consists of an IP address, a pound sign (#), and a number indicating the length of the range.
 For example, `10.20.60.4#10, 10.20.61.0#2` indicates that the IPv4 addresses can range from 10.20.60.4 to 10.20.60.13 and 10.20.61.0 to 10.20.61.1.
 The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** text box cannot include the gateway address.
 - f) Click **Next**.
6. On the **IPv6** page, configure the relevant IPv6 settings.
 - a) In the **Subnet** and the **Gateway** text boxes, enter the IP subnet and gateway.
 - b) Select the **DHCP Present** radio button to indicate that the DHCP server is available on this network.
 - c) In the **DNS server addresses**, enter the DNS server information.
 - d) Enable the **IP Pool** option to specify an IP pool range.
 - e) If you enable IP pools, enter a comma-separated list of host address ranges in the **IP pool range** text box.
 A range consists of an IP address, a pound sign (#), and a number indicating the length of the range.
 For example, assume that you specify the following IP pool range: `fe80:0:0:0:2bff:fe59:5a:2b#10, fe80:0:0:0:2bff:fe59:5f:b1#2`. Then the addresses are in this range:
`fe80:0:0:0:2bff:fe59:5a:2b - fe80:0:0:0:2bff:fe59:5a:34`
 and
`fe80:0:0:0:2bff:fe59:5f:b1 - fe80:0:0:0:2bff:fe59:5f:b2`.
 The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** text box cannot include the gateway address.
 - f) Click **Next**.
7. On the **Other network configurations** page, specify additional network configurations.
 - a) Enter the DNS domain.
 - b) Enter the host prefix.
 - c) Enter the DNS search path.
 The search paths are specified as a list of DNS domains separated by commas, semi-colons, or spaces.
 - d) Enter the server name and port number for the proxy server.
 The server name must include a colon and a port number. For example, `web-proxy:3912` is a valid proxy server.
 - e) Click **Next**.
8. On the **Name and Network Assignment** page, review the settings and click **Finish**.

Related Links

[Associate a Port Group with a Network Protocol Profile on page 2388](#)

To apply the range of IP addresses from a network protocol profile to a virtual machine that is a part of a vApp or has vApp functionality enabled, associate the profile with a port group that controls the networking of the virtual machine.

[Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine or vApp on page 2389](#)

After you associate a network protocol profile with a port group of a standard switch or a distributed switch, you can use the profile to dynamically allocate IP addresses to a virtual machine that is within a vApp.

Select the Network Protocol Profile Name and Network

Name the network protocol profile and select the network that should use it.

1. Type the name of the network protocol profile.
2. Select the networks that use this network protocol profile.
A network can be associated with one network protocol profile at a time.
3. Click **Next**.

Specify Network Protocol Profile IPv4 Configuration

A network protocol profile contains a pool of IPv4 and IPv6 addresses for use by vApps. When you create a network protocol profile, you set up its IPv4 configuration.

You can configure network protocol profile ranges for IPv4, IPv6, or both. vCenter Server uses these ranges to dynamically allocate IP addresses to virtual machines when a vApp is set up to use transient IP allocation.

1. Enter the **IP Subnet** and **Gateway** in their respective fields.
2. Select **DHCP Present** to indicate that the DHCP server is available on this network.
3. Enter the DNS server information.
Specify the servers by IP addresses separated by a comma, semicolon, or space.
4. Select the **Enable IP Pool** check box to specify an IP pool range.
5. If you enable IP Pools, enter a comma-separated list of host address ranges in the **IP pool range** field.
A range consists of an IP address, a pound sign (#), and a number indicating the length of the range.
The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** field cannot include the gateway address.
For example, 10.20.60.4#10, 10.20.61.0#2 indicates that the IPv4 addresses can range from 10.20.60.4 to 10.20.60.13 and 10.20.61.0 to 10.20.61.1.
6. Click **Next**.

Specify Network Protocol Profile IPv6 Configuration

A network protocol profile contains a pool of IPv4 and IPv6 addresses for use by vApps. When you create a network protocol profile, you set up its IPv6 configuration.

You can configure network protocol profile ranges for IPv4, IPv6, or both. vCenter Server uses these ranges to dynamically allocate IP addresses to virtual machines when a vApp is set up to use transient IP allocation.

1. Enter the **IP Subnet** and **Gateway** in their respective fields.
2. Select **DHCP Present** to indicate that the DHCP server is available on this network.
3. Enter the DNS server information.
Specify the servers by IP addresses separated by a comma, semicolon, or space.

4. Select the **Enable IP Pool** check box to specify an IP pool range.
5. If you enable IP Pools, enter a comma-separated list of host address ranges in the **IP pool range** field.
A range consists of an IP address, a pound sign (#), and a number indicating the length of the range. For example, assume that you specify the following IP pool range:
fe80:0:0:0:2bff:fe59:5a:2b#10,fe80:0:0:0:2bff:fe59:5f:b1#2
Then the addresses are in this range:
fe80:0:0:0:2bff:fe59:5a:2b - fe80:0:0:0:2bff:fe59:5a:34
and
fe80:0:0:0:2bff:fe59:5f:b1 - fe80:0:0:0:2bff:fe59:5f:b2
The gateway and the ranges must be within the subnet. The ranges that you enter in the **IP pool range** field cannot include the gateway address.
6. Click **Next**.

Specify Network Protocol Profile DNS and Other Configuration

When you create a network protocol profile, you can specify the DNS domain, DNS search path, a host prefix, and HTTP proxy.

1. Enter the DNS domain.
2. Enter the host prefix.
3. Enter the DNS search path.
The search paths are specified as a list of DNS domains separated by commas, semi-colons, or spaces.
4. Enter the server name and port number for the proxy server.
The server name can optionally include a colon and a port number.
For example, web-proxy:3912 is a valid proxy server.
5. Click **Next**.

Complete the Network Protocol Profile Creation

Review the settings and click **Finish** to complete adding the network protocol profile.

Associate a Port Group with a Network Protocol Profile

To apply the range of IP addresses from a network protocol profile to a virtual machine that is a part of a vApp or has vApp functionality enabled, associate the profile with a port group that controls the networking of the virtual machine.

1. Navigate to a data center that is associated with a vApp.
2. On the **Configure** tab, select **More > Network Protocol Profiles**.
Existing network protocol profiles are listed.
3. Select a network profile from the list and click **Assign**.
The **Assign Networks** dialog box opens.

4. Select a port group or a network to assign to the network profile.
 - On the **Distributed portgroups** tab, you see a list of the distributed port groups.
 - On the **Networks** tab, you see a list of the port groups of standard switches.

You can select multiple port groups before you close the dialog box.

5. Click **Save**.

Related Links

[Add a Network Protocol Profile on page 2385](#)

A network protocol profile contains a pool of IPv4 and IPv6 addresses that vCenter Server assigns to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

[Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine or vApp on page 2389](#)

After you associate a network protocol profile with a port group of a standard switch or a distributed switch, you can use the profile to dynamically allocate IP addresses to a virtual machine that is within a vApp.

Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine or vApp

After you associate a network protocol profile with a port group of a standard switch or a distributed switch, you can use the profile to dynamically allocate IP addresses to a virtual machine that is within a vApp.

Verify that the virtual machine is connected to a port group that is associated with the network protocol profile.

Select your task.

Option	Description
Use a Network Protocol Profile to Allocate IP Addresses to a Virtual Machine	1. Navigate to a virtual machine in the vCenter Server inventory. 2. On the Configure tab, expand Settings and select vApp Options. 3. Click the Edit button. The Edit vApp options dialog box opens. 4. If vApp options are not enabled, select the Enable vApp options check box. 5. Click the IP Allocation tab. 6. In the Authoring section, select OVF environment as an IP allocation scheme. 7. In the Deployment section, set the IP allocation to Transient - IP Pool or Static - IP Pool. 8. Click OK.
Use a Network Protocol Profile to Allocate IP Addresses to a vApp	1. Navigate to a vApp in the vCenter Server inventory 2. Right-click the vApp and select Edit Settings. The Edit vApp dialog box opens. 3. Click the IP Allocation tab. 4. In the Authoring section, select OVF environment as an IP allocation scheme. 5. In the Deployment section, set the IP allocation to Transient - IP Pool or Static - IP Pool. 6. Click OK.

Both the **Static - IP Pool** and **Transient - IP Pool** options allocate an IP address from the range defined in the network protocol profile that is associated with the port group. If you select **Static - IP Pool**, the IP address is assigned the first time the virtual machine or vApp is powered on. The assigned IP address persists across restarts. If you select **Transient - IP Pool**, an IP address is assigned every the virtual machine or vApp is powered on.

When the virtual machine is powered on, the adapters connected to the port group receive IP addresses from the range in the protocol profile. When the virtual machine is powered off, the IP addresses are released.

Related Links

[Add a Network Protocol Profile on page 2385](#)

A network protocol profile contains a pool of IPv4 and IPv6 addresses that vCenter Server assigns to vApps or to virtual machines with vApp functionality that are connected to port groups associated with the profile.

[Associate a Port Group with a Network Protocol Profile on page 2388](#)

To apply the range of IP addresses from a network protocol profile to a virtual machine that is a part of a vApp or has vApp functionality enabled, associate the profile with a port group that controls the networking of the virtual machine.

What is Multicast Filtering

In vSphere 6.0 and later, vSphere Distributed Switch supports basic and snooping models for filtering of multicast packets that are related to individual multicast groups. Choose a model according to the number of multicast groups to which the virtual machines on the switch subscribe.

Multicast Filtering Modes

In addition to the default basic mode for filtering multicast traffic, vSphere Distributed Switch 6.0.0 and later releases support multicast snooping that forwards multicast traffic in a more precise way based on the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) messages from virtual machines.

Basic Multicast Filtering

In basic multicast filtering mode, a vSphere Standard Switch or vSphere Distributed Switch forwards multicast traffic for virtual machines according to the destination MAC address of the multicast group. When joining a multicast group, the guest operating system pushes the multicast MAC address of the group down to the network through the switch. The switch saves the mapping between the port and the destination multicast MAC address in a local forwarding table.

The switch does not interpret the IGMP messages that a virtual machine sends to join or leave a group. The switch sends them directly to the local multicast router, which then interprets them to join the virtual machine to or remove it from the group.

The basic mode has the following restrictions:

- A virtual machine might receive packets from groups that it is not subscribed for because the switch forwards packets according to the destination MAC address of a multicast group, which can be potentially mapped up to 32 IP multicast groups.
- A virtual machine that is subscribed for traffic from more than 32 multicast MAC addresses receives packets that it is not subscribed for because of a limitation in the forwarding model.
- The switch does not filter packets according to source address as defined in IGMP version 3.

Multicast Snooping

In multicast snooping mode, a vSphere Distributed Switch provides IGMP and MLD snooping according to RFC 4541. The switch dispatches multicast traffic more precisely by using IP addresses. This mode supports IGMPv1, IGMPv2, and IGMPv3 for IPv4 multicast group addresses, and MLDv1 and MLDv2 for IPv6 multicast group addresses.

The switch dynamically detects the membership of a virtual machine. When a virtual machine sends a packet which contains IGMP or MLD membership information through a switch port, the switch creates a record about the destination IP address of the group, and in the case of IGMPv3, about a source IP address that the virtual machine prefers to receive traffic from. If a virtual machine does not renew its membership to a group within a certain period of time, the switch removes the entry for the group from the lookup records.

In multicast snooping mode of a distributed switch, a virtual machine can receive multicast traffic on a single switch port from up to 512 groups and 10 sources.

Note: In vSphere 6.7, the default multicast filtering mode is Basic. In vSphere 7.0, the default multicast filtering mode is IGMP/MLD snooping. When the DVS is upgraded to 7.0, the default multicast filtering mode will change from basic to IGMP/MLD snooping.

Enable Multicast Snooping on a vSphere Distributed Switch

Use multicast snooping on a vSphere Distributed Switch to forward traffic in a precise manner according to Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) membership information that virtual machines send to subscribe for multicast traffic.

Verify that vSphere Distributed Switch is version 6.0.0 and later.

Use multicast snooping if virtualized workloads on the switch subscribe to more than 32 multicast groups or must receive traffic from specific source nodes. For information about the multicast filtering modes of vSphere Distributed Switch, see [Multicast Filtering Modes](#).

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. From the **Actions** menu, select **Settings > Edit Settings**.
3. In the dialog box that displays the settings of the switch, click **Advanced**.
4. From the **Multicast filtering mode** drop-down menu, select **IGMP/MLD snooping**, and click **OK**.

Multicast snooping becomes active on hosts running ESXi 6.0 and later.

Edit the Query Time Interval for Multicast Snooping

When IGMP or MLD multicast snooping is enabled on a vSphere Distributed Switch, the switch sends general queries about the membership of virtual machines in case a snooping querier is not configured on the physical switch. On ESXi hosts that are attached to the distributed switch, you can edit the time interval in which the switch sends general queries.

The default time interval for sending snooping queries is 125 seconds.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **System** and select **Advanced System Settings**.
3. Click **Edit**.
4. Locate the `Net.IGMPQueryInterval` system setting and enter a new value in seconds for the setting.
5. Click **OK**.

Edit the Number of Source IP Addresses for IGMP and MLD

When you enable IGMP or MLD multicast snooping on a vSphere Distributed Switch, you can edit the maximum number of IP sources from which the members of a multicast group receive packets.

1. In the vSphere Client, navigate to the host.
2. On the **Configure** tab, expand **System** and select **Advanced System Settings**.
3. Click **Edit**.
4. Locate the `Net.IGMPV3MaxSrcIPNum` or `Net.MLDV2MaxSrcIPNum` system setting and enter a new value between 1 and 32 for the setting.
5. Click **OK**.

What is Stateless Network Deployment

Stateless is a mode of execution for ESXi hosts with no local storage that formerly would save configuration or state. Configurations are abstracted into a host profile, which is a template that applies to a class of machines. Stateless allows easy replacement, removal, and addition of failed hardware, and improves the ease of scaling a hardware deployment.

Every stateless ESXi boot is like a first boot. The ESXi host boots with networking connectivity to vCenter Server through the built-in standard switch. If the host profile specifies distributed switch membership, vCenter Server joins the ESXi host to VMware distributed switches.

When planning the network setup for stateless ESXi hosts, you should keep the configuration as generic as possible and avoid host-specific items. Currently the design has no hooks to reconfigure physical switches when deploying a new host. Any such requirement would need special handling.

To set up stateless deployment, one ESXi host must be installed in the standard fashion. Then find and record the following network-related information to save in the host profile:

- vSphere standard switch instances and settings (port groups, uplinks, MTU, and so forth)
- Distributed switch instances
- Selection rules for uplinks and uplink port or port groups
- vNIC information:
 - Address information (IPv4 or IPv6, static or DHCP, gateway)
 - Port groups and distributed port groups assigned to the physical network adapter (`vmknics`)
 - If there are distributed switches, record VLAN, physical NICs bound to the `vmknics`, and if `Etherchannel` is configured

The recorded information is used as a template for the host profile. Once the host profile virtual switch information has been extracted and placed in the host profile, you have the opportunity to change any of the information. Modifications are offered for both standard and distributed switches in these sections: uplink selection policy, based on either vmnic name or device number, and auto discovery based on VLAN ID. The (possibly modified) information is stored by the stateless boot infrastructure and applied to a stateless ESXi host on its next boot. During network initialization, a generic network plug-in interprets the recorded host profile setting and does the following:

- Loads appropriate physical NIC drivers.
- Creates all standard switch instances, along with port groups. It selects uplinks based on policy. If the policy is based on the VLAN ID, there is a probing process to gather relevant information.
- For VMkernel network adapters connected to the standard switch, it creates VMkernel network adapters and connects them to port groups.
- For each VMkernel network adapter connected to a distributed switch, it creates a temporary standard switch (as needed) with uplinks bound to the VMkernel network adapter. It creates a temporary port group with VLAN and teaming policies based on recorded information. Specifically, IP-hash is used if Etherchannel was used in the distributed switch.
- Configures all VMkernel network adapter settings (assigns address, gateway, MTU, and so forth).

Basic connectivity is functioning, and the networking setup is complete if there is no distributed switch present.

If there is a distributed switch present, the system stays in maintenance mode until distributed switch remediation is complete. No virtual machines are started at this time. Because distributed switches requires vCenter Server, the boot process continues until vCenter Server connectivity is established, and vCenter Server notices that the host should be part of a distributed switch. It issues a distributed switch host join, creating a distributed switch proxy standard switch on the host, selects appropriate uplinks, and migrates the vmknics from the standard switch to the distributed switch. When this operation is complete, it deletes the temporary standard switch and port groups.

At the end of the remediation process, the ESXi host is taken out of maintenance mode, and HA or DRS can start virtual machines on the host.

In the absence of a host profile, a temporary standard switch is created with “default networking” logic, which creates a management network switch (with no VLAN tag) whose uplink corresponds to the PXE booting vNIC. A vmknics is created on the management network port group with the same MAC address as the PXE booting vNIC. This logic was previously used for PXE booting. If there is a host profile, but the networking host profile is deactivated or fatally incomplete, vCenter Server falls back to default networking so that the ESXi host can be managed remotely. This triggers a compliance failure, so vCenter Server then initiates recovery actions.

vSphere Networking Best Practices

Consider these best practices when you configure your network.

- To ensure a stable connection between , , and other products and services, do not set connection limits and timeouts between the products. Setting limits and timeouts can affect the packet flow and cause services interruption.
- Isolate from one another the networks for host management, vSphere vMotion, vSphere FT, and so on, to improve security and performance.
- Dedicate a separate physical NIC to a group of virtual machines, or use Network I/O Control and traffic shaping to guarantee bandwidth to the virtual machines. This separation also enables distributing a portion of the total networking workload across multiple CPUs. The isolated virtual machines can then better handle application traffic, for example, from a vSphere Client.
- To physically separate network services and to dedicate a particular set of NICs to a specific network service, create a vSphere Standard Switch or vSphere Distributed Switch for each service. If this is not possible, separate network services on a single switch by attaching them to port groups with different VLAN IDs. In either case, verify with your network administrator that the networks or VLANs you choose are isolated from the rest of your environment and that no routers connect them.

- Keep the vSphere vMotion connection on a separate network. When migration with vMotion occurs, the contents of the guest operating system's memory is transmitted over the network. You can do this either by using VLANs to segment a single physical network or by using separate physical networks (the latter is preferable).
For migration across IP subnets and for using separate pools of buffer and sockets, place traffic for vMotion on the vMotion TCP/IP stack, and traffic for migration of powered-off virtual machines and cloning on the Provisioning TCP/IP stack. See [VMkernel Networking Layer](#).
- Multihoming in VMkernel networking is not supported. For more information, see <http://kb.vmware.com/kb/2010877>.
- You can add and remove network adapters from a standard or distributed switch without affecting the virtual machines or the network service that is running behind that switch. If you remove all the running hardware, the virtual machines can still communicate among themselves. If you leave one network adapter intact, all the virtual machines can still connect with the physical network.
- To protect your most sensitive virtual machines, deploy firewalls in virtual machines that route between virtual networks with uplinks to physical networks and pure virtual networks with no uplinks.
- For best performance, use VMXNET 3 virtual machine NICs.
- Physical network adapters connected to the same vSphere Standard Switch or vSphere Distributed Switch should also be connected to the same physical network.
- Configure the same MTU on all VMkernel network adapters in a vSphere Distributed Switch. If several VMkernel network adapters, configured with different MTUs, are connected to vSphere distributed switches, you might experience network connectivity problems.

Troubleshoot vSphere Networking

The troubleshooting topics about networking in vSphere provide solutions to potential problems that you might encounter with the connectivity of ESXi hosts, vCenter Server and virtual machines.

Guidelines for Troubleshooting Your Implementation of vSphere

To troubleshoot your implementation of vSphere, identify the symptoms of the problem, determine which of the components are affected, and test possible solutions.

Identifying Symptoms

A number of potential causes might lead to the under-performance or nonperformance of your implementation. The first step in efficient troubleshooting is to identify exactly what is going wrong.

Defining the Problem Space

After you have isolated the symptoms of the problem, you must define the problem space. Identify the software or hardware components that are affected and might be causing the problem and those components that are not involved.

Testing Possible Solutions

When you know what the symptoms of the problem are and which components are involved, test the solutions systematically until the problem is resolved.

How Do You Identify the Symptoms

Before you attempt to resolve a problem in your vSphere implementation, you must identify precisely how it is failing.

The first step in the troubleshooting process is to gather information that defines the specific symptoms of what is happening. You might ask these questions when gathering this information:

- What is the task or expected behavior that is not occurring?
- Can the affected task be divided into subtasks that you can evaluate separately?
- Is the task ending in an error? Is an error message associated with it?
- Is the task completing but in an unacceptably long time?
- Is the failure consistent or sporadic?

- What has changed recently in the software or hardware that might be related to the failure?

How Do You Define the Problem Space

After you identify the symptoms of the problem in your vSphere implementation, determine which components in your setup are affected, which components might be causing the problem, and which components are not involved.

To define the problem space in an implementation of vSphere, be aware of the components present. In addition to VMware software, consider third-party software in use and which hardware is being used with the VMware virtual hardware.

Recognizing the characteristics of the software and hardware elements and how they can impact the problem, you can explore general problems that might be causing the symptoms.

- Misconfiguration of software settings
- Failure of physical hardware
- Incompatibility of components

Break down the process and consider each piece and the likelihood of its involvement separately. For example, a case that is related to a virtual disk on local storage is probably unrelated to third-party router configuration. However, a local disk controller setting might be contributing to the problem. If a component is unrelated to the specific symptoms, you can probably eliminate it as a candidate for solution testing.

Think about what changed in the configuration recently before the problems started. Look for what is common in the problem. If several problems started at the same time, you can probably trace all the problems to the same cause.

How Do You Test Possible Solutions

After you know the symptoms of the problem in your vSphere implementation and which software or hardware components are most likely involved, you can systematically test solutions until you resolve the problem.

With the information that you have gained about the symptoms and affected components, you can design tests for pinpointing and resolving the problem. These tips might make this process more effective.

- Generate ideas for as many potential solutions as you can.
- Verify that each solution determines unequivocally whether the problem is fixed. Test each potential solution but move on promptly if the fix does not resolve the problem.
- Develop and pursue a hierarchy of potential solutions based on likelihood. Systematically eliminate each potential problem from the most likely to the least likely until the symptoms disappear.
- When testing potential solutions, change only one thing at a time. If your setup works after many things are changed at once, you might not be able to discern which of those things made a difference.
- If the changes that you made for a solution do not help resolve the problem, return the implementation to its previous status. If you do not return the implementation to its previous status, new errors might be introduced.
- Find a similar implementation that is working and test it in parallel with the implementation that is not working properly. Make changes on both systems at the same time until few differences or only one difference remains between them.

Troubleshooting with vCenter Server Logs

You can often obtain valuable troubleshooting information by looking at the logs provided by the various services and agents that your vSphere implementation is using.

Most logs are located in `/var/log/vmware/<service_name>` for vCenter Server deployments.

Common Logs

The following logs are common to all vCenter Server deployments.

Table 204: Common Log Directories

Log Directory	Description
../firstboot	Stores first boot logs
applmgmt and applmgmt-audit	Stores logs related to the VMware Appliance Management Service
cloudvm	Stores logs for allotment and distribution of resources between services
rhttpproxy	Stores logs for the VMware HTTP Reverse Proxy service
sca	Stores logs for the VMware Service Control Agent service
vapi	Stores logs for the VMware vAPI Endpoint service
vmafdd	Stores logs for the VMware Authentication Framework - LDAP service
vmldird	Stores logs for the VMware Directory Service - LDAP service
vmon	Stores logs for the VMware Service Lifecycle Manager service

Management Node Logs

The following logs are available if a management node deployment is selected.

Table 205: Management Node Log Directories

Log Directory	Service
rbd	VMware vSphere Auto Deploy
content-library	VMware Content Library Service
eam	VMware ESX Agent Manager
netdumper	VMware vSphere ESXi Dump Collector
perfcharts	VMware Performance Charts Service
vmcam	VMware vSphere Authentication Proxy
vmldird	VMware Directory Service - LDAP
vmware-sps	VMware vSphere Profile-Driven Storage Service
vpxd	VMware vCenter Server
vpostgres	VMware Postgres service
vcha	VMware vCenter High Availability service

Troubleshoot MAC Address Allocation

In vSphere, certain restrictions on the range of MAC addresses that can be assigned to virtual machines might cause lost of connectivity or inability to power on workloads.

Duplicate MAC Addresses of Virtual Machines on the Same Network

You encounter loss of packets and connectivity because virtual machines have duplicate MAC addresses generated by vCenter Server.

The MAC addresses of virtual machines on the same broadcast domain or IP subnet are in conflict, or vCenter Server generates a duplicate MAC address for a newly created virtual machine.

A virtual machine powers on and functions properly, but shares a MAC address with another virtual machine. This situation might cause packet loss and other problems.

Virtual machines might have duplicate MAC addresses due to several reasons.

- Two vCenter Server instances with identical IDs generate overlapping MAC addresses for virtual machine network adapters.
Each vCenter Server instance has an ID between 0 and 63 that is randomly generated at installation time, but can be reconfigured after installation. vCenter Server uses the instance ID to generate MAC addresses for the network adapters of the machine.
- A virtual machine has been transferred in power-off state from one vCenter Server instance to another in the same network, for example, by using shared storage, and a new virtual machine network adapter on the first vCenter Server receives the freed MAC address.
- Change the MAC address of a virtual machine network adapter manually.
If you have an existing virtual machine with a conflicting MAC address, you must provide a unique MAC address in the **Virtual Hardware** settings.
 - Power off the virtual machine, configure the adapter to use a manual MAC address, and type the new address.
 - If you cannot power the virtual machine off for configuration, re-create the network adapter that is in conflict with enabled manual MAC address assignment and type the new address. In the guest operating system, set the same static IP address to the re-added adapter as before.

For information about configuring the network adapters of virtual machines, see the *vSphere Networking* and *vSphere Virtual Machine Administration* documentation.

- If the vCenter Server instance generates the MAC addresses of virtual machines according to the default allocation, VMware OUI, change the vCenter Server instance ID or use another allocation method to resolve conflicts.

Note: Changing the vCenter Server instance ID or switching to a different allocation scheme does not resolve MAC address conflicts in existing virtual machines. Only virtual machines created or network adapters added after the change receive addresses according to the new scheme.

For information about MAC address allocation schemes and setup, see the *vSphere Networking* documentation.

Solution	Description
Change the vCenter Server ID	You can keep using the VMware OUI allocation scheme if your deployment contains a small number of vCenter Server instances. According to this scheme, a MAC address has the following format: <pre>00:50:56:XX:YY:ZZ</pre> where 00:50:56 represents the VMware OUI, <i>XX</i> is calculated as (80 + vCenter Server ID), and <i>YY:ZZ</i> is a random number. To change the vCenter Server ID, configure the vCenter Server unique ID option in the Runtime Settings section from the General settings of the vCenter Server instance and restart it. The VMware OUI allocation works with up to 64 vCenter Server instances and is suitable for small scale deployments.

Solution	Description
Switch to prefix-based allocation	You can use a custom OUI. For example, for a 02:12:34 locally administered address range, MAC addresses have the form 02:12:34:XX:YY:ZZ. You can use the fourth octet <i>xx</i> to distribute the OUI address space between the vCenter Server instances. This structure results in 255 address clusters, each cluster managed by a vCenter Server instance, and in about 65000 MAC addresses per vCenter Server. For example, 02:12:34:01:YY:ZZ for vCenter Server A, 02:12:34:02:YY:ZZ for vCenter Server B, and so on. Prefix-based allocation is suitable for deployments of a larger scale. For globally unique MAC addresses, the OUI must be registered in IEEE.

- a) Configure MAC address allocation.
- b) Apply the new MAC address allocation scheme to an existing virtual machine in its **Virtual Hardware** settings.
 - Power off a virtual machine, configure the adapter to use a manual MAC address, revert to automatic MAC address allocation, and power on the virtual machine.
 - If the virtual machine is in production and you cannot power it off for configuration, after you change the vCenter Server ID or the address allocation scheme, re-create the network adapter in conflict with enabled automatic MAC address assignment. In the guest operating system, set the same static IP address to the re-added adapter as before.
- Enforce MAC address regeneration when transferring a virtual machine between vCenter Server instances by using the virtual machine files from a datastore.
 - a) Power off a virtual machine, remove it from the inventory, and in its configuration file (`.vmx`), set the `ethernetX.addressType` parameter to `generated`.
`X` next to `ethernet` stands for the sequence number of the virtual NIC in the virtual machine.
 - b) Import the virtual machine from one vCenter Server system to another by registering the virtual machine from a datastore in the target vCenter Server.
 The virtual machine files can reside in a datastore that is shared between the two vCenter Server instances or can be uploaded to a datastore that is accessible only from the target vCenter Server system.
 For information about registering a virtual machine from a datastore, see *vSphere Virtual Machine Administration*.
 - c) Power on the virtual machines for the first time.
 While the virtual machine is starting up, an information icon appears on the virtual machine in the vSphere Client.
 - d) Right-click the virtual machine and select **Guest OS > Answer Question**.
 - e) Select the **I Copied It** option.
 The target vCenter Server re-generates the MAC address of the virtual machine. The new MAC address starts with the VMware OUI 00:0c:29 and is based on the BIOS UUID of the virtual machine. The BIOS UUID of the virtual machine is calculated from the BIOS UUID of the host.
- If the vCenter Server and hosts are version 6.0 and later and the vCenter Server instances are connected in Enhanced Linked Mode, migrate virtual machines by using vMotion across vCenter Server systems.
 When a virtual machine is migrated across vCenter Server systems, the source vCenter Server adds the MAC address of the virtual machine to a denylist and does not assign them to other virtual machines.

Attempt to Power On a Virtual Machine Fails Due to a MAC Address Conflict

After you set a certain static MAC address to a virtual machine adapter, you cannot power on the virtual machine.

In the vSphere Web Client, after you assign a MAC address within the range 00:50:56:40:YY:ZZ – 00:50:56:7F:YY:ZZ to a virtual machine, attempts to power the virtual machine on fail with a status message that the MAC address is in conflict.

00:50:56:XX:YY:ZZ is not a valid static Ethernet address. It conflicts with VMware reserved MACs for other usage.

You attempt to assign a MAC address which starts with the VMware OUI 00:50:56 and is within the address range allocated for host VMkernel adapters on the vCenter Server system.

If you want to preserve the VMware OUI prefix, set a static MAC address within the range 00:50:56:00:00:00 – 00:50:56:3F:FF:FF. Otherwise, set an arbitrary MAC address whose prefix is different from the VMware OUI one. For information about the ranges available for static MAC addresses that have the VMware OUI prefix, see the *vSphere Networking* documentation.

Unable to Remove a Host from a vSphere Distributed Switch

Under certain conditions, you might be unable to remove a host from the vSphere distributed switch.

- Attempts to remove a host from a vSphere distributed switch fail, and you receive a notification that resources are still in use. The notification that you receive might look like the following:

```
The resource '16' is in use.
vDS DSwitch port 16 is still on host 10.23.112.2 connected to MyVM nic=4000 type=vmVnic
```
- Attempts to remove a host proxy switch that still exists on the host from a previous networking configuration fail. For example, you moved the host to a different data center or vCenter Server system, or upgraded the ESXi and vCenter Server software, and created new networking configuration. When trying to remove the host proxy switch, the operation fails because resources on the proxy switch are still in use.

You cannot remove the host from the distributed switch or delete the host proxy switch because of the following reasons.

- There are VMkernel adapters on the switch that are in use.
- There are virtual machine network adapters connected to the switch.

Problem	Solution
Cannot remove a host from a distributed switch	- In the vSphere Client, navigate to the distributed switch. - On the Configure tab, select More > Ports. - Locate all ports that are still in use and check which VMkernel or virtual machine network adapters on the host are still attached to the ports. - Migrate or delete the VMkernel and virtual machine network adapters that are still connected to the switch. - Use the Add and Manage Hosts wizard in the vSphere Client to remove the host from the switch. After the host is removed, the host proxy switch is deleted automatically.
Cannot remove a host proxy switch	- In the vSphere Client, navigate to the host. - Delete or migrate any VMkernel or virtual machine network adapters that are still connected to the host proxy switch. - Delete the host proxy switch from the Networking view on the host.

Hosts on a vSphere Distributed Switch Lose Connectivity to vCenter Server

Hosts on a vSphere Distributed Switch cannot connect to vCenter Server after a port group configuration.

After you change the networking configuration of a port group on a vSphere Distributed Switch that contains the VMkernel adapters for the management network, the hosts on the switch lose connectivity to vCenter Server. In the vSphere Client the status of the hosts is nonresponsive.

On a vSphere Distributed Switch in the vCenter Server that has networking rollback deactivated, the port group containing the VMkernel adapters for the management network is misconfigured in vCenter Server and the invalid configuration is propagated to the hosts on the switch.

Note: In vSphere networking rollback is enabled by default. However, you can activate or deactivate rollbacks at the vCenter Server level. For more information, see the *vSphere Networking* documentation.

1. From the Direct Console User Interface (DCUI) to an affected host, use the **Restore vDS** option from the **Network Restore Options** menu to configure the uplinks and the ID of the VLAN for the management network.

The DCUI creates a local ephemeral port and applies the VLAN and uplink configuration to the port. The DCUI changes the VMkernel adapter for the management network to use the new host local port to restore connectivity to vCenter Server.

After the host re-connects to vCenter Server, the vSphere Client displays a warning that some hosts on the switch have different networking configuration from the configuration stored in vSphere distributed switch.

2. In the vSphere Client, configure the distributed port group for the management network with correct settings.

Situation	Solution
You have altered the port group configuration only once	You can roll the configuration of the port group back one step. Right-click the port group, click Restore Configuration , and select Restore to previous configuration .
You have backed up a valid configuration of the port group	You can restore the configuration of the port group by using the backup file. Right-click the port group, click Restore Configuration , and select Restore configuration from a file . You can also restore the configuration for the entire switch, including the port group, from a backup file for the switch.
You have performed more than one configuration step and you do not have a backup file	You must provide valid settings for the port group manually.

For information about networking rollback, recovery, and restore, see the *vSphere Networking* documentation.

3. Migrate the VMkernel adapter for the management network from the host local ephemeral port to a distributed port on the switch by using the **Add and Manage Hosts** wizard.

Unlike distributed ports, the ephemeral local port of the VMKernel has a non-numeric ID.

For information about handling VMkernel adapters through the **Add and Manage Hosts** wizard, see the *vSphere Networking* documentation.

4. Apply the configuration of the distributed port group and VMkernel adapter from vCenter Server to the host.
 - Push the correct configuration of the distributed port group and VMkernel adapter from vCenter Server to the host.
 - a. In the vSphere Client, navigate to the host.
 - b. On the **Configure** tab, click **Networking**.
 - c. From the **Virtual switches** list, select the distributed switch and click **Rectify the state of the selected distributed switch on the host**.
 - Wait until vCenter Server applies the settings within the next 24 hours.

Alarm for Loss of Network Redundancy on a Host

An alarm reports a loss of uplink redundancy on a vSphere standard or a distributed switch for a host.

No redundant physical NICs for a host are connected to a particular standard or a distributed switch, and the following alarm appears:

Host name or IP Network uplink redundancy lost

Only one physical NIC on the host is connected to a certain standard or a distributed switch. The redundant physical NICs are either down or are not assigned to the switch.

For example, assume that a host in your environment has physical NICs `vmnic0` and `vmnic1` connected to `vSwitch0`, and the physical NIC `vmnic1` goes offline, leaving only `vmnic0` connected to `vSwitch0`. As a result, the uplink redundancy for `vSwitch0` is lost on the host.

Check which switch has lost uplink redundancy on the host. Connect at least one more physical NIC on the host to this switch and reset the alarm to green. You can use the vSphere Client or the ESXi Shell.

If a physical NIC is down, try to bring it back up by using the ESXi Shell on the host.

For information about using the networking commands in the ESXi Shell, see *ESXCLI Reference*. For information about configuring networking on a host in the vSphere Client, see *vSphere Networking*.

Virtual Machines Lose Connectivity After Changing the Uplink Failover Order of a Distributed Port Group

Changes in the failover NIC order on a distributed port group cause the virtual machines associated with the group to disconnect from the external network.

After you rearrange the uplinks in the failover groups for a distributed port group in vCenter Server, for example, by using the vSphere Client, some virtual machines in the port group can no longer access the external network.

After changing the failover order, many reasons might cause virtual machines to lose connectivity to the external network.

- The host that runs the virtual machines does not have physical NICs associated with the uplinks that are set to active or standby. All uplinks that are associated with physical NICs from the host for the port group are moved to unused.
- A Link Aggregation Group (LAG) that has no physical NICs from the host is set as the only active uplink according to the requirements for using LACP in vSphere.
- If the virtual machine traffic is separated in VLANs, the host physical adapters for the active uplinks might be connected to trunk ports on the physical switch that do not handle traffic from these VLANs.
- If the port group is configured with IP hash load balancing policy, an active uplink adapter is connected to a physical switch port that might not be in an EtherChannel.

You can examine the connectivity of the virtual machines in the port group to associated host uplinks and uplink adapters from the central topology diagram of the distributed switch or from the proxy switch diagram for the host.

- Restore the failover order with the uplink that is associated with a single physical NIC on the host back to active.
- Create a port group with identical settings, make it use the valid uplink number for the host, and migrate the virtual machine networking to the port group.
- Move the NIC to an uplink that participates in the active failover group.

You can use the vSphere Client to move the host physical NIC to another uplink.

- Use the **Add and Manage Hosts** wizard on the distributed switch.
 - a. Navigate to the distributed switch in the vSphere Client.
 - b. From the **Actions** menu select **Add and Manage Hosts**.
 - c. On the **Select task** page, select the **Manage host networking** option and select the host.
 - d. To assign the NIC of the host to an active uplink, navigate to the **Manage physical network adapters** page and associate the NIC to the switch uplink.

- Move the NIC at the level of the host.
 - a. Navigate to the host in the vSphere Client, and on the **Configure** tab, expand the **Networking** menu.
 - b. Select **Virtual Switches** and select the distributed proxy switch.
 - c. Click **Manage the physical network adapters connected to the selected switch**, and move the NIC to the active uplink

Unable to Add a Physical Adapter to a vDS That Has Network I/O Control

You might be unable to add a physical adapter with low speed, for example, 1 Gbps, to a vSphere Distributed Switch that has vSphere Network I/O Control version 3 configured.

You try to add a physical adapter with low speed, for example, 1 Gbps, to a vSphere Distributed Switch that is connected to physical adapters with high speed, for example, 10 Gbps. Network I/O Control version 3 is enabled on the switch and bandwidth reservations exist for one or more system traffic types, such as vSphere management traffic, vSphere vMotion traffic, vSphere NFS traffic, and so on. The task for adding the physical adapter fails with a status message that a parameter is incorrect.

```
A specified parameter was not correct: spec.host[].backing.pnicSpec[]
```

Network I/O Control aligns the bandwidth that is available for reservation to the 10-Gbps speed of the individual physical adapters that are already connected to the distributed switch. After you reserve a part of this bandwidth, adding a physical adapter whose speed is less than 10 Gbps might not meet the potential needs of a system traffic type.

For information about Network I/O Control version 3, see the *vSphere Networking* documentation.

1. In the , navigate to the host.
2. On the **Configure** tab, expand the **System** group of settings.
3. Select **Advanced System Settings** and click **Edit**.
4. Type the physical adapters that you want to use outside the scope of Network I/O Control as a comma-separated list for the `Net.IOControlPnicOptOut` parameter.
For example: `vmnic2,vmnic3`
5. Click **OK** to apply the changes.
6. In the vSphere Client, add the physical adapter to the distributed switch.

Troubleshoot SR-IOV Enabled Workloads

Under certain conditions, you might experience connectivity or power-on problems with virtual machines that use SR-IOV to send data to physical network adapters.

SR-IOV Enabled Workload Cannot Communicate After You Change Its MAC Address

After you change the MAC address in the guest operating system of an SR-IOV enabled virtual machine, the virtual machine loses connectivity.

When you connect the network adapter of a virtual machine to an SR-IOV virtual function (VF), you create a passthrough network adapter for the virtual machine. After the (VF) driver in the guest operating system modifies the MAC address for the passthrough network adapter, the guest operating system shows that the change is successful but the VM network adapter loses connectivity. Although the guest operating system shows that the new MAC address is enabled, a log message in the `/var/log/vmkernel.log` file indicates that the operation has failed.

```
Requested mac address change to new MAC address on port VM NIC port number, disallowed by vswitch policy.
```

where

- `new MAC address` is the MAC address in the guest operation system.

- *VM NIC port number* is the port number of the VM network adapter in hexadecimal format.

The default security policy on the port group to which the passthrough network adapter is connected does not allow changes in the MAC address in the guest operating system. As a result, the networking interface in the guest operating system cannot acquire an IP address and loses connectivity.

In the guest operating system, reset the interface to cause the passthrough network adapter to regain its valid MAC address. If the interface is configured to use DHCP for address assignment, the interface acquires an IP address automatically.

For example, on a Linux virtual machine run the `ifconfig` console command.

```
ifconfig ethX down
ifconfig ethX up
```

where *X* in `ethX` represents the sequence number of the virtual machine network adapter in the guest operating system.

Virtual Machine That Uses an SR-IOV Virtual Function Fails to Power On Because the Host Is Out of Interrupt Vectors

On an ESXi host, one or more virtual machines that use SR-IOV virtual functions (VFs) for networking are powered off.

On an ESXi host, one or more virtual machines that use SR-IOV virtual functions (VFs) for networking fail to power on if the total number of assigned virtual functions is close to the maximum number of VFs specified in the *vSphere Configuration Maximums* guide.

The virtual machine log file `vmware.log` contains the following message about the VF:

```
PCIPassthruChangeIntrSettings: vf_name failed to register interrupt (error code 195887110)
```

The VMkernel log file `vmkernel.log` contains the following messages about the VF assigned to the virtual machine:

```
VMKPCIPassthru: 2565: BDF = vf_name intrType = 4 numVectors: 3
WARNING: IntrVector: 233: Out of interrupt vectors
```

The number of allocatable interrupt vectors scales up with the number of physical CPUs on an ESXi host. An ESXi host that has 32 CPUs can provide a total of 4096 interrupt vectors. When the host boots, devices on the host such as storage controllers, physical network adapters, and USB controllers consume a subset of the 4096 vectors. If these devices require more than 1024 vectors, the maximum number of potentially supported VFs is reduced.

When a virtual machine powers on and the guest operating system VF driver starts, interrupt vectors are consumed. If the required number of interrupt vectors is not available, the guest operating system shuts down unexpectedly without any error messages.

No rule presently exists to determine the number of interrupt vectors consumed or available on a host. This number depends on the hardware configuration of the host.

- To be able to power on the virtual machines, reduce the total number of VFs assigned to virtual machines on the host. For example, change the SR-IOV network adapter of a virtual machine to an adapter that is connected to a vSphere Standard Switch or vSphere Distributed Switch.

A Virtual Machine that Runs a VPN Client Causes Denial of Service for Virtual Machines on the Host or Across a vSphere HA Cluster

A virtual machine sending Bridge Protocol Data Unit (BPDU) frames, for example, a VPN client, causes some virtual machines connected to the same port group to lose connectivity. The transmission of BPDU frames might also break the connection of the host or of the parent vSphere HA cluster.

A virtual machine that is expected to send BPDU frames causes the traffic to the external network of the virtual machines in the same port group to be blocked.

If the virtual machine runs on a host that is a part of a vSphere HA cluster, and the host becomes network-isolated under certain conditions, you observe Denial of Service (DoS) on the hosts in the cluster.

As a best practice, a physical switch port that is connected to an ESXi host has the Port Fast and BPDU guard enabled to enforce the boundary of the Spanning Tree Protocol (STP). A standard or distributed switch does not support STP, and it does not send any BPDU frames to the switch port. However, if any BPDU frame from a compromised virtual machine arrives at a physical switch port facing an ESXi host, the BPDU guard feature deactivates the port to stop the frames from affecting the Spanning Tree Topology of the network.

In certain cases a virtual machine is expected to send BPDU frames, for example, when deploying VPN that is connected through a Windows bridge device or through a bridge function. If the physical switch port paired with the physical adapter that handles the traffic from this virtual machine has the BPDU guard on, the port is error-disabled, and the virtual machines and VMkernel adapters using the host physical adapter cannot communicate with the external network anymore.

If the teaming and failover policy of the port group contains more active uplinks, the BPDU traffic is moved to the adapter for the next active uplink. The new physical switch port becomes deactivated, and more workloads become unable to exchange packets with the network. Eventually, almost all entities on the ESXi host might become unreachable.

If the virtual machine runs on a host that is a part of a vSphere HA cluster, and the host becomes network-isolated because most of the physical switch ports connected to it are deactivated, the active primary host in the cluster moves the BPDU sender virtual machine to another host. The virtual machine starts disabling the physical switch ports connected to the new host. The migration across the vSphere HA cluster eventually leads to accumulated DoS across the entire cluster.

- If the VPN software must continue its work on the virtual machine, allow the traffic out of the virtual machine and configure the physical switch port individually to pass the BPDU frames.

Network Device	Configuration
Distributed or standard switch	Set the Forged Transmit security property on the port group to Accept to allow BPDU frames to leave the host and reach the physical switch port. You can isolate the settings and the physical adapter for the VPN traffic by placing the virtual machine in a separate port group and assigning the physical adapter to the group. CAUTION: Setting the Forged Transmit security property to Accept to enable a host to send BPDU frames carries a security risk because a compromised virtual machine can perform spoofing attacks.
Physical switch	• Keep the Port Fast enabled. • Enable the BPDU filter on the individual port. When a BPDU frame arrives at the port, it is filtered out. Note: Do not enable the BPDU filter globally. If the BPDU filter is enabled globally, the Port Fast mode becomes deactivated and all physical switch ports perform the full set of STP functions.

- To deploy a bridge device between two virtual machine NICs connected to the same Layer 2 network, allow the BPDU traffic out of the virtual machines and deactivate Port Fast and BPDU loop prevention features.

Network Device	Configuration
Distributed or standard switch	Set the Forged Transmit property of the security policy on the port groups to Accept to allow BPDU frames to leave the host and reach the physical switch port. You can isolate the settings and one or more physical adapters for the bridge traffic by placing the virtual machine in a separate port group and assigning the physical adapters to the group.

Network Device	Configuration
	CAUTION: Setting the Forged Transmit security property to Accept to enable bridge deployment carries a security risk because a compromised virtual machine can perform spoofing attacks.
Physical switch	- Deactivate Port Fast on the ports to the virtual bridge device to run STP on them. - Deactivate BPDU guard and filter on the ports facing the bridge device.

- Protect the environment from DoS attacks in any case by activating the BPDU filter on the ESXi host or on the physical switch.
 - On a host that does not have the Guest BPDU filter implemented enable the BPDU filter on the physical switch port to the virtual bridge device.

Network Device	Configuration
Distributed or standard switch	Set the Forged Transmit property of the security policy on the port group to Reject .
Physical switch	- Keep the Port Fast configuration. - Enable the BPDU filter on the individual physical switch port. When a BPDU frame arrives at the physical port, it is filtered out. Note: Do not enable the BPDU filter globally. If the BPDU filter is enabled globally, the Port Fast mode is deactivated and all physical switch ports perform the full set of STP functions.

Low Throughput for UDP Workloads on Windows Virtual Machines

When a Windows virtual machine in vSphere transmits large UDP packets, the throughput is lower than expected or is oscillating even when other traffic is negligible.

When a Windows virtual machine transmits UDP packets larger than 1024 bytes, you experience lower than expected or oscillating throughput even when other traffic is negligible. In case of a video streaming server, video playback pauses.

For every UDP packet larger than 1024 bytes, the Windows network stack waits for a transmit completion interrupt before sending the next packet. vSphere does not provide a transparent workaround of the situation.

- Increase the threshold in bytes at which Windows changes its behavior for UDP packets by modifying the registry of the Windows guest OS.
 - Locate the `HKLM\System\CurrentControlSet\Services\Afd\Parameters` registry key.
 - Add a value with the name `FastSendDatagramThreshold` of type `DWORD` equal to 1500.

For information about fixing this issue in the Windows registry, see <http://support.microsoft.com/kb/235257>.

- Modify the coalescing settings of the virtual machine NIC.

If the Windows virtual machine has a VMXNET3 vNIC adapter, configure one of the following parameters in the `.vmx` file of the virtual machine. Use the vSphere Client, or directly modify the `.vmx` file.

Action	Parameter	Value
Increase the interrupt rate of the virtual machine to a higher rate than expected packet rate. For example, if the expected packet rate is 15000 interrupts per second, set the interrupt rate to 16000 interrupts per second. Set the <code>ethernetX.coalescingScheme</code> parameter to <code>rbc</code> and the <code>ethernetX.coalescingParams</code> parameter to 16000. The default interrupt rate is 4000 interrupts per second.	<code>ethernetX.coalescingScheme</code> <code>ethernetX.coalescingParams</code>	<code>rbc</code> 16000
Deactivate coalescing for low throughput or latency-sensitive workloads. For information about configuring low-latency workloads, see Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs .	<code>ethernetX.coalescingScheme</code>	<code>disabled</code>
Revert to the coalescing algorithm from earlier ESXi releases. Note: The ability to revert to the earlier algorithm will not be available in later vSphere releases.	<code>ethernetX.coalescingScheme</code>	<code>calibrate</code>

`X` next to `ethernet` stands for the sequence number of the vNIC in the virtual machine.

For more information about configuring parameters in the `.vmx` file, see the *vSphere Virtual Machine Administration* documentation.

- Modify ESXi host coalescing settings.

This approach affects all virtual machines and all virtual machine NICs on the host.

You can edit the advanced system settings list for the host in the vSphere Client, or by using a vCLI console command on the host from the ESXi Shell.

Action	Parameter in the vSphere Client	Parameter for the <code>esxcli system settings advanced set</code> Command	Value
Set a default interrupt rate higher than the expected packet rate. For example, set the interrupt rate to 16000 if 15000 interrupts are expected per second.	<code>Net.CoalesceScheme</code> <code>Net.CoalesceParams</code>	<code>/Net/CoalesceScheme</code> <code>/Net/CoalesceParams</code>	<code>rbc</code> 16000

Action	Parameter in the vSphere Client	Parameter for the <code>esxcli system settings advanced set</code> Command	Value
Deactivate coalescing for low throughput or latency-sensitive workloads. For information about configuring low-latency workloads, see Best Practices for Performance Tuning of Latency-Sensitive Workloads in vSphere VMs .	<code>Net.CoalesceDefaultOn</code>	<code>/Net/CoalesceDefaultOn</code>	0
Revert to the coalescing scheme from earlier ESXi releases. Note: The ability to revert to the earlier algorithm will not be available in later vSphere releases.	<code>Net.CoalesceScheme</code>	<code>/Net/CoalesceScheme</code>	calibrate

For information about configuring a host from the vSphere Client, see the *vCenter Server and Host Management* documentation. For information about setting host properties by using a vCLI command, refer to the *ESXCLI Reference* documentation.

VMs on the Same Distributed Port Group are Unable to Communicate

Under certain conditions, the virtual machines that are on the same distributed port group but on different hosts cannot communicate with each other.

Virtual machines that reside on different hosts and on the same port group are unable to communicate. Pings from one virtual machine to another have no effect. You cannot migrate the virtual machines between the hosts by using vMotion.

- There are no physical NICs on some of the hosts assigned to active or standby uplinks in the teaming and failover order of the distributed port group.
- The physical NICs on the hosts that are assigned to the active or standby uplinks reside in different VLANs on the physical switch. The physical NICs in different VLANs cannot see each other and thus cannot communicate with each other.
- In the topology of the distributed switch, check which host does not have physical NICs assigned to an active or standby uplink on the distributed port group. Assign at least one physical NIC on that host to an active uplink on the port group.
- In the topology of the distributed switch, check the VLAN IDs of the physical NICs that are assigned to the active uplinks on the distributed port group. On all hosts, assign physical NICs that are from the same VLAN to an active uplink on the distributed port group.
- To verify that there is no problem at the physical layer, migrate the virtual machines to the same host and check the communication between them. Verify that inbound and outbound ICMP traffic is enabled in the guest OS. By default ICMP traffic is deactivated in Windows Server 2008 and Windows Server 2012.

Attempt to Power On a Migrated vApp Fails Because the Associated Protocol Profile Is Missing

You cannot power on a vApp or virtual machine that you transferred to a data center or a vCenter Server system because a network protocol profile is missing.

After you cold migrate a vApp or a virtual machine to another data center or vCenter Server system, an attempt to power it on fails. An error message states that a property cannot be initialized or allocated because the network of the vApp or virtual machine does not have an associated network protocol profile.

```
Cannot initialize property 'property'. Network 'port group' has no associated network protocol profile.
```

```
Cannot allocate IP address for property 'property'. Network 'port group' has no associated network protocol profile.
```

By using the OVF environment, the vApp or virtual machine retrieves network settings from a network protocol profile that is associated with the port group of the vApp or virtual machine.

vCenter Server creates such a network protocol profile for you when you install the OVF of a vApp and associates the profile with the port group that you specify during the installation.

The mapping between the protocol profile and port group is valid only in the scope of a data center. When you move the vApp, the protocol profile is not transferred to the target data center because of the following reasons:

- The network settings of the protocol profile might not be valid in the network environment of the target data center.
- A port group that has the same name and is associated with another protocol profile might already exist in the target data center, and vApps and virtual machines might be connected to this group. Replacing the protocol profiles for the port group might affect the connectivity of these vApp and virtual machines.
- Create a network protocol profile on the target data center or vCenter Server system with the required network settings and associate the protocol profile with the port group to which the vApp or virtual machine is connected. For example, this approach is suitable if the vApp or virtual machine is a vCenter Server extension that uses the vCenter Extension vService.

For information about providing network settings to a vApp or virtual machine from a network protocol profile, see the *vSphere Networking* documentation.

- Use the vSphere Client to export the OVF file of the vApp or virtual machine from the source data center or vCenter Server system and deploy it on the target data center or vCenter Server system.

When you use the vSphere Client to deploy the OVF file, the target vCenter Server system creates the network protocol profile for the vApp.

For information about managing OVF files in the vSphere Client, see the *vSphere Virtual Machine Administration* documentation.

Networking Configuration Operation Is Rolled Back and a Host Is Disconnected from vCenter Server

When you attempt to add or configure networking on a vSphere Distributed Switch on a host, the operation is rolled back and the host is disconnected from vCenter Server.

An attempt to perform a networking configuration operation on a vSphere Distributed Switch on a host, such as creating a virtual machine adapter or a port group, causes the host to disconnect from vCenter Server and results in the error message `Transaction has rolled back on the host`.

Under stressful conditions on a host, that is, if many concurrent networking operations compete for limited resources, the time to perform some of the operations might exceed the default timeout for rollback of network configuration operations on the distributed switch. As a result, these operations are rolled back.

For example, such a condition might come up when you create a VMkernel adapter on a host that has a very high number of switch ports or virtual adapters, all of which consume system resources on the host.

The default timeout to roll an operation back is 30 seconds.

- Use the vSphere Client to increase the timeout for rollback on vCenter Server.
If you encounter the same problem again, increase the rollback timeout with 60 seconds incrementally until the operation has enough time to succeed.
 - a) On the **Configure** tab of a vCenter Server instance, expand **Settings**.
 - b) Select **Advanced Settings** and click **Edit**.
 - c) If the property is not present, add the `config.vpxd.network.rollbackTimeout` parameter to the settings.
 - d) Type a new value, in seconds, for the `config.vpxd.network.rollbackTimeout` parameter
 - e) Click **OK**.
 - f) Restart the vCenter Server system to apply the changes.
- Increase the timeout for rollback by editing the `vpxd.cfg` configuration file.
If you encounter the same problem again, increase the rollback timeout with 60 seconds incrementally until the operation has enough time to succeed.
 - a) On the host machine of , navigate to the directory `/etc/vmware-vpx`.
 - b) Open the `vpxd.cfg` file for editing.
 - c) Under the `<network>` section, increase the timeout, in the `<rollbackTimeout>` element.

```
<config>
  <vpxd>
    <network>
      <rollbackTimeout>60</rollbackTimeout>
    </network>
  </vpxd>
</config>
```
 - d) Save and close the file.
 - e) Restart the vCenter Server system to apply the changes.

vSphere Storage

vSphere Storage describes virtualized and software-defined storage technologies that VMware ESXi™ and VMware vCenter Server® offer, and explains how to configure and use these technologies.

Intended Audience

This information is for experienced system administrators who are familiar with the virtual machine and storage virtualization technologies, data center operations, and SAN storage concepts.

Updated Information

This *vSphere Storage* is updated with each release of the product or when necessary.

This table provides the update history of *vSphere Storage*.

Revision	Description
15 SEP 2025	Updated Checking Metadata Consistency with vSphere On-disk Metadata Analyzer with new VOMA command options.
05 MAR 2025	Added information about the datastore status. See Display Datastore Information .
14 AUG 2024	Added information about creating an NFS user account to be used by the NFS 4.1 for Kerberos authentication to Enable Kerberos Authentication in Active Directory .
29 JUL 2024	Added a required privilege to Upload Files or Folders to Datastores .
25 JUN 2024	Initial release.

Introduction to Storage in vSphere Environment

vSphere supports various storage options and functionalities in traditional and software-defined storage environments. A high-level overview of vSphere storage elements and aspects helps you plan a proper storage strategy for your virtual data center.

Traditional Storage Virtualization Models in vSphere Environment

Generally, storage virtualization refers to a logical abstraction of physical storage resources and capacities from virtual machines and their applications. ESXi provides host-level storage virtualization.

In vSphere environment, a traditional model is built around the following storage technologies and ESXi and vCenter Server virtualization functionalities.

Local and Networked Storage

In traditional storage environments, the ESXi storage management process starts with storage space that your storage administrator preallocates on different storage systems. ESXi supports local storage and networked storage.

See [What Types of Physical Storage Does vSphere Support](#).

Storage Area Networks

A storage area network (SAN) is a specialized high-speed network that connects computer systems, or ESXi hosts, to high-performance storage systems. ESXi can use Fibre Channel or iSCSI protocols to connect to storage systems.

Fibre Channel

See [Using ESXi with a SAN](#).

Fibre Channel (FC) is a storage protocol that the SAN uses to transfer data traffic from ESXi host servers to shared storage. The protocol packages SCSI commands into FC frames. To connect to the FC SAN, your host uses Fibre Channel host bus adapters (HBAs).

Internet SCSI

See [Using ESXi with Fibre Channel SAN](#).

Internet iSCSI (iSCSI) is a SAN transport that can use Ethernet connections between computer systems, or ESXi hosts, and high-performance storage systems. To connect to the storage systems, your hosts use hardware iSCSI adapters or software iSCSI initiators with standard network adapters.

Storage Device or LUN

See [Using ESXi with iSCSI SAN](#).

In the ESXi context, the terms storage device and LUN are used interchangeably. Typically, both terms mean a storage volume that is presented to the host from a block storage system and is available for formatting.

See [Target and Device Representations](#) and [Managing ESXi Storage Devices](#).

Virtual Disks

A virtual machine on an ESXi host uses a virtual disk to store its operating system, application files, and other data associated with its activities. Virtual disks are large physical files, or sets of files, that can be copied, moved, archived, and backed up as any other files. You can configure virtual machines with multiple virtual disks. To access virtual disks, a virtual machine uses virtual NVMe or SCSI controllers. These virtual controllers include BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, VMware Paravirtual, NVMe, and others.

Each virtual disk resides on a datastore that is deployed on physical storage. From the standpoint of the virtual machine, each virtual disk appears as if it were a SCSI or NVMe drive connected to a SCSI or NVMe controller. Whether the physical storage is accessed through storage or network adapters on the host is typically transparent to the VM guest operating system and applications.

For information about configuring controllers for VMs, see [SCSI, SATA, and NVMe Storage Controller Conditions, Limitations, and Compatibility](#).

VMware vSphere® VMFS

The datastores that you deploy on block storage devices use the native vSphere Virtual Machine File System (VMFS) format. It is a special high-performance file system format that is optimized for storing virtual machines.

See [vSphere VMFS Datastore Concepts and Operations](#).

NFS

An NFS client built into ESXi uses the Network File System (NFS) protocol over TCP/IP to access an NFS volume that is located on a NAS server. The ESXi host can mount the volume and use it as an NFS datastore.

See [NFS Datastore Concepts and Operations in vSphere Environment](#).

Raw Device Mapping

In addition to virtual disks, vSphere offers a mechanism called raw device mapping (RDM). RDM is useful when a guest operating

system inside a virtual machine requires direct access to a storage device. For information about RDMS, see [Raw Device Mapping in vSphere](#).

Software-Defined Storage Models in vSphere Environment

In addition to abstracting underlying storage capacities from vSphere VMs, as traditional storage models do, software-defined storage abstracts storage capabilities.

With the software-defined storage model, a virtual machine becomes a unit of storage provisioning and can be managed through a flexible policy-based mechanism. The model involves the following vSphere technologies.

VMware vSphere® Virtual Volumes™ (vVols)	The Virtual Volumes functionality changes the storage management paradigm from managing space inside datastores to managing abstract storage objects handled by storage arrays. With Virtual Volumes, an individual virtual machine, not the datastore, becomes a unit of storage management. And storage hardware gains complete control over virtual disk content, layout, and management. See Working with VMware vSphere Virtual Volumes .
VMware vSAN	vSAN is a distributed layer of software that runs natively as a part of the hypervisor. vSAN aggregates local or direct-attached capacity devices of an ESXi host cluster and creates a single storage pool shared across all hosts in the vSAN cluster. See the <i>Administering VMware vSAN</i> documentation.
Storage Policy Based Management	Storage Policy Based Management (SPBM) is a framework that provides a single control panel across various data services and storage solutions, including vSAN and Virtual Volumes. Using storage policies, the framework aligns application demands of your virtual machines with capabilities provided by storage entities. See Storage Policy Based Management in vSphere .
I/O Filters	I/O filters are software components that can be installed on ESXi hosts and can offer additional data services to virtual machines. Depending on implementation, the services might include replication, encryption, caching, and so on. See Filtering Virtual Machine I/O in vSphere .

vSphere Storage APIs

Storage APIs is a family of APIs used by third-party hardware, software, and storage providers to develop components that enhance several vSphere features and solutions.

This Storage publication describes several Storage APIs that contribute to your storage environment. For information about other APIs from this family, including vSphere APIs - Data Protection, see the VMware website.

vSphere APIs for Storage Awareness

Also known as VASA, these APIs, either supplied by third-party vendors or offered by VMware, enable communications between vCenter Server and underlying storage. Through VASA, storage entities can inform vCenter Server about their configurations, capabilities, and storage health and events. In return, VASA can deliver VM storage requirements from vCenter Server to a storage entity and ensure that the storage layer meets the requirements.

VASA becomes essential when you work with Virtual Volumes, vSAN, vSphere APIs for I/O Filtering (VAIO), and storage VM policies. See [Using Storage Providers in vSphere](#).

vSphere APIs for Array Integration

These APIs, also known as VAAI, include the following components:

- **Hardware Acceleration APIs.** Help arrays to integrate with vSphere, so that vSphere can offload certain storage operations to the array. This integration significantly reduces CPU overhead on the host. See [Storage Hardware Acceleration in vSphere](#).
- **Array Thin Provisioning APIs.** Help to monitor space use on thin-provisioned storage arrays to prevent out-of-space conditions, and to perform space reclamation. See [ESXi and Array Thin Provisioning](#).

vSphere APIs for Multipathing

Known as the Pluggable Storage Architecture (PSA), these APIs allow storage partners to create and deliver multipathing and load-balancing plug-ins that are optimized for each array. Plug-ins communicate with storage arrays and determine the best path selection strategy to increase I/O performance and reliability from the ESXi host to the storage array. For more information, see [Using Pluggable Storage Architecture and Path Management with ESXi](#).

Getting Started with a Traditional Storage Model in vSphere Environment

Setting up your ESXi storage in traditional environments, includes configuring your storage systems and devices, enabling storage adapters, and creating datastores.

What Types of Physical Storage Does ESXi Support

In traditional storage environments, the ESXi storage management process starts with storage space that your storage administrator preallocates on different storage systems. ESXi supports local storage and networked storage.

Local Storage

Local storage can be internal hard disks located inside your ESXi host. It can also include external storage systems located outside and connected to the host directly through protocols such as SAS or SATA.

Local storage does not require a storage network to communicate with your host. You need a cable connected to the storage unit and, when required, a compatible HBA in your host.

The following illustration depicts a virtual machine using local SCSI storage.

Figure 89: Local Storage

In this example of a local storage topology, the ESXi host uses a single connection to a storage device. On that device, you can create a VMFS datastore, which you use to store virtual machine disk files.

Although this storage configuration is possible, it is not a best practice. Using single connections between storage devices and hosts creates single points of failure (SPOF) that can cause interruptions when a connection becomes unreliable or fails. However, because most of local storage devices do not support multiple connections, you cannot use multiple paths to access local storage.

ESXi supports various local storage devices, including SCSI, IDE, SATA, USB, SAS, flash, and NVMe devices.

Note: You cannot use IDE/ATA or USB drives to store virtual machines.

Local storage does not support sharing across multiple hosts. Only one host has access to a datastore on a local storage device. As a result, although you can use local storage to create VMs, you cannot use VMware features that require shared storage, such as HA and vMotion.

However, if you use a cluster of hosts that have just local storage devices, you can implement vSAN. vSAN transforms local storage resources into software-defined shared storage. With vSAN, you can use features that require shared storage. For details, see the *Administering VMware vSAN* documentation.

Networked Storage

Networked storage consists of external storage systems that your ESXi host uses to store virtual machine files remotely. Typically, the host accesses these systems over a high-speed storage network.

Networked storage devices are shared. Datastores on networked storage devices can be accessed by multiple hosts concurrently. ESXi supports multiple networked storage technologies.

In addition to traditional networked storage that this topic covers, VMware supports virtualized shared storage, such as vSAN. vSAN transforms internal storage resources of your ESXi hosts into shared storage that provides such capabilities as High Availability and vMotion for virtual machines. For details, see the *Administering VMware vSAN* documentation.


Note: The same LUN cannot be presented to an ESXi host or multiple hosts through different storage protocols. To access the LUN, hosts must always use a single protocol, for example, either Fibre Channel only or iSCSI only.

Fibre Channel (FC)

Stores virtual machine files remotely on an FC storage area network (SAN). FC SAN is a specialized high-speed network that connects your hosts to high-performance storage devices. The network uses Fibre Channel protocol to transport SCSI or NVMe traffic from virtual machines to the FC SAN devices.

To connect to the FC SAN, your host should be equipped with Fibre Channel host bus adapters (HBAs). Unless you use Fibre Channel direct connect storage, you need Fibre Channel switches to route storage traffic. If your host contains FCoE (Fibre Channel over Ethernet) adapters, you can connect to your shared Fibre Channel devices by using an Ethernet network.

Fibre Channel Storage depicts virtual machines using Fibre Channel storage.

Figure 90: Fibre Channel Storage

In this configuration, a host connects to a SAN fabric, which consists of Fibre Channel switches and storage arrays, using a Fibre Channel adapter. LUNs from a storage array become available to the host. You can access the LUNs and create datastores for your storage needs. The datastores use the VMFS format.

For specific information on setting up the Fibre Channel SAN, see [Using ESXi with Fibre Channel SAN](#).

Internet SCSI (iSCSI)

Stores virtual machine files on remote iSCSI storage devices. iSCSI packages SCSI storage traffic into the TCP/IP protocol, so that it can travel through standard TCP/IP networks instead of the specialized FC network. With an iSCSI connection, your host serves as the initiator that communicates with a target, located in remote iSCSI storage systems.

ESXi offers the following types of iSCSI connections:

Hardware iSCSI


Your host connects to storage through a third-party adapter capable of offloading the iSCSI and network processing. Hardware adapters can be dependent and independent.

Software iSCSI

Your host uses a software-based iSCSI initiator in the VMkernel to connect to storage. With this type of iSCSI connection, your host needs only a standard network adapter for network connectivity.

You must configure iSCSI initiators for the host to access and display iSCSI storage devices.

iSCSI Storage depicts different types of iSCSI initiators.

Figure 91: iSCSI Storage

In the left example, the host uses the hardware iSCSI adapter to connect to the iSCSI storage system.

In the right example, the host uses a software iSCSI adapter and an Ethernet NIC to connect to the iSCSI storage.

iSCSI storage devices from the storage system become available to the host. You can access the storage devices and create VMFS datastores for your storage needs.


For specific information on setting up the iSCSI SAN, see [Using ESXi with iSCSI SAN](#).

Network-attached Storage (NAS)

Stores virtual machine files on remote file servers accessed over a standard TCP/IP network. The NFS client built into ESXi uses Network File System (NFS) protocol version 3 and 4.1 to communicate with the NAS/NFS servers. For network connectivity, the host requires a standard network adapter.

You can mount an NFS volume directly on the ESXi host. You then use the NFS datastore to store and manage virtual machines in the same way that you use the VMFS datastores.

NFS Storage depicts a virtual machine using the NFS datastore to store its files. In this configuration, the host connects to the NAS server, which stores the virtual disk files, through a regular network adapter.

Figure 92: NFS Storage

For specific information on setting up NFS storage, see [NFS Datastore Concepts and Operations in vSphere Environment](#).

Shared Serial Attached SCSI (SAS)

Stores virtual machines on direct-attached SAS storage systems that offer shared access to multiple hosts. This type of access permits multiple hosts to access the same VMFS datastore on a LUN.

NVMe over Fabrics Storage

VMware NVMe over Fabrics (NVMe-oF) provides a distance connectivity between a host and a target storage device on a shared storage array. VMware supports different technologies, including NVMe over RDMA (with RoCE v2 technology), NVMe over Fibre Channel, and NVMe over TCP/IP. For more information, see [About VMware NVMe Storage](#).

Comparing Types of Storage

Whether certain vSphere functionality is supported might depend on the storage technology that you use.

The following table compares networked storage technologies that ESXi supports.

Table 206: Networked Storage that ESXi Supports

Technology	Protocols	Transfers	Interface
Fibre Channel	FC/SCSI, FC/NVMe	Block access of data/LUN	FC HBA
Fibre Channel over Ethernet	FCoE/SCSI	Block access of data/LUN	Converged Network Adapter (hardware FCoE)
iSCSI	IP/SCSI	Block access of data/LUN	- iSCSI HBA or iSCSI-enabled NIC (hardware iSCSI) - Network adapter (software iSCSI)
NAS	IP/NFS	File (no direct LUN access)	Network adapter

The following table compares the vSphere features that different types of storage support.

Table 207: vSphere Features Supported by Storage


Storage Type	Boot VM	vMotion	Datastore	RDM	VM Cluster	vSphere HA and DRS	Storage APIs - Data Protection
Local Storage	Yes	No	VMFS	No	Yes	No	Yes
Fibre Channel	Yes	Yes	VMFS	Yes	Yes	Yes	Yes
iSCSI	Yes	Yes	VMFS	Yes	Yes	Yes	Yes
NAS over NFS	Yes	Yes	NFS 3 and NFS 4.1	No	No	Yes	Yes

Note: Local storage supports a cluster of virtual machines on a single host (also known as a cluster in a box). A shared virtual disk is required. For more information about this configuration, see the *vSphere Resource Management* documentation.

Target and Device Representations

In the ESXi context, the term target identifies a single storage unit that the host can access. The terms storage device and LUN describe a logical volume that represents storage space on a target. In the ESXi context, both terms also mean a storage volume that is presented to the host from a storage target and is available for formatting. Storage device and LUN are often used interchangeably.

Different storage vendors present the storage systems to ESXi hosts in different ways. Some vendors present a single target with multiple storage devices or LUNs on it, while others present multiple targets with one LUN each.

Figure 93: Target and LUN Representations

In this illustration, three LUNs are available in each configuration. In one case, the host connects to one target, but that target has three LUNs that can be used. Each LUN represents an individual storage volume. In the other example, the host detects three different targets, each having one LUN.

Targets that are accessed through the network have unique names that are provided by the storage systems. The iSCSI targets use iSCSI names. Fibre Channel targets use World Wide Names (WWNs).

Note: ESXi does not support accessing the same LUN through different transport protocols, such as iSCSI and Fibre Channel.


A device, or LUN, is identified by its UUID name. If a LUN is shared by multiple hosts, it must be presented to all hosts with the same UUID.

How Virtual Machines Access Storage

When a virtual machine communicates with its virtual disk stored on a datastore, it issues SCSI or NVMe commands. Because datastores can exist on various types of physical storage, these commands are encapsulated into other forms, depending on the protocol that the ESXi host uses to connect to a storage device.

Regardless of the type of storage device your host uses, the virtual disk always appears to the virtual machine as a mounted SCSI or NVMe device. The virtual disk hides a physical storage layer from the virtual machine's operating system. This allows you to run operating systems that are not certified for specific storage equipment, such as SAN, inside the virtual machine.

The following graphic depicts five virtual machines using different types of storage to illustrate the differences between each type.

Figure 94: Virtual machines accessing different types of storage

Note: This diagram is for conceptual purposes only. It is not a recommended configuration.

View Storage Adapters Available on an ESXi Host

Storage adapters provide connectivity for your ESXi host to a specific storage unit or network. ESXi supports different classes of adapters, including SCSI, NVMe, iSCSI, RAID, Fibre Channel, Fibre Channel over Ethernet (FCoE), and Ethernet. ESXi accesses the adapters directly through device drivers in the VMkernel.

You must enable certain adapters, for example software iSCSI, before you can view their information. To configure adapters, see [Configuring iSCSI and iSER Adapters and Storage with ESXi](#).

Depending on the type of storage you use, you might need to enable and configure a storage adapter on your host.

You can display details for the available storage adapters and review their information.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**.
4. Use the icons to perform storage adapter tasks.

Availability of specific icons depends on the storage configuration.

Icon	Description
Add Software Adapter	Add a storage adapter. Applies to software iSCSI.
Refresh	Refresh information about storage adapters, topology, and file systems on the host.
Rescan Storage	Rescan all storage adapters on the host to discover newly added storage devices or VMFS datastores.
Rescan Adapter	Rescan the selected adapter to discover newly added storage devices.

5. To view details for a specific adapter, select the adapter from the list.
6. Use tabs under Adapter Details to access additional information and modify properties for the selected adapter.

Tab	Description
Properties	Review general adapter properties that typically include a name and model of the adapter and unique identifiers formed according to specific storage standards. For some types of adapters, you can use this tab to configure additional properties, for example, authentication.
Devices	View storage devices the adapter can access. Use the tab to perform basic device management tasks. See Display Storage Devices for an Adapter .
Paths	List and manage all paths the adapter uses to access storage devices.
Targets (Fibre Channel and iSCSI)	Review and manage targets accessed through the adapter.
Network Port Binding (iSCSI only)	Configure port binding for software and dependent hardware iSCSI adapters.
Advanced Options (iSCSI only)	Configure advanced parameters for iSCSI.

Characteristics of ESXi Datastores

Datastores are logical containers, analogous to file systems, that hide specifics of each storage device and provide a uniform model for storing virtual machine files. You can display all datastores available to your ESXi hosts and analyze their properties.

Datastores are added to vCenter Server in the following ways:

- You can create a VMFS datastore, an NFS version 3 or 4.1 datastore, or a Virtual Volumes datastore using the New Datastore wizard. A vSAN datastore is automatically created when you enable vSAN.
- When you add an ESXi host to vCenter Server, all datastores on the host are added to vCenter Server.

The following table describes datastore details that you can see when you review datastores through the vSphere Client. Certain characteristic might not be available or applicable to all types of datastores.

Table 208: Datastore Information

Datastore Information	Applicable Datastore Type	Description
Name	VMFS NFS vSAN vVol	Editable name that you assign to a datastore. For information on renaming a datastore, see Change Datastore Name .
Type	VMFS NFS vSAN vVol	File system that the datastore uses. For information about VMFS and NFS datastores and how to manage them, see Working with Datastores in vSphere Storage Environment . For information about vSAN datastores, see the <i>Administering VMware vSAN</i> documentation. For information about Virtual Volumes, see Working with VMware vSphere Virtual Volumes .
Device Backing	VMFS NFS vSAN	Information about underlying storage, such as a storage device on which the datastore is deployed (VMFS), server and folder (NFS), or disk groups (vSAN).
Protocol Endpoints	vVol	Information about corresponding protocol endpoints. See Static Protocol Endpoints .
Extents	VMFS	Individual extents that the datastore spans and their capacity.
Drive Type	VMFS	Type of the underlying storage device, such as a flash drive or a regular HDD drive. For details, see Using Flash Devices with ESXi Hosts .
Capacity	VMFS NFS vSAN vVol	Includes total capacity, provisioned space, and free space.
Mount Point	VMFS NFS vSAN vVol	A path to the datastore in the <code>/vmfs/volumes/</code> directory of the host.
Capability Sets	VMFS Note: A multi-extent VMFS datastore assumes capabilities of only one of its extents. NFS vSAN vVol	Information about storage data services that the underlying storage entity provides. You cannot modify them.
Storage I/O Control	VMFS NFS	Information on whether cluster-wide storage I/O prioritization is enabled. See the <i>vSphere Resource Management</i> documentation.
Hardware Acceleration	VMFS NFS vSAN vVol	Information on whether the underlying storage entity supports hardware acceleration. The status can be Supported, Not Supported, or Unknown. For details, see Storage Hardware Acceleration in vSphere . Note: NFS 4.1 does not support Hardware Acceleration.
Tags	VMFS NFS vSAN vVol	Datastore capabilities that you define and associate with datastores in a form of tags. For information, see Assign Tags to Datastores .
Connectivity with Hosts	VMFS	Hosts where the datastore is mounted.

Datastore Information	Applicable Datastore Type	Description
	NFS vVol	
Multipathing	VMFS vVol	Path selection policy the host uses to access storage. For more information, see Understanding Multipathing and Failover in the ESXi Environment .

Display Datastore Information

Access the Datastores view with the vSphere Client navigator.

Use the Datastores view to list all datastores available in the vSphere infrastructure inventory, analyze the information, and modify properties.

1. Navigate to any inventory object that is a valid parent object of a datastore, such as a host, a cluster, or a data center, and click the **Datastores** tab.

Datastores that are available in the inventory appear in the center panel.

You can check the status of the datastore.

Normal	The datastore is healthy and without problems.
Warning	The datastore experiences a potential problem that might require investigation. For example, the datastore is nearing capacity.
Critical	The datastore experiences a serious problem that requires immediate attention. For example, data corruption or connectivity issues.
Unknown	vCenter Server cannot determine the status due to datastore connectivity issues or other problems.

2. Use the options from a datastore right-click menu to perform basic tasks for a selected datastore.

Availability of specific options depends on the type of the datastore and its configuration.

Option	Description
Register VM	Register an existing virtual machine in the inventory. See the <i>vSphere Virtual Machine Administration</i> documentation.
Increase Datastore Capacity	Increase the capacity of the VMFS datastore or add an extent. See Increase Capacity of vSphere VMFS Datastore .
Browse Files	Navigate to the datastore file browser. See Navigate to the Datastore Browser .
Rename	Change the datastore name. See Change Datastore Name .
Mount Datastore	Mount the datastore to certain hosts. See Mount Datastores .
Unmount Datastore	Unmount the datastore from certain hosts. See Unmount Datastores .
Maintenance Mode	Use datastore maintenance mode. See the <i>vSphere Resource Management</i> documentation.
Configure Storage I/O Control (VMFS)	Enable Storage I/O Control for the VMFS datastore. See the <i>vSphere Resource Management</i> documentation.
Edit Space Reclamation (VMFS)	Change space reclamation settings for the VMFS datastore. See Configure Fixed Reclamation for the Existing VMFS6 Datastore .

Option	Description
Delete Datastore (VMFS)	Remove the VMFS datastore. See Remove VMFS Datastores .
Tags & Custom Attributes	Use tags to encode information about the datastore. See Assign Tags to Datastores .

3. To view specific datastore details, click a selected datastore.
4. Use tabs to access additional information and modify datastore properties.

Tab	Description
Summary	View statistics and configuration for the selected datastore.
Monitor	View alarms, performance data, resource allocation, events, and other status information for the datastore.
Configure	View and modify datastore properties. Menu items that you can see depend on the datastore type.
Permissions	Assign or edit permissions for the selected datastore.
Files	Navigate to the datastore file browser.
Hosts	View hosts where the datastore is mounted.
VMs	View virtual machines that reside on the datastore.

Using Persistent Memory Devices with ESXi

ESXi supports next generation persistent memory devices, also known as Non-Volatile Memory (NVM) devices. These devices combine performance and speed of memory with the persistence of traditional storage. They can retain stored data through reboots or power source failures.

Virtual machines that require high bandwidth, low latency, and persistence can benefit from this technology. Examples include VMs with acceleration databases and analytics workload.

To use persistent memory with your ESXi host, you must be familiar with the following concepts.

PMem Datastore

After you add persistent memory to your ESXi host, the host detects the hardware, and then formats and mounts it as a local PMem datastore. ESXi uses VMFS-L as a file system format. Only one local PMem datastore per host is supported.

Note: When you manage physical persistent memory, make sure to evacuate all VMs from the host and place the host into maintenance mode.

To reduce administrative overhead, the PMem datastore offers a simplified management model. Traditional datastore tasks do not generally apply to the datastore because the host automatically performs all the required operations on the background. As an administrator, you cannot display the datastore in the Datastores view of the vSphere Client, or perform other regular datastore actions. The only operation available to you is monitoring statistics for the PMem datastore.

The PMem datastore is used to store virtual NVDIMM devices and traditional virtual disks of a VM. The VM home directory with the `vmx` and `vmware.log` files cannot be placed on the PMem datastore.

PMem Access Modes

ESXi exposes persistent memory to a VM in two different modes. PMem-aware VMs can have direct access to persistent memory.

Direct-Access Mode

Traditional VMs can use fast virtual disks stored on the PMem datastore.

In this mode, also called virtual PMem (vPMem) mode, a PMem region can be presented to a VM as a virtual non-volatile dual in-line memory module (NVDIMM) module. The VM uses the NVDIMM module as a standard byte-addressable memory that can persist across power cycles.

You can add one or several NVDIMM modules when provisioning the VM.

The VMs must be of the hardware version ESXi 6.7 or later and have a PMem-aware guest OS. The NVDIMM device is compatible with latest guest OSes that support persistent memory, for example, Windows 2016.

Each NVDIMM device is automatically stored on the PMem datastore.

Virtual Disk Mode

This mode, also called virtual PMem disks (vPMemDisk) mode, is available to any traditional VM and supports any hardware version, including all legacy versions. VMs are not required to be PMem-aware. When you use this mode, you create a regular SCSI virtual disk and attach a PMem VM storage policy to the disk. The policy automatically places the disk on the PMem datastore.


PMem Storage Policy

To place the virtual disk on the PMem datastore, you must apply the host-local PMem default storage policy to the disk. The policy is not editable.

The policy can be applied only to virtual disks. Because the VM home directory does not reside on the PMem datastore, make sure to place it on any standard datastore.

After you assign the PMem storage policy to the virtual disk, you cannot change the policy through the **VM Edit Setting** dialog box. To change the policy, migrate or clone the VM.

The following graphic illustrates how the persistent memory components interact.



For information about how to configure and manage VMs with NVDIMMs or virtual persistent memory disks, see the *vSphere Resource Management* documentation and *vSphere Virtual Machine Administration*.

Monitor PMem Datastore Statistics

You can use the vSphere Client and the `esxcli` command to review the capacity of the PMem datastore and some of its other attributes.

However, unlike regular datastores, such as VMFS or vVol, the PMem datastore does not appear in the Datastores view of the vSphere Client. Regular datastore administrative tasks do not apply to it.

Review PMem datastore information.

Option	Description
vSphere Client	- Navigate to the ESXi host and click Summary. - In the Hardware panel, verify that Persistent Memory is displayed and review its capacity.
esxcli command	Use the <code>esxcli storage filesystem list</code> to list the PMem datastore.

Viewing the PMem Datastore

The following sample output appears when you use the `esxcli storage filesystem list` command to list the datastore.

```
# esxcli storage filesystem list
```

Mount Point	Volume Name	UUID	Mounted	Type	Size	Free
-------------	-------------	------	---------	------	------	------

```

-----
-----
/vmfs/volumes/5xxx...    ds01-102      5xxx...      true      VMFS-6  14227079168  12718178304
/vmfs/volumes/59ex...    ds02-102      59ex...      true      VMFS-6  21206401024  19697500160
/vmfs/volumes/59bx...    59bx...      true      vfat      4293591040  4274847744
/vmfs/volumes/pmem:5ax... PMemDS-56ax... pmem:5a0x... true      PMEM      12880707584  11504975872

```

Using ESXi with a SAN

Using ESXi with a SAN improves flexibility, efficiency, and reliability. Using ESXi with a SAN also supports centralized management, failover, and load balancing technologies.

The following are benefits of using ESXi with a SAN:

- You can store data securely and configure multiple paths to your storage, eliminating a single point of failure.
- An ESXi host can access storage devices presented from multiple storage arrays, including arrays from different vendors.
- Using a SAN with ESXi systems extends failure resistance to the server. When you use SAN storage, all applications can instantly be restarted on another host after the failure of the original host.
- You can perform live migration of virtual machines using VMware vMotion.
- Use VMware High Availability (HA) in conjunction with a SAN to restart virtual machines in their last known state on a different server if their host fails.
- Use VMware Fault Tolerance (FT) to replicate protected virtual machines on two different hosts. Virtual machines continue to function without interruption on the secondary host if the primary one fails.
- Use VMware Distributed Resource Scheduler (DRS) to migrate virtual machines from one host to another for load balancing. Because storage is on a shared SAN array, applications continue running seamlessly.
- If you use VMware DRS clusters, put an ESXi host into maintenance mode to have the system migrate all running virtual machines to other ESXi hosts. You can then perform upgrades or other maintenance operations on the original host.

The portability and encapsulation of VMware virtual machines complements the shared nature of this storage. When virtual machines are located on SAN-based storage, you can quickly shut down a virtual machine on one server and power it up on another server, or suspend it on one server and resume operation on another server on the same network. This ability allows you to migrate computing resources while maintaining consistent shared access.

ESXi and SAN Use Cases

When used with a SAN, ESXi can benefit from multiple vSphere features, including Storage vMotion, Distributed Resource Scheduler (DRS), High Availability, and so on.

Using ESXi with a SAN is effective for the following tasks:

Storage consolidation and simplification of storage layout

If you are working with multiple hosts, and each host is running multiple virtual machines, the storage on the hosts is no longer sufficient. You might need to use external storage. The SAN can provide a simple system architecture and other benefits.

Maintenance with zero downtime

When performing ESXi host or infrastructure maintenance, use vMotion to migrate virtual machines to other host. If shared storage is on the SAN, you can perform maintenance without interruptions to the users of the virtual machines. Virtual machine working processes continue throughout a migration.

Load balancing

You can add a host to a DRS cluster, and the host's resources become part of the cluster's resources. The distribution and use of CPU and memory resources for all hosts and virtual machines in the cluster are continuously monitored. DRS compares these

Disaster recovery

metrics to an ideal resource use. The ideal use considers the attributes of the cluster's resource pools and virtual machines, the current demand, and the imbalance target. If needed, DRS performs or recommends virtual machine migrations.

You can use VMware High Availability to configure multiple ESXi hosts as a cluster. The cluster provides rapid recovery from outages and cost-effective high availability for applications running in virtual machines.

Simplified array migrations and storage upgrades

When you purchase new storage systems, use Storage vMotion to perform live migrations of virtual machines from existing storage to their new destinations. You can perform the migrations without interruptions of the virtual machines.

Specifics of Using SAN Storage with ESXi

Using a SAN with an ESXi host differs from a traditional SAN use in several ways.

- You cannot use SAN administration tools to access operating systems of virtual machines that reside on the storage. With traditional tools, you can monitor only the VMware ESXi operating system. You use the vSphere Client to monitor virtual machines.
- The HBA visible to the SAN administration tools is part of the ESXi system, not part of the virtual machine.
- Typically, your ESXi system performs multipathing for you.

Making LUN Decisions

You must plan how to set up storage for your ESXi systems before you format LUNs with VMFS datastores.

When you make your LUN decision, the following considerations apply:

- Each LUN must have the correct RAID level and storage characteristic for the applications running in virtual machines that use the LUN.
- Each LUN must contain only one VMFS datastore.
- If multiple virtual machines access the same VMFS, use disk shares to prioritize virtual machines.

You might want fewer, larger LUNs for the following reasons:

- More flexibility to create virtual machines without asking the storage administrator for more space.
- More flexibility for resizing virtual disks, doing snapshots, and so on.
- Fewer VMFS datastores to manage.

You might want more, smaller LUNs for the following reasons:

- Less wasted storage space.
- Different applications might need different RAID characteristics.
- More flexibility, as the multipathing policy and disk shares are set per LUN.
- Use of Microsoft Cluster Service requires that each cluster disk resource is in its own LUN.
- Better performance because there is less contention for a single volume.

When the storage characterization for a virtual machine is unavailable, it might not be easy to determine the number and size of LUNs to provision. You can experiment using either a predictive or adaptive scheme.

Use the Predictive Scheme to Make LUN Decisions

Experiment using the predictive scheme.

1. Provision several LUNs with different storage characteristics.
2. Create a VMFS datastore on each LUN, labeling each datastore according to its characteristics.
3. Create virtual disks to contain the data for virtual machine applications in the VMFS datastores created on LUNs with the appropriate RAID level for the applications' requirements.
4. Use disk shares to distinguish high-priority from low-priority virtual machines.

Note: Disk shares are relevant only within a given host. The shares assigned to virtual machines on one host have no effect on virtual machines on other hosts.

5. Run the applications to determine whether virtual machine performance is acceptable.

Use the Adaptive Scheme to Make LUN Decisions

You can experiment using the adaptive scheme.

1. Provision a large LUN (RAID 1+0 or RAID 5), with write caching enabled.
2. Create a VMFS on that LUN.
3. Create four or five virtual disks on the VMFS.
4. Run the applications to determine whether disk performance is acceptable.

If performance is acceptable, you can place additional virtual disks on the VMFS. If performance is not acceptable, create a new, large LUN, possibly with a different RAID level, and repeat the process. Use migration so that you do not lose virtual machines data when you recreate the LUN.

Selecting Virtual Machine Locations

When you are working on optimizing performance for your virtual machines, storage location is an important factor. Depending on your storage needs, you might select storage with high performance and high availability, or storage with lower performance.

Storage can be divided into different tiers depending on several factors:

- High Tier. Offers high performance and high availability. Might offer built-in snapshots to facilitate backups and point-in-time (PiT) restorations. Supports replication, full storage processor redundancy, and SAS drives. Uses high-cost spindles.
- Mid Tier. Offers mid-range performance, lower availability, some storage processor redundancy, and SCSI or SAS drives. Might offer snapshots. Uses medium-cost spindles.
- Lower Tier. Offers low performance, little internal storage redundancy. Uses low-end SCSI drives or SATA.

Not all VMs must be on the highest-performance and most-available storage throughout their entire life cycle.

When you decide where to place a virtual machine, the following considerations apply:

- Criticality of the VM
- Performance and availability requirements
- PiT restoration requirements
- Backup and replication requirements

A virtual machine might change tiers throughout its life cycle because of changes in criticality or changes in technology. Criticality is relative and might change for various reasons, including changes in the organization, operational processes, regulatory requirements, disaster planning, and so on.

Third-Party Management Applications

You can use third-party management applications with your ESXi host.

Most SAN hardware is packaged with storage management software. In many cases, this software is a Web application that can be used with any Web browser connected to your network. In other cases, this software typically runs on the storage system or on a single server, independent of the servers that use the SAN for storage.

Use this third-party management software for the following tasks:

- Storage array management, including LUN creation, array cache management, LUN mapping, and LUN security.
- Setting up replication, check points, snapshots, or mirroring.

If you run the SAN management software on a virtual machine, you gain the benefits of a virtual machine, including failover with vMotion and VMware HA. Because of the additional level of indirection, however, the management software might not see the SAN. In this case, you can use an RDM.

Note: Whether a virtual machine can run management software successfully depends on the particular storage system.

SAN Storage Backup Considerations

Having a proper backup strategy is one of the most important aspects of SAN management. In the SAN environment, backups have two goals. The first goal is to archive online data to offline media. This process is repeated periodically for all online data on a time schedule. The second goal is to provide access to offline data for recovery from a problem. For example, database recovery often requires a retrieval of archived log files that are not currently online.

Scheduling a backup depends on several factors:

- Identification of critical applications that require more frequent backup cycles within a given period.
- Recovery point and recovery time goals. Consider how precise your recovery point must be, and how long you are willing to wait for it.
- The rate of change (RoC) associated with the data. For example, if you are using synchronous/asynchronous replication, the RoC affects the amount of bandwidth required between the primary and secondary storage devices.
- Overall impact on a SAN environment, storage performance, and other applications.
- Identification of peak traffic periods on the SAN. Backups scheduled during those peak periods can slow the applications and the backup process.
- Time to schedule all backups within the data center.
- Time it takes to back up an individual application.
- Resource availability for archiving data, such as offline media access.

Include a recovery-time objective for each application when you design your backup strategy. That is, consider the time and resources necessary to perform a backup. For example, if a scheduled backup stores so much data that recovery requires a considerable amount of time, examine the scheduled backup. Perform the backup more frequently, so that less data is backed up at a time and the recovery time decreases.

If an application requires recovery within a certain time frame, the backup process must provide a time schedule and specific data processing to meet the requirement. Fast recovery can require the use of recovery volumes that reside on online storage. This process helps to minimize or eliminate the need to access slow offline media for missing data components.

Using Third-Party Backup Packages

You can use third-party backup solutions to protect system, application, and user data in your virtual machines.

The Storage APIs - Data Protection that VMware offers can work with third-party products. When using the APIs, third-party software can perform backups without loading ESXi hosts with the processing of backup tasks.

The third-party products using the Storage APIs - Data Protection can perform the following backup tasks:

- Perform a full, differential, and incremental image backup and restore of virtual machines.
- Perform a file-level backup of virtual machines that use supported Windows and Linux operating systems.
- Ensure data consistency by using Microsoft Volume Shadow Copy Services (VSS) for virtual machines that run supported Microsoft Windows operating systems.

Because the Storage APIs - Data Protection use the snapshot capabilities of VMFS, backups do not require that you stop virtual machines. These backups are nondisruptive, can be performed at any time, and do not need extended backup windows.

For information about the Storage APIs - Data Protection and integration with backup products, see the [Broadcom KB article 1021175](#).

Using ESXi with Fibre Channel SAN

ESXi supports Fibre Channel (FC), a storage protocol that the SAN uses to transfer data traffic from hosts to shared storage. This section provides introductory information about how to use ESXi with Fibre Channel SAN. For more information, check your vendor documentation.

Fibre Channel SAN Concepts

If you are a vSphere administrator planning to set up hosts to work with SANs, you must have a working knowledge of SAN concepts. You can find information about SANs in print and on the Internet. Because this industry changes constantly, check these resources frequently.

Storage Area Network (SAN)

A storage area network (SAN) is a specialized high-speed network that connects host servers to high-performance storage subsystems. The SAN components include host bus adapters (HBAs) in the host servers, switches that help route storage traffic, cables, storage processors (SPs), and storage disk arrays.

SAN Fabric

A SAN topology with at least one switch present on the network forms a SAN fabric.

Fibre Channel (FC) Protocol

To transfer traffic from host servers to shared storage, the SAN uses the Fibre Channel (FC) protocol that packages SCSI or NVMe commands into Fibre Channel frames.

Zoning

To restrict server access to storage arrays not allocated to that server, the SAN uses zoning. Typically, zones are created for each group of servers that access a shared group of storage devices and LUNs. Zones define which HBAs can connect to which SPs. Devices outside a zone are not visible to the devices inside the zone.

Zoning has the following effects:

- Reduces the number of targets and LUNs presented to a host.
- Controls and isolates paths in a fabric.
- Can prevent non-ESXi systems from accessing a particular storage system, and from possibly destroying VMFS data.

- Can be used to separate different environments, for example, a test from a production environment.

With ESXi hosts, use a single-initiator zoning or a single-initiator-single-target zoning. The latter is a preferred zoning practice. Using the more restrictive zoning prevents problems and misconfigurations that can occur on the SAN.

For detailed instructions and best zoning practices, contact storage array or switch vendors.

LUN Masking

Zoning is similar to LUN masking, which is commonly used for permission management. LUN masking is a process that makes a LUN available to some hosts and unavailable to other hosts.

Multipathing

When transferring data between the host server and storage, the SAN uses a technique known as multipathing. Multipathing allows you to have more than one physical path from the ESXi host to a LUN on a storage system.

Path Failover

Generally, a single path from a host to a LUN consists of an HBA, switch ports, connecting cables, and the storage controller port. If any component of the path fails, the host selects another available path for I/O. The process of detecting a failed path and switching to another is called path failover.

Ports in Fibre Channel SAN

In the context of this document, a port is the connection from a device into the SAN. Each node in the SAN, such as a host, a storage device, or a fabric component has one or more ports that connect it to the SAN. Ports are identified in a number of ways.

WWPN (World Wide Port Name)

A globally unique identifier for a port that allows certain applications to access the port. The FC switches discover the WWPN of a device or host and assign a port address to the device.

Port_ID (or port address)

Within a SAN, each port has a unique port ID that serves as the FC address for the port. This unique ID enables routing of data through the SAN to that port. The FC switches assign the port ID when the device logs in to the fabric. The port ID is valid only while the device is logged on.

When N-Port ID Virtualization (NPIV) is used, a single FC HBA port (N-port) can register with the fabric by using several WWPNs. This method allows an N-port to claim multiple fabric addresses, each of which appears as a unique entity. When ESXi hosts use a SAN, these multiple, unique identifiers allow the assignment of WWNs to individual virtual machines as part of their configuration.

Fibre Channel Storage Array Types

ESXi supports different storage systems and arrays. They generally fall into these categories.

Active-active storage system

Supports access to the LUNs simultaneously through all the storage ports that are available without significant performance degradation. All the paths are active, unless a path fails.

Active-passive storage system

A system in which one storage processor is actively providing access to a given LUN. The other processors act as a backup for the LUN and can be actively providing access to other LUN I/O. I/O can be successfully sent only to an active port for a given LUN.

Asymmetrical storage system

If access through the active storage port fails, one of the passive storage processors can be activated by the servers accessing it. Supports Asymmetric Logical Unit Access (ALUA). ALUA-compliant storage systems provide different levels of access per port. With ALUA, the host can determine the states of target ports and prioritize paths. The host uses some of the active paths as primary, and uses others as secondary.

How Virtual Machines Access Data on a Fibre Channel SAN

ESXi stores a virtual machine's disk files within a VMFS datastore that resides on a SAN storage device. When virtual machine guest operating systems send SCSI or NVMe commands to their virtual disks, the SCSI or NVMe virtualization layer translates these commands to VMFS file operations.

When a virtual machine interacts with its virtual disk stored on a SAN, the following process takes place:

1. When the guest operating system in a virtual machine reads or writes to a SCSI or NVMe disk, it sends SCSI or NVMe commands to the virtual disk.
2. Device drivers in the virtual machine's operating system communicate with the virtual SCSI or NVMe controllers.
3. The virtual SCSI or NVMe controller forwards the command to the VMkernel.
4. The VMkernel performs the following tasks.
 - a. Locates the appropriate virtual disk file in the VMFS volume.
 - b. Maps the requests for the blocks on the virtual disk to blocks on the appropriate physical device.
 - c. Sends the modified I/O request from the device driver in the VMkernel to the physical HBA.
5. The physical HBA performs the following tasks.
 - a. Packages the I/O request according to the rules of the FC protocol.
 - b. Transmits the request to the SAN.
6. Depending on a port the HBA uses to connect to the fabric, one of the SAN switches receives the request. The switch routes the request to the appropriate storage device.

Setting Up ESXi Fibre Channel Environment

This topic provides an overview of setup steps that you need to follow when configuring your Fibre Channel SAN environment to work with ESXi. In preparation for configuring your Fibre Channel SAN and setting up the ESXi system, review the requirements and recommendations.

ESXi Fibre Channel SAN Requirements

- Make sure that ESXi systems support the SAN storage hardware and firmware combinations you use. For an up-to-date list, see the [VMware Compatibility Guide](#).
- Configure your system to have only one VMFS volume per LUN.
- Unless you are using diskless servers, do not set up the diagnostic partition on a SAN LUN. If you use diskless servers that boot from a SAN, a shared diagnostic partition is appropriate.
- Use RDMS to access raw disks. For information, see [Raw Device Mapping in vSphere](#).
- Make sure that the storage device driver specifies a large enough queue. You can set the queue depth for the physical HBA during a system setup.
- VMtools increases the default disk timeout SCSI `TimeoutValue` parameter to 180 seconds. On virtual machines without VMtools, increase the value of the SCSI `TimeoutValue` parameter. With this parameter set up, the VMs can better tolerate delayed I/O resulting from a path failover. For information, see [Understanding Multipathing and Failover in the ESXi Environment](#) or the [KB 1009465](#).

ESXi Fibre Channel SAN Restrictions

When you use ESXi with a SAN, certain restrictions apply.

- ESXi does not support FC connected tape devices.
- You cannot use multipathing software inside a virtual machine to perform I/O load balancing to a single physical LUN. However, when your Microsoft Windows virtual machine uses dynamic disks, this restriction does not apply. For information about configuring dynamic disks, see [Set Up Dynamic Disk Mirroring for Virtual Machines on vSphere VMFS](#).

Setting LUN Allocations

This topic provides general information about how to allocate LUNs when your ESXi works with SAN.

When you set LUN allocations, be aware of the following points:

Storage provisioning

To ensure that the ESXi system recognizes the LUNs at startup time, provision all LUNs to the appropriate HBAs before you connect the SAN to the ESXi system.

Provision all LUNs to all ESXi HBAs at the same time. HBA failover works only if all HBAs see the same LUNs.

vMotion and VMware DRS

When you use vCenter Server and vMotion or DRS, make sure that the LUNs for the virtual machines are provisioned to all ESXi hosts. This action provides the most ability to move virtual machines.

Active-active compared to active-passive arrays

When you use vMotion or DRS with an active-passive SAN storage device, make sure that all ESXi systems have consistent paths to all storage processors. Not doing so can cause path thrashing when a vMotion migration occurs.

For active-passive storage arrays not listed in Storage/SAN Compatibility, VMware does not support storage port failover. In those cases, you must connect the server to the active port on the storage array. This configuration ensures that the LUNs are presented to the ESXi host.

Setting Fibre Channel HBAs

Typically, FC HBAs that you use on your ESXi host work correctly with the default configuration settings.

You should follow the configuration guidelines provided by your storage array vendor. During FC HBA setup, consider the following issues.

- Do not mix FC HBAs from different vendors in a single host. Having different models of the same HBA is supported, but a single LUN cannot be accessed through two different HBA types, only through the same type.
- Ensure that the firmware level on each HBA is the same.
- Set the timeout value for detecting a failover. To ensure optimal performance, do not change the default value.
- ESXi supports 32 Gbps end-to-end Fibre Channel connectivity.

Configuring Fibre Channel over Ethernet

To access Fibre Channel storage, an ESXi host can use the Fibre Channel over Ethernet (FCoE) protocol. The FCoE protocol encapsulates Fibre Channel frames into Ethernet frames. As a result, your host does not need special Fibre Channel links to connect to Fibre Channel storage. The host can use 10 Gbit lossless Ethernet to deliver Fibre Channel traffic.

Fibre Channel over Ethernet Adapters

VMware supports hardware FCoE adapters, a specialized offloaded Converged Network Adapters (CNAs) that contain network and Fibre Channel functionalities on the same card.

When such adapter is installed, your host detects and can use both CNA components. In the vSphere Client, the networking component appears as a standard network adapter (vmnic) and the Fibre Channel component as a FCoE adapter (vmhba).

For information on adapters that can be used with VMware FCoE, see the *VMware Compatibility Guide*.

Setup Workflow

This topic provides an overview of installation and setup steps that you need to follow when configuring your Fibre Channel SAN environment to work with ESXi.

Follow these steps to configure your ESXi SAN environment.

1. Design your SAN if it is not already configured. Most existing SANs require only minor modification to work with ESXi.
2. Check that all SAN components meet requirements.
3. Perform any necessary storage array modification.
Most vendors have vendor-specific documentation for setting up a SAN to work with VMware ESXi.
4. Set up the HBAs for the hosts you have connected to the SAN.
5. Install ESXi on the hosts.
6. Create virtual machines and install guest operating systems.
7. (Optional) Set up your system for VMware HA failover or for using Microsoft Clustering Services.
8. Upgrade or modify your environment as needed.

Booting ESXi from Fibre Channel SAN

When you set up your ESXi host to boot from a SAN, your host's boot image is stored on one or more LUNs in the SAN storage system. When the host starts, it boots from the LUN on the SAN rather than from its local disk.

ESXi supports booting through a Fibre Channel host bus adapter (HBA) or a Fibre Channel over Ethernet (FCoE) converged network adapter (CNA).

Boot from SAN Benefits

Boot from SAN can provide numerous benefits to your ESXi environment. However, in certain cases, boot from SAN is not compatible with your hosts. Before you set up your system for boot from SAN, decide whether it is appropriate for your environment.

CAUTION: When you use boot from SAN with multiple ESXi hosts, each host must have its own boot LUN. If you configure multiple hosts to share the boot LUN, ESXi image corruption might occur.

If you use boot from SAN, the benefits for your environment include the following:

- Cheaper servers. Servers can be more dense and run cooler without internal storage.
- Easier server replacement. You can replace servers and have the new server point to the old boot location.
- Less wasted space. Servers without local disks often take up less space.
- Easier backup processes. You can back up the system boot images in the SAN as part of the overall SAN backup procedures. Also, you can use advanced array features such as snapshots on the boot image.
- Improved management. Creating and managing the operating system image is easier and more efficient.
- Better reliability. You can access the boot disk through multiple paths, which protects the disk from being a single point of failure.

Requirements and Considerations when Booting from Fibre Channel SAN

Your ESXi boot configuration must meet specific requirements.

Table 209: Boot from SAN Requirements

Requirement	Description
ESXi system requirements	Follow vendor recommendations for the server booting from a SAN.
Adapter requirements	Configure the adapter, so it can access the boot LUN. See your vendor documentation.
Access control	- Each host must have access to its own boot LUN only, not the boot LUNs of other hosts. Use storage system software to make sure that the host accesses only the designated LUNs. - Multiple servers can share a diagnostic partition. You can use array-specific LUN masking to achieve this configuration.
Multipathing support	Multipathing to a boot LUN on active-passive arrays is not supported because the BIOS does not support multipathing and is unable to activate a standby path.
SAN considerations	If the array is not certified for a direct connect topology, the SAN connections must be through a switched topology. If the array is certified for the direct connect topology, the SAN connections can be made directly to the array. Boot from SAN is supported for both switched topology and direct connect topology.
Hardware-specific considerations	If you are running an IBM eServer BladeCenter and use boot from SAN, you must deactivate IDE drives on the blades.

Getting Ready to Boot ESXi from SAN

When you prepare your ESXi host to boot from a SAN, you perform several tasks.

This workflow describes the generic boot-from-SAN enablement process on the rack-mounted servers. For information on enabling the boot from SAN option on blade servers or other types of hosts, refer to partner documentation.

For more information, see <https://www.vmware.com/resources/compatibility/search.php>.

Configure SAN Components and Storage System

Before you set up your ESXi host to boot from a SAN LUN, configure SAN components and a storage system.

Because configuring the SAN components is vendor-specific, refer to the product documentation for each item.

1. Connect network cable, referring to any cabling guide that applies to your setup.
Check the switch wiring, if there is any.
2. Configure the storage array.
 - a) From the SAN storage array, make the ESXi host visible to the SAN. This process is often called creating an object.
 - b) From the SAN storage array, set up the host to have the WWPNs of the host's adapters as port names or node names.
 - c) Create and assign LUNs.
 - d) Record the IP addresses of the switches and storage arrays.
 - e) Record the WWPN for each SP.

CAUTION: If you use a scripted installation process to install ESXi in boot from SAN mode, take special steps to avoid unintended data loss.

Configure Storage Adapter to Boot from SAN

When you set up your host to boot from SAN, you enable the boot adapter in the host BIOS. You then configure the boot adapter to initiate a primitive connection to the target boot LUN.

Determine the WWPN for the storage adapter.

Configure the storage adapter to boot from SAN.

Because configuring boot adapters is vendor-specific, consult your vendor documentation.

Set Up Your System to Boot from Installation Media

When setting up your host to boot from SAN, you first boot the host from the VMware installation media. To boot from the installation media, change the system boot sequence in the BIOS setup.

Because changing the boot sequence in the BIOS is vendor-specific, refer to vendor documentation for instructions. The following procedure explains how to change the boot sequence on an IBM host.

1. Power on your system and enter the system BIOS Configuration/Setup Utility.
2. Select **Startup Options** and press Enter.
3. Select **Startup Sequence Options** and press Enter.
4. Change the **First Startup Device** to **[CD-ROM]**.

You can now install ESXi.

Best Practices for Fibre Channel Storage with ESXi

When you use ESXi with Fibre Channel SAN, follow recommendations to avoid performance problems.

The vSphere Client offers extensive facilities for collecting performance information. The information is graphically displayed and frequently updated.

You can also use the `resxtop` or `esxtop` command-line utilities. The utilities provide a detailed look at how ESXi uses resources. For more information, see the *vSphere Resource Management* documentation.

Check with your storage representative if your storage system supports Storage API - Array Integration hardware acceleration features. If it does, refer to your vendor documentation to enable hardware acceleration support on the storage system side. For more information, see [Storage Hardware Acceleration in vSphere](#).

Preventing Fibre Channel SAN Problems

When you use ESXi with a Fibre Channel SAN, follow specific guidelines to avoid SAN problems.

To prevent problems with your SAN configuration, observe these tips:

- Place only one VMFS datastore on each LUN.
- Do not change the path policy the system sets for you unless you understand the implications of making such a change.
- Document everything. Include information about zoning, access control, storage, switch, server and FC HBA configuration, software and firmware versions, and storage cable plan.
- Plan for failure:
 - Make several copies of your topology maps. For each element, consider what happens to your SAN if the element fails.
 - Verify different links, switches, HBAs, and other elements to ensure that you did not miss a critical failure point in your design.

- Ensure that the Fibre Channel HBAs are installed in the correct slots in the host, based on slot and bus speed. Balance PCI bus load among the available buses in the server.
- Become familiar with the various monitor points in your storage network, at all visibility points, including host's performance charts, FC switch statistics, and storage performance statistics.
- Be cautious when changing IDs of the LUNs that have VMFS datastores being used by your ESXi host. If you change the ID, the datastore becomes inactive and its virtual machines fail. Resignature the datastore to make it active again. See [vSphere VMFS Datastore Copies and Datastore Resignaturing](#).
After you change the ID of the LUN, rescan the storage to reset the ID on your host. For information on using the rescan, see [Rescan Operations for ESXi Storage](#).

Deactivate Automatic ESXi Host Registration

Certain storage arrays require that ESXi hosts register with the arrays. ESXi performs automatic host registration by sending the host's name and IP address to the array. If you prefer to perform manual registration using storage management software, deactivate the ESXi auto-registration feature.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. Under Advanced System Settings, select the **Disk.EnableNaviReg** parameter and click the **Edit** icon.
5. Change the value to 0.

This operation deactivates the automatic host registration that is activated by default.

Optimizing Fibre Channel SAN Storage Performance

Several factors contribute to optimizing a typical SAN environment.

If the environment is properly configured, the SAN fabric components (particularly the SAN switches) are only minor contributors because of their low latencies relative to servers and storage arrays. Make sure that the paths through the switch fabric are not saturated, that is, that the switch fabric is running at the highest throughput.

Storage Array Performance

Storage array performance is one of the major factors contributing to the performance of the entire SAN environment.

If you encounter any problems with storage array performance, consult your storage array vendor documentation for any relevant information.

To improve the array performance in the vSphere environment, follow these general guidelines:

- When assigning LUNs, remember that several hosts might access the LUN, and that several virtual machines can run on each host. One LUN used by a host can service I/O from many different applications running on different operating systems. Because of this diverse workload, the RAID group containing the ESXi LUNs typically does not include LUNs used by other servers that are not running ESXi.
- Make sure that the read/write caching is available.
- SAN storage arrays require continual redesign and tuning to ensure that I/O is load-balanced across all storage array paths. To meet this requirement, distribute the paths to the LUNs among all the SPs to provide optimal load-balancing. Close monitoring indicates when it is necessary to rebalance the LUN distribution.
Tuning statically balanced storage arrays is a matter of monitoring the specific performance statistics, such as I/O operations per second, blocks per second, and response time. Distributing the LUN workload to spread the workload across all the SPs is also important.

Note: Dynamic load-balancing is not currently supported with ESXi.

Server Performance with Fibre Channel

You must consider several factors to ensure optimal server performance.

Each server application must have access to its designated storage with the following conditions:

- High I/O rate (number of I/O operations per second)
- High throughput (megabytes per second)
- Minimal latency (response times)

Because each application has different requirements, you can meet these goals by selecting an appropriate RAID group on the storage array.

To achieve performance goals, follow these guidelines:

- Place each LUN on a RAID group that provides the necessary performance levels. Monitor the activities and resource use of other LUNs in the assigned RAID group. A high-performance RAID group that has too many applications doing I/O to it might not meet performance goals required by an application running on the ESXi host.
- Ensure that each host has enough HBAs to increase throughput for the applications on the host for the peak period. I/O spread across multiple HBAs provides faster throughput and less latency for each application.
- To provide redundancy for a potential HBA failure, make sure that the host is connected to a dual redundant fabric.
- When allocating LUNs or RAID groups for ESXi systems, remember that multiple operating systems use and share that resource. The LUN performance required by the ESXi host might be much higher than when you use regular physical machines. For example, if you expect to run four I/O intensive applications, allocate four times the performance capacity for the ESXi LUNs.
- When you use multiple ESXi systems in with vCenter Server, the performance requirements for the storage subsystem increase correspondingly.
- The number of outstanding I/Os needed by applications running on the ESXi system must match the number of I/Os the HBA and storage array can handle.

Using Fibre Channel NPIV with vSphere Virtual Machines

Virtual N-Port ID Virtualization (NPIV) is an ANSI T11 standard that describes how a single Fibre Channel HBA port can register with the fabric using several worldwide port names (WWPNs). This allows a fabric-attached N-port to claim multiple fabric addresses. Each address appears as a unique entity on the Fibre Channel fabric. You can configure your vSphere virtual machines to use Fibre Channel NPIV.

How NPIV-Based LUN Access Works

NPIV enables a single FC HBA port to register several unique World Wide Name (WWN) identifiers with the fabric, each of which can be assigned to an individual virtual machine. When using NPIV, a SAN administrator can monitor and route storage access per a virtual machine.

Only virtual machines with RDMs can have WWN assignments, and they use these assignments for all RDM traffic.

When a virtual machine has a WWN assigned to it, the virtual machine's configuration file (.vmx) is updated to include a WWN pair. The WWN pair consists of a World Wide Port Name (WWPN) and a World Wide Node Name (WWNN). When that virtual machine is powered on, the VMkernel creates a virtual port (VPORT) on the physical HBA which is used to access the LUN. The VPORT is a virtual HBA that appears to the FC fabric as a physical HBA. As its unique identifier, the VPORT uses the WWN pair that was assigned to the virtual machine.

Each VPORT is specific to the virtual machine. The VPORT is destroyed on the host and no longer appears to the FC fabric when the virtual machine is powered off. When a virtual machine is migrated from one host to another, the VPORT closes on the first host and opens on the destination host.

When virtual machines do not have WWN assignments, they access storage LUNs with the WWNs of their host's physical HBAs.

Requirements for Using NPIV

If you plan to enable NPIV on your virtual machines, you should be aware of certain requirements.

- NPIV can be used only for virtual machines with RDM disks. Virtual machines with regular virtual disks use the WWNs of the host's physical HBAs.
- HBAs on your host must support NPIV.
For information, see the *VMware Compatibility Guide* and refer to your vendor documentation.
 - Use HBAs of the same type. VMware does not support heterogeneous HBAs on the same host accessing the same LUNs.
 - If a host uses multiple physical HBAs as paths to the storage, zone all physical paths to the virtual machine. This is required to support multipathing even though only one path at a time will be active.
 - Make sure that physical HBAs on the host can detect all LUNs that are to be accessed by NPIV-enabled virtual machines running on that host.
- The switches in the fabric must be NPIV-aware.
- When configuring a LUN for NPIV access at the storage level, make sure that the NPIV LUN number and NPIV target ID match the physical LUN and Target ID.
- Zone the NPIV WWPNS so that they connect to all storage systems the cluster hosts can access, even if the VM does not use the storage. If you add any new storage systems to a cluster with one or more NPIV-enabled VMs, add the new zones, so the NPIV WWPNS can detect the new storage system target ports.

NPIV Capabilities and Limitations

Learn about specific capabilities and limitations of the use of NPIV with ESXi.

ESXi with NPIV supports the following items:

- NPIV supports vMotion. When you use vMotion to migrate a virtual machine it retains the assigned WWN.
If you migrate an NPIV-enabled virtual machine to a host that does not support NPIV, VMkernel reverts to using a physical HBA to route the I/O.
- If your FC SAN environment supports concurrent I/O on the disks from an active-active array, the concurrent I/O to two different NPIV ports is also supported.

When you use ESXi with NPIV, the following limitations apply:

- Because the NPIV technology is an extension to the FC protocol, it requires an FC switch and does not work on the direct attached FC disks.
- When you clone a virtual machine or template with a WWN assigned to it, the clones do not retain the WWN.
- NPIV does not support Storage vMotion.
- Deactivating and then re-activating the NPIV capability on an FC switch while virtual machines are running can cause an FC link to fail and I/O to stop.

Configure or Modify WWN Assignments

Assign WWN settings to a virtual machine. You can later modify the WWN assignments.

- Before configuring WWN, ensure that the ESXi host can access the storage LUN access control list (ACL) configured on the array side.
- If you want to edit the existing WWNs, power off the virtual machine.

You can create from 1 to 16 WWN pairs, which can be mapped to the first 1 to 16 physical FC HBAs on the host.

Typically, you do not need to change existing WWN assignments on your virtual machine. In certain circumstances, for example, when manually assigned WWNs are causing conflicts on the SAN, you might need to change or remove WWNs.

1. Right-click the virtual machine in the inventory and select **Edit Settings**.
2. Click **VM Options** and expand **Fibre Channel NPIV**.
3. Create or edit the WWN assignments by selecting one of the following options:

Option	Description
Temporarily disable NPIV for this virtual machine	Deactivate but do not remove the existing WWN assignments for the virtual machine.
Leave unchanged	Retain the existing WWN assignments. The read-only WWN assignments section displays the node and port values of any existing WWN assignments.
Generate new WWNs	Generate new WWNs, overwriting any existing WWNs. The WWNs of the HBA are not affected. Specify the number of WWNNs and WWPNS. A minimum of two WWPNS are required to support failover with NPIV. Typically only one WWNN is created for each virtual machine.
Remove WWN assignment	Remove the WWNs assigned to the virtual machine. The virtual machine uses the HBA WWNs to access the storage LUN.

4. Click **OK** to save your changes.

Register newly created WWNs in the fabric.

Using ESXi with iSCSI SAN


ESXi can connect to external SAN storage using the Internet SCSI (iSCSI) protocol. In addition to traditional iSCSI, ESXi also supports iSCSI Extensions for RDMA (iSER).

With the iSER protocol, the host can use the same iSCSI framework, but replaces the TCP/IP transport with the Remote Direct Memory Access (RDMA) transport.

About iSCSI SAN

iSCSI SAN uses Ethernet connections between hosts and high-performance storage subsystems.

On the host side, the iSCSI SAN components include iSCSI host bus adapters (HBAs) or Network Interface Cards (NICs). The iSCSI network also includes switches and routers that transport the storage traffic, cables, storage processors (SPs), and storage disk systems.



The iSCSI SAN uses a client-server architecture.

The client, called iSCSI initiator, operates on your ESXi host. It initiates iSCSI sessions by issuing SCSI commands and transmitting them, encapsulated into the iSCSI protocol, to an iSCSI server. The server is known as an iSCSI target. Typically, the iSCSI target represents a physical storage system on the network.

The target can also be a virtual iSCSI SAN, for example, an iSCSI target emulator running in a virtual machine. The iSCSI target responds to the initiator's commands by transmitting required iSCSI data.

iSCSI Multipathing

When transferring data between the host server and storage, the SAN uses a technique known as multipathing. With multipathing, your ESXi host can have more than one physical path to a LUN on a storage system.

Generally, a single path from a host to a LUN consists of an iSCSI adapter or NIC, switch ports, connecting cables, and the storage controller port. If any component of the path fails, the host selects another available path for I/O. The process of detecting a failed path and switching to another is called path failover.

For more information on multipathing, see [Understanding Multipathing and Failover in the ESXi Environment](#).

Nodes and Ports in the iSCSI SAN

A single discoverable entity on the iSCSI SAN, such as an initiator or a target, represents an iSCSI node.

Each node has a node name. ESXi uses several methods to identify a node.

IP Address

Each iSCSI node can have an IP address associated with it so that routing and switching equipment on your network can establish the connection between the host and storage. This address is like the

iSCSI Name

IP address that you assign to your computer to get access to your company's network or the Internet.

A worldwide unique name for identifying the node. iSCSI uses the iSCSI Qualified Name (IQN) and Extended Unique Identifier (EUI).

By default, ESXi generates unique iSCSI names for your iSCSI initiators, for example, `iqn.1998-01.com.vmware:iscsitestox-68158ef2`. Usually, you do not have to change the default value, but if you do, make sure that the new iSCSI name you enter is worldwide unique.

iSCSI Alias

A more manageable name for an iSCSI device or port used instead of the iSCSI name. iSCSI aliases are not unique and are intended to be a friendly name to associate with a port.

Each node has one or more ports that connect it to the SAN. iSCSI ports are end-points of an iSCSI session.

iSCSI Naming Conventions

iSCSI uses a special unique name to identify an iSCSI node, either target or initiator.

iSCSI names are formatted in two different ways. The most common is the IQN format.

For more details on iSCSI naming requirements and string profiles, see RFC 3721 and RFC 3722 on the IETF website.

iSCSI Qualified Name Format

The iSCSI Qualified Name (IQN) format takes the form

`iqn.yyyy-mm.naming-authority:unique name`, where:

- *yyyy-mm* is the year and month when the naming authority was established.
- *naming-authority* is the reverse syntax of the Internet domain name of the naming authority. For example, the `iscsi.vmware.com` naming authority can have the iSCSI qualified name form of `iqn.1998-01.com.vmware.iscsi`. The name indicates that the `vmware.com` domain name was registered in January of 1998, and `iscsi` is a subdomain, maintained by `vmware.com`.
- *unique name* is any name you want to use, for example, the name of your host. The naming authority must make sure that any names assigned following the colon are unique, such as:
 - `iqn.1998-01.com.vmware.iscsi:name1`
 - `iqn.1998-01.com.vmware.iscsi:name2`
 - `iqn.1998-01.com.vmware.iscsi:name999`

Enterprise Unique Identifier Format

The Enterprise Unique Identifier (EUI) format takes the form

`eui.16_hex_digits`.

For example, `eui.0123456789ABCDEF`.

The 16-hexadecimal digits are text representations of a 64-bit number of an IEEE EUI (extended unique identifier) format. The top 24 bits are a company ID that IEEE registers with a particular company. The remaining 40 bits are assigned by the entity holding that company ID and must be unique.

iSCSI Initiators

To access iSCSI targets, your ESXi host uses iSCSI initiators.

The initiator is a software or hardware installed on your ESXi host. The iSCSI initiator originates communication between your host and an external iSCSI storage system and sends data to the storage system.

In the ESXi environment, iSCSI adapters configured on your host play the role of initiators. ESXi supports several types of iSCSI adapters.

For information on configuring and using iSCSI adapters, see [Configuring iSCSI and iSER Adapters and Storage with ESXi](#).

Software iSCSI Adapter

A software iSCSI adapter is a VMware code built into the VMkernel. Using the software iSCSI adapter, your host can connect to the iSCSI storage device through standard network adapters. The software iSCSI adapter handles iSCSI processing while communicating with the network adapter. With the software iSCSI adapter, you can use iSCSI technology without purchasing specialized hardware.

Hardware iSCSI Adapter

A hardware iSCSI adapter is a third-party adapter that offloads iSCSI and network processing from your host. Hardware iSCSI adapters are divided into categories.

- **Dependent Hardware iSCSI Adapter.** Depends on VMware networking, and iSCSI configuration and management interfaces provided by VMware.

This type of adapter can be a card that presents a standard network adapter and iSCSI offload functionality for the same port. The iSCSI offload functionality depends on the host's network configuration to obtain the IP, MAC, and other parameters used for iSCSI sessions. An example of a dependent adapter is the iSCSI licensed Broadcom 5709 NIC.

- **Independent Hardware iSCSI Adapter.** Implements its own networking and iSCSI configuration and management interfaces.

Typically, an independent hardware iSCSI adapter is a card that either presents only iSCSI offload functionality or iSCSI offload functionality and standard NIC functionality. The iSCSI offload functionality has independent configuration management that assigns the IP, MAC, and other parameters used for the iSCSI sessions. An example of an independent adapter is the QLogic QLA4052 adapter.

Hardware iSCSI adapters might need to be licensed. Otherwise, they might not appear in the client or vSphere CLI. Contact your vendor for licensing information.

Using iSER Protocol with ESXi

In addition to traditional iSCSI, ESXi supports the iSCSI Extensions for RDMA (iSER) protocol. When the iSER protocol is enabled, the iSCSI framework on the ESXi host can use the Remote Direct Memory Access (RDMA) transport instead of TCP/IP.

The traditional iSCSI protocol carries SCSI commands over a TCP/IP network between an iSCSI initiator on a host and an iSCSI target on a storage device. The iSCSI protocol encapsulates the commands and assembles that data in packets for the TCP/IP layer. When the data arrives, the iSCSI protocol disassembles the TCP/IP packets, so that the SCSI commands can be differentiated and delivered to the storage device.

iSER differs from traditional iSCSI as it replaces the TCP/IP data transfer model with the Remote Direct Memory Access (RDMA) transport. Using the direct data placement technology of the RDMA, the iSER protocol can transfer data

directly between the memory buffers of the ESXi host and storage devices. This method eliminates unnecessary TCP/IP processing and data copying, and can also reduce latency and the CPU load on the storage device.

In the iSER environment, iSCSI works exactly as before, but uses an underlying RDMA fabric interface instead of the TCP/IP-based interface.


Because the iSER protocol preserves the compatibility with iSCSI infrastructure, the process of enabling iSER on the ESXi host is similar to the iSCSI process. See [Configure iSER with ESXi](#).

Establishing iSCSI Connections

In the ESXi context, the term target identifies a single storage unit that your host can access. The terms storage device and LUN describe a logical volume that represents storage space on a target. Typically, the terms device and LUN, in the ESXi context, mean a SCSI volume presented to your host from a storage target and available for formatting.

Different iSCSI storage vendors present storage to hosts in different ways. Some vendors present multiple LUNs on a single target. Others present multiple targets with one LUN each.

Figure 95: Target Compared to LUN Representations



In these examples, three LUNs are available in each of these configurations. In the first case, the host detects one target but that target has three LUNs that can be used. Each of the LUNs represents individual storage volume. In the second case, the host detects three different targets, each having one LUN.

Host-based iSCSI initiators establish connections to each target. Storage systems with a single target containing multiple LUNs have traffic to all the LUNs on a single connection. With a system that has three targets with one LUN each, the host uses separate connections to the three LUNs.

This information is useful when you are trying to aggregate storage traffic on multiple connections from the host with multiple iSCSI adapters. You can set the traffic for one target to a particular adapter, and use a different adapter for the traffic to another target.

iSCSI Storage System Types

The types of ESXi storage that your host supports include active-active, active-passive, and ALUA-compliant.

Active-active storage system

Supports access to the LUNs simultaneously through all the storage ports that are available without significant performance degradation. All the paths are always active, unless a path fails.

Active-passive storage system

A system in which one storage processor is actively providing access to a given LUN. The other processors act as a backup for the LUN and can be actively providing access to other LUN I/O. I/O can be successfully sent only to an active port for a given LUN. If access through the active storage port fails, one of the passive storage processors can be activated by the servers accessing it.

Asymmetrical storage system

Supports Asymmetric Logical Unit Access (ALUA). ALUA-compliant storage systems provide different levels of access per

Virtual port storage system

port. With ALUA, hosts can determine the states of target ports and prioritize paths. The host uses some of the active paths as primary and others as secondary.

Supports access to all available LUNs through a single virtual port. Virtual port storage systems are active-active storage devices, but hide their multiple connections through a single port. ESXi multipathing does not make multiple connections from a specific port to the storage by default. Some storage vendors supply session managers to establish and manage multiple connections to their storage. These storage systems handle port failovers and connection balancing transparently. This capability is often called transparent failover.

Discovery, Authentication, and Access Control

You can use several mechanisms to discover your storage and to limit access to it.

You must configure your host and the iSCSI storage system to support your storage access control policy.

Discovery

A discovery session is part of the iSCSI protocol. It returns the set of targets you can access on an iSCSI storage system. The two types of discovery available on ESXi are dynamic and static. Dynamic discovery obtains a list of accessible targets from the iSCSI storage system. Static discovery can access only a particular target by target name and address.

For more information, see [Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host](#).

Authentication

iSCSI storage systems authenticate an initiator by a name and key pair. ESXi supports the CHAP authentication protocol. To use CHAP authentication, the ESXi host and the iSCSI storage system must have CHAP enabled and have common credentials.

For information on enabling CHAP, see [Configuring CHAP Parameters for iSCSI or iSER Storage Adapters on ESXi Host](#).

Access Control

Access control is a policy set up on the iSCSI storage system. Most implementations support one or more of three types of access control:

- By initiator name
- By IP address
- By the CHAP protocol

Only initiators that meet all rules can access the iSCSI volume.

Using only CHAP for access control can slow down rescans because the ESXi host can discover all targets, but then fails at the authentication step. iSCSI rescans work faster if the host discovers only the targets it can authenticate.

How Virtual Machines Access Data on an iSCSI SAN

ESXi stores a virtual machine's disk files within a VMFS datastore that resides on a SAN storage device. When virtual machine guest operating systems send SCSI commands to their virtual disks, the SCSI virtualization layer translates these commands to VMFS file operations.

When a virtual machine interacts with its virtual disk stored on a SAN, the following process takes place:

1. When the guest operating system in a virtual machine reads or writes to SCSI disk, it sends SCSI commands to the virtual disk.
2. Device drivers in the virtual machine's operating system communicate with the virtual SCSI controllers.
3. The virtual SCSI controller forwards the commands to the VMkernel.
4. The VMkernel performs the following tasks.
 - a. Locates an appropriate virtual disk file in the VMFS volume.
 - b. Maps the requests for the blocks on the virtual disk to blocks on the appropriate physical device.
 - c. Sends the modified I/O request from the device driver in the VMkernel to the iSCSI initiator, hardware or software.
5. If the iSCSI initiator is a hardware iSCSI adapter, independent or dependent, the adapter performs the following tasks.
 - a. Encapsulates I/O requests into iSCSI Protocol Data Units (PDUs).
 - b. Encapsulates iSCSI PDUs into TCP/IP packets.
 - c. Sends IP packets over Ethernet to the iSCSI storage system.
6. If the iSCSI initiator is a software iSCSI adapter, the following takes place.
 - a. The iSCSI initiator encapsulates I/O requests into iSCSI PDUs.
 - b. The initiator sends iSCSI PDUs through TCP/IP connections.
 - c. The VMkernel TCP/IP stack relays TCP/IP packets to a physical NIC.
 - d. The physical NIC sends IP packets over Ethernet to the iSCSI storage system.
7. Ethernet switches and routers on the network carry the request to the appropriate storage device.

Error Correction

To protect the integrity of iSCSI headers and data, the iSCSI protocol offers error correction methods known as header digests and data digests.

Both methods are supported by the ESXi host, and you can activate them. These methods are used to check the header and SCSI data transferred between iSCSI initiators and targets, in both directions.

Header and data digests check the noncryptographic data integrity beyond the integrity checks that other networking layers, such as TCP and Ethernet, provide. The digests check the entire communication path, including all elements that can change the network-level traffic, such as routers, switches, and proxies.

The existence and type of the digests are negotiated when an iSCSI connection is established. When the initiator and target agree on a digest configuration, this digest must be used for all traffic between them.

Enabling header and data digests requires additional processing for both the initiator and the target, and can affect throughput and CPU use performance.

Note: Systems that use the Intel Nehalem processors offload the iSCSI digest calculations, as a result, reducing the impact on performance.

For information on enabling header and data digests, see [Configuring Advanced Parameters for iSCSI on ESXi Host](#).

Bootting from iSCSI SAN with ESXi

When you set up your ESXi host to boot from a SAN, your host's boot image is stored on one or more LUNs in the SAN storage system. When the ESXi host starts, it boots from the LUN on the SAN rather than from its local disk.

You can use boot from the SAN if you do not want to handle maintenance of local storage or have diskless hardware configurations, such as blade systems.

ESXi supports different methods of booting from the iSCSI SAN.

Table 210: Boot from iSCSI SAN support

Independent Hardware iSCSI	Software iSCSI
Configure the iSCSI HBA to boot from the SAN. For information on configuring the HBA, see Configure Independent Hardware iSCSI Adapter for SAN Boot .	Use the software iSCSI adapter and a network adapter that supports the iSCSI Boot Firmware Table (iBFT) format. For information, see <i>VMware ESXi Installation and Setup</i> .

General Recommendations for Boot from iSCSI SAN

If you plan to set up and use an iSCSI LUN as the boot device for your host, follow certain general guidelines.

The following guidelines apply to booting from the independent hardware iSCSI and iBFT.

- Review any vendor recommendations for the hardware you use in your boot configuration.
- For installation prerequisites and requirements, review *vSphere Installation and Setup*.
- Use static IP addresses to reduce the chances of DHCP conflicts.
- Use different LUNs for VMFS datastores and boot partitions.
- Configure proper ACLs on your storage system.
 - The boot LUN must be visible only to the host that uses the LUN. No other host on the SAN is permitted to see that boot LUN.
 - If a LUN is used for a VMFS datastore, multiple hosts can share the LUN.
- Configure a diagnostic partition.
 - With the independent hardware iSCSI only, you can place the diagnostic partition on the boot LUN. If you configure the diagnostic partition in the boot LUN, this LUN cannot be shared across multiple hosts. If a separate LUN is used for the diagnostic partition, multiple hosts can share the LUN.
 - If you boot from SAN using iBFT, you cannot set up a diagnostic partition on a SAN LUN. To collect your host's diagnostic information, use the vSphere ESXi Dump Collector on a remote server. For information about the ESXi Dump Collector, see *vCenter Server Installation and Setup* and *vSphere Networking*.

Prepare the iSCSI SAN

Before you configure your host to boot from an iSCSI LUN, prepare and configure your storage area network.

CAUTION: If you use scripted installation to install ESXi when booting from a SAN, you must take special steps to avoid unintended data loss.

1. Connect network cables, referring to any cabling guide that applies to your setup.
2. Ensure IP connectivity between your storage system and server.

Verify configuration of any routers or switches on your storage network. Storage systems must be able to ping the iSCSI adapters in your hosts.

3. Configure the storage system.
 - a) Create a volume (or LUN) on the storage system for your host to boot from.
 - b) Configure the storage system so that your host has access to the assigned LUN.

This step might involve updating ACLs with the IP addresses, iSCSI names, and the CHAP authentication parameter you use on your host. On some storage systems, in addition to providing access information for the ESXi host, you must also explicitly associate the assigned LUN with the host.
 - c) Ensure that the LUN is presented to the host correctly.
 - d) Ensure that no other system has access to the configured LUN.
 - e) Record the iSCSI name and IP addresses of the targets assigned to the host.

You must have this information to configure your iSCSI adapters.

Configure Independent Hardware iSCSI Adapter for SAN Boot

If your ESXi host uses an independent hardware iSCSI adapter, such as QLogic HBA, you can configure the adapter to boot from the SAN.

This procedure discusses how to enable the QLogic iSCSI HBA to boot from the SAN. For more information and more up-to-date details about QLogic adapter configuration settings, see the QLogic website.

1. Start the installation media and reboot the host.
2. Use the BIOS to set the host to boot from the installation media first.
3. During server POST, press Ctrl+q to enter the QLogic iSCSI HBA configuration menu.
4. Select the I/O port to configure.

By default, the Adapter Boot mode is set to Disable.
5. Configure the HBA.
 - a) From the **Fast!UTIL Options** menu, select **Configuration Settings > Host Adapter Settings**.
 - b) Optional: Configure the following settings for your host adapter: initiator IP address, subnet mask, gateway, initiator iSCSI name, and CHAP.
6. Configure iSCSI settings.

See [Configure iSCSI Boot Settings](#).

7. Save your changes and restart the system.

Configure iSCSI Boot Settings

Configure iSCSI boot parameters, so that your ESXi host can boot from an iSCSI LUN.

1. From the **Fast!UTIL Options** menu, select **Configuration Settings > iSCSI Boot Settings**.
2. Before you can set SendTargets, set Adapter Boot mode to **Manual**.
3. Select **Primary Boot Device Settings**.
 - a) Enter the discovery **Target IP** and **Target Port**.
 - b) Configure the **Boot LUN** and **iSCSI Name** parameters.
 - If only one iSCSI target and one LUN are available at the target address, leave **Boot LUN** and **iSCSI Name** blank.
After your host reaches the target storage system, these text boxes are populated with appropriate information.
 - If more than one iSCSI target and LUN are available, supply values for **Boot LUN** and **iSCSI Name**.
 - c) Save changes.
4. From the **iSCSI Boot Settings** menu, select the primary boot device.
An auto rescan of the HBA finds new target LUNs.
5. Select the iSCSI target.
If more than one LUN exists within the target, you can select a specific LUN ID by pressing **Enter** after you locate the iSCSI device.
6. Return to the **Primary Boot Device Setting** menu. After the rescan, **Boot LUN** and **iSCSI Name** are populated.
Change the value of **Boot LUN** to the appropriate LUN ID.

Best Practices for iSCSI Storage with ESXi

When you use ESXi with iSCSI SAN, follow recommendations that VMware offers to avoid problems.

Check with your storage representative if your storage system supports Storage API - Array Integration hardware acceleration features. If it does, refer to your vendor documentation to enable hardware acceleration support on the storage system side. For more information, see [Storage Hardware Acceleration in vSphere](#).

Preventing iSCSI SAN Problems

When using ESXi with a SAN, you must follow specific guidelines to avoid SAN problems.

Observe the following tips:

- Place only one VMFS datastore on each LUN.
- Do not change the path policy the system sets for you unless you understand the implications of making such a change.
- Document everything. Include information about configuration, access control, storage, switch, server and iSCSI HBA configuration, software and firmware versions, and storage cable plan.
- Plan for failure:
 - Make several copies of your topology maps. For each element, consider what happens to your SAN if the element fails.
 - Cross off different links, switches, HBAs, and other elements to ensure that you did not miss a critical failure point in your design.
- Ensure that the iSCSI HBAs are installed in the correct slots in the ESXi host, based on slot and bus speed. Balance PCI bus load among the available buses in the server.

- Become familiar with the various monitor points in your storage network, at all visibility points, including ESXi performance charts, Ethernet switch statistics, and storage performance statistics.
- Change LUN IDs only when VMFS datastores deployed on the LUNs have no running virtual machines. If you change the ID, virtual machines running on the VMFS datastore might fail.
After you change the ID of the LUN, you must rescan your storage to reset the ID on your host. For information on using the rescan, see [Rescan Operations for ESXi Storage](#).
- If you change the default iSCSI name of your iSCSI adapter, make sure that the name you enter is worldwide unique and properly formatted. To avoid storage access problems, never assign the same iSCSI name to different adapters, even on different hosts.

Optimizing iSCSI SAN Storage Performance

Several factors contribute to optimizing a typical SAN environment.

If the network environment is properly configured, the iSCSI components provide adequate throughput and low enough latency for iSCSI initiators and targets. If the network is congested and links, switches or routers are saturated, iSCSI performance suffers and might not be adequate for ESXi environments.

Storage System Performance

Storage system performance is one of the major factors contributing to the performance of the entire iSCSI environment.

If issues occur with storage system performance, consult your storage system vendor's documentation for any relevant information.

When you assign LUNs, remember that you can access each shared LUN through a number of hosts, and that a number of virtual machines can run on each host. One LUN used by the ESXi host can service I/O from many different applications running on different operating systems. Because of this diverse workload, the RAID group that contains the ESXi LUNs should not include LUNs that other hosts use that are not running ESXi for I/O intensive applications.

Enable read caching and write caching.

Load balancing is the process of spreading server I/O requests across all available SPs and their associated host server paths. The goal is to optimize performance in terms of throughput (I/O per second, megabytes per second, or response times).

SAN storage systems require continual redesign and tuning to ensure that I/O is load balanced across all storage system paths. To meet this requirement, distribute the paths to the LUNs among all the SPs to provide optimal load balancing. Close monitoring indicates when it is necessary to manually rebalance the LUN distribution.

Tuning statically balanced storage systems is a matter of monitoring the specific performance statistics (such as I/O operations per second, blocks per second, and response time) and distributing the LUN workload to spread the workload across all the SPs.

Server Performance with iSCSI

To ensure optimal ESXi host performance, consider several factors.

Each server application must have access to its designated storage with the following conditions:

- High I/O rate (number of I/O operations per second)
- High throughput (megabytes per second)
- Minimal latency (response times)

Because each application has different requirements, you can meet these goals by selecting an appropriate RAID group on the storage system.

To achieve performance goals, follow these guidelines:


- Place each LUN on a RAID group that provides the necessary performance levels. Monitor the activities and resource use of other LUNS in the assigned RAID group. A high-performance RAID group that has too many applications doing I/O to it might not meet performance goals required by an application running on the ESXi host.
- To achieve maximum throughput for all the applications on the host during the peak period, install enough network adapters or iSCSI hardware adapters. I/O spread across multiple ports provides faster throughput and less latency for each application.
- To provide redundancy for software iSCSI, make sure that the initiator is connected to all network adapters used for iSCSI connectivity.
- When allocating LUNs or RAID groups for ESXi systems, remember that multiple operating systems use and share that resource. The LUN performance required by the ESXi host might be much higher than when you use regular physical machines. For example, if you expect to run four I/O intensive applications, allocate four times the performance capacity for the ESXi LUNs.
- When you use multiple ESXi systems with vCenter Server, the storage performance requirements increase.
- The number of outstanding I/Os needed by applications running on an ESXi system must match the number of I/Os the SAN can handle.

Network Performance

A typical SAN consists of a collection of computers connected to a collection of storage systems through a network of switches. Several computers often access the same storage.


The following graphic shows several computer systems connected to a storage system through an Ethernet switch. In this configuration, each system is connected through a single Ethernet link to the switch. The switch is connected to the storage system through a single Ethernet link.

Figure 96: Single Ethernet Link Connection to Storage



When systems read data from storage, the storage responds with sending enough data to fill the link between the storage systems and the Ethernet switch. It is unlikely that any single system or virtual machine gets full use of the network speed. However, this situation can be expected when many systems share one storage device.

When writing data to storage, multiple systems or virtual machines might attempt to fill their links. As a result, the switch between the systems and the storage system might drop network packets. The data drop might occur because the switch has more traffic to send to the storage system than a single link can carry. The amount of data the switch can transmit is limited by the speed of the link between it and the storage system.

Figure 97: Dropped Packets

Recovering from dropped network packets results in large performance degradation. In addition to time spent determining that data was dropped, the retransmission uses network bandwidth that can otherwise be used for current transactions.

iSCSI traffic is carried on the network by the Transmission Control Protocol (TCP). TCP is a reliable transmission protocol that ensures that dropped packets are retried and eventually reach their destination. TCP is designed to recover from dropped packets and retransmits them quickly and seamlessly. However, when the switch discards packets with any regularity, network throughput suffers. The network becomes congested with requests to resend data and with the resent packets. Less data is transferred than in a network without congestion.

Most Ethernet switches can buffer, or store, data. This technique gives every device attempting to send data an equal chance to get to the destination. The ability to buffer some transmissions, combined with many systems limiting the number of outstanding commands, reduces transmissions to small bursts. The bursts from several systems can be sent to a storage system in turn.

If the transactions are large and multiple servers are sending data through a single switch port, an ability to buffer can be exceeded. In this case, the switch drops the data it cannot send, and the storage system must request a retransmission of the dropped packet. For example, if an Ethernet switch can buffer 32 KB, but the server sends 256 KB to the storage device, some of the data is dropped.

Most managed switches provide information on dropped packets, similar to the following:

```
*: interface is up
IHQ: pkts in input hold queue      IQD: pkts dropped from input queue
OHQ: pkts in output hold queue    OQD: pkts dropped from output queue
RXBS: rx rate (bits/sec)          RXPS: rx rate (pkts/sec)
TXBS: tx rate (bits/sec)          TXPS: tx rate (pkts/sec)
TRTL: throttle count
```

Table 211: Sample Switch Information

Interface	IHQ	IQD	OHQ	OQD	RXBS	RXPS	TXBS	TXPS	TRTL
* GigabitEthernet0/1	3	9922	0	0	476303000	62273	477840000	63677	0


In this example from a Cisco switch, the bandwidth used is 476303000 bits/second, which is less than half of wire speed. The port is buffering incoming packets, but has dropped several packets. The final line of this interface summary indicates that this port has already dropped almost 10,000 inbound packets in the IQD column.

Configuration changes to avoid this problem involve making sure several input Ethernet links are not funneled into one output link, resulting in an oversubscribed link. When several links transmitting near capacity are switched to a smaller number of links, oversubscription becomes possible.

Generally, applications or systems that write much data to storage must avoid sharing Ethernet links to a storage device. These types of applications perform best with multiple connections to storage devices.

Multiple Connections from Switch to Storage shows multiple connections from the switch to the storage.

Figure 98: Multiple Connections from Switch to Storage



Using VLANs or VPNs does not provide a suitable solution to the problem of link oversubscription in shared configurations. VLANs and other virtual partitioning of a network provide a way of logically designing a network. However, they do not change the physical capabilities of links and trunks between switches. When storage traffic and other network traffic share physical connections, oversubscription and lost packets might become possible. The same is true of VLANs that share interswitch trunks. Performance design for a SAN must consider the physical limitations of the network, not logical allocations.

Checking Ethernet Switch Statistics

Many Ethernet switches provide different methods for monitoring switch health.

Switches that have ports operating near maximum throughput much of the time do not provide optimum performance. If you have ports in your iSCSI SAN running near the maximum, reduce the load. If the port is connected to an ESXi system or iSCSI storage, you can reduce the load by using manual load balancing.

If the port is connected between multiple switches or routers, consider installing additional links between these components to handle more load. Ethernet switches also commonly provide information about transmission errors, queued packets, and dropped Ethernet packets. If the switch regularly reports any of these conditions on ports being used for iSCSI traffic, performance of the iSCSI SAN will be poor.

Configuring iSCSI and iSER Adapters and Storage with ESXi

Before ESXi can work with iSCSI SAN, you must set up your iSCSI environment. To work properly, your environment must follow specific recommendations. In addition, several restrictions exist when you use ESXi with iSCSI SAN.

iSCSI Storage Recommendations

- Verify that your ESXi host supports the iSCSI SAN storage hardware and firmware. For an up-to-date list, see *VMware Compatibility Guide*.
- To ensure that the host recognizes LUNs at startup time, configure all iSCSI storage targets so that your host can access them and use them. Configure your host so that it can discover all available iSCSI targets.
- Unless you are using diskless servers, set up a diagnostic partition on local storage. If you have diskless servers that boot from iSCSI SAN, see [General Recommendations for Boot from iSCSI SAN](#) for information about diagnostic partitions with iSCSI.
- Set the SCSI controller driver in the guest operating system to a large enough queue.
- VMtools increases the default disk timeout SCSI `TimeoutValue` parameter to 180 seconds. On virtual machines without VMtools, increase the value of the SCSI `TimeoutValue` parameter. With this parameter set up, the VMs can better tolerate delayed I/O resulting from a path failover. For information, see [Path Failover and Virtual Machines](#) or the [KB 1009465](#).
- Configure your environment to have only one VMFS datastore for each LUN.

iSCSI Storage Restrictions

- ESXi does not support iSCSI-connected tape devices.
- You cannot use virtual-machine multipathing software to perform I/O load balancing to a single physical LUN.
- ESXi does not support multipathing when you combine independent hardware adapters with either software or dependent hardware adapters.

Setting Up iSCSI ESXi Environment

The process of preparing your iSCSI environment involves the following steps.

Step	Details
1. Set up iSCSI storage	For information, see your storage vendor documentation. In addition, follow recommendations in Best Practices for iSCSI Storage with ESXi .
2. Configure iSCSI/iSER adapters	Use an appropriate workflow to configure your adapter: • Set Up Independent Hardware iSCSI Adapters with ESXi • Configure Dependent Hardware iSCSI Adapters with ESXi • Configure the Software iSCSI Adapter with ESXi • Configure iSER with ESXi
3. Create a datastore on iSCSI storage	Creating vSphere Datastores

iSCSI Networking

For certain types of iSCSI adapters, you must configure VMkernel networking.

You can verify the network configuration by using the `vmkping` utility.

The independent hardware iSCSI adapter does not require VMkernel networking. You can configure network parameters, such as an IP address, subnet mask, and default gateway on the independent hardware iSCSI adapter.

All types of iSCSI adapters support IPv4 and IPv6 protocols.

iSCSI Adapter (vmhba)	Description	VMkernel Networking	Adapter Network Settings
Independent Hardware iSCSI Adapter	Third-party adapter that offloads the iSCSI and network processing and management from your host.	Not required.	For information, see Edit Network Settings for Hardware iSCSI .
Software iSCSI Adapter	Uses standard NICs to connect your host to a remote iSCSI target on the IP network.	Required. For information, see Setting Up Network for iSCSI and iSER with ESXi .	N/A
Dependent Hardware iSCSI Adapter	Third-party adapter that depends on VMware networking and iSCSI configuration and management interfaces.	Required For information, see Setting Up Network for iSCSI and iSER with ESXi .	N/A
VMware iSER Adapter	Uses an RDMA-capable network adapter to connect your host to a remote iSCSI target.	Required For information, see Setting Up Network for iSCSI and iSER with ESXi .	N/A

Discovery Methods

For all types of iSCSI adapters, you must set the dynamic discovery address or static discovery address. In addition, you must provide a target name of the storage system. For software iSCSI and dependent hardware iSCSI, the address is pingable using `vmkping`.

See [Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host](#).

CHAP Authentication

Enable the CHAP parameter on the initiator and the storage system side. After authentication is enabled, it applies to all targets that are not yet discovered. It does not apply to targets that are already discovered.

See [Configuring CHAP Parameters for iSCSI or iSER Storage Adapters on ESXi Host](#).

Set Up Independent Hardware iSCSI Adapters with ESXi

An independent hardware iSCSI adapter is a specialized third-party adapter capable of accessing iSCSI storage over TCP/IP. This iSCSI adapter handles all iSCSI and network processing and management for your ESXi system.

- Verify whether the adapter must be licensed.
- Install the adapter on your ESXi host.

For information about licensing, installation, and firmware updates, see vendor documentation.

The process of setting up the independent hardware iSCSI adapter includes these steps.

Step	Description
View Independent Hardware iSCSI Adapters	View an independent hardware iSCSI adapter and verify that it is correctly installed and ready for configuration.
Modify General Properties for iSCSI or iSER Adapters on ESXi Hosts	If needed, change the default iSCSI name and alias assigned to your iSCSI adapters. For the independent hardware iSCSI adapters, you can also change the default IP settings.
Edit Network Settings for Hardware iSCSI	Change default network settings so that the adapter is configured properly for the iSCSI SAN.
Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host	Set up dynamic discovery. With dynamic discovery, each time the initiator contacts a specified iSCSI storage system, it sends the SendTargets request to the system. The iSCSI system responds by supplying a list of available targets to the initiator. In addition to the dynamic discovery method, you can use static discovery and manually enter information for the targets.
Set Up CHAP for iSCSI or iSER Storage Adapter	If your iSCSI environment uses the Challenge Handshake Authentication Protocol (CHAP), configure it for your adapter.
Enable Jumbo Frames for Independent Hardware iSCSI	If your iSCSI environment supports Jumbo Frames, enable them for the adapter.

View Independent Hardware iSCSI Adapters

On the ESXi host, view an independent hardware iSCSI adapter and verify that it is correctly installed and ready for configuration.

Required privilege: **Host > Configuration > Storage Partition Configuration**

After you install an independent hardware iSCSI adapter on the host, it appears on the list of storage adapters available for configuration. You can view its properties.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**.
If installed, the hardware iSCSI adapter appears on the list of storage adapters.
4. Select the adapter to view.
The default details for the adapter appear.

Adapter Information	Description
Model	Model of the adapter.
iSCSI Name	Unique name formed according to iSCSI standards that identifies the iSCSI adapter. You can edit the iSCSI name.
iSCSI Alias	A friendly name used instead of the iSCSI name. You can edit the iSCSI alias.
IP Address	Address assigned to the iSCSI HBA.
Targets	Number of targets accessed through the adapter.
Devices	All storage devices or LUNs the adapter can access.
Paths	All paths the adapter uses to access storage devices.

Edit Network Settings for Hardware iSCSI

After you install an independent hardware iSCSI adapter on an ESXi host, you might need to change its default network settings so that the adapter is configured properly for the iSCSI SAN.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
4. Click the **Network Settings** tab and click **Edit**.
5. In the IPv4 settings section, deactivate IPv4 or select the method to obtain IP addresses.

Note: The automatic DHCP option and static option are mutually exclusive.

Option	Description
No IPv4 settings	Deactivate IPv4.
Obtain IPv4 settings automatically	Use DHCP to obtain IP settings.
Use static IPv4 settings	Enter the IPv4 IP address, subnet mask, and default gateway for the iSCSI adapter.

6. In the IPv6 settings section, deactivate IPv6 or select an appropriate option for obtaining IPv6 addresses.

Note: Automatic options and the static option are mutually exclusive.

Option	Description
No IPv6 settings	Deactivate IPv6.
Enable IPv6	Select an option for obtaining IPv6 addresses.

Option	Description
Obtain IPv6 addresses automatically through DHCP	Use DHCP to obtain IPv6 addresses.
Obtain IPv6 addresses automatically through Router Advertisement	Use router advertisement to obtain IPv6 addresses.
Override Link-local address for IPv6	Override the link-local IP address by configuring a static IP address.
Static IPv6 addresses	1. Click Add to add a new IPv6 address. 2. Enter the IPv6 address and subnet prefix length, and click OK.

7. In the DNS settings section, provide IP addresses for a preferred DNS server and an alternate DNS server.
You must provide both values.

Configure Dependent Hardware iSCSI Adapters with ESXi

A dependent hardware iSCSI adapter in your ESXi system is a third-party adapter that depends on VMware networking, and iSCSI configuration and management interfaces provided by VMware.

An example of a dependent iSCSI adapter is a Broadcom 5709 NIC. When installed on a host, it presents its two components, a standard network adapter and an iSCSI engine, to the same port. The iSCSI engine appears on the list of storage adapters as an iSCSI adapter (vmhba).

The iSCSI adapter is enabled by default. To make it functional, you must connect it, through a virtual VMkernel adapter (vmk), to a physical network adapter (vmnic) associated with it. You can then configure the iSCSI adapter.

After you configure the dependent hardware iSCSI adapter, the discovery and authentication data is passed through the network connection. The iSCSI traffic goes through the iSCSI engine, bypassing the network.

The entire setup and configuration process for the dependent hardware iSCSI adapters involves several steps.

Step	Description
View Dependent Hardware iSCSI Adapters	View a dependent hardware iSCSI adapter to verify that it is correctly loaded.
Modify General Properties for iSCSI or iSER Adapters on ESXi Hosts	If needed, change the default iSCSI name and alias assigned to your adapter.
Determine Association Between iSCSI and Network Adapters	You must create network connections to bind dependent iSCSI and physical network adapters. To create the connections correctly, determine the name of the physical NIC with which the dependent hardware iSCSI adapter is associated.
Configure Port Binding for iSCSI or iSER on ESXi	Configure connections for the traffic between the iSCSI component and the physical network adapters. The process of configuring these connections is called port binding.
Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host	Set up dynamic discovery. With dynamic discovery, each time the initiator contacts a specified iSCSI storage system, it sends the SendTargets request to the system. The iSCSI system responds by supplying a list of available targets to the initiator. In addition to the dynamic discovery method, you can use static discovery and manually enter information for the targets.
Set Up CHAP for iSCSI or iSER Storage Adapter	If your iSCSI environment uses the Challenge Handshake Authentication Protocol (CHAP), configure it for your adapter.
Set Up CHAP for Target	You can also configure different CHAP credentials for each discovery address or static target.

Step	Description
Enable Jumbo Frames for Networking	If your iSCSI environment supports Jumbo Frames, enable them for the adapter.

Dependent Hardware iSCSI Considerations

When you use dependent hardware iSCSI adapters with ESXi, certain considerations apply.

- When you use any dependent hardware iSCSI adapter, performance reporting for a NIC associated with the adapter might show little or no activity, even when iSCSI traffic is heavy. This behavior occurs because the iSCSI traffic bypasses the regular networking stack.
- If you use a third-party virtual switch deactivate automatic pinning. Use manual pinning instead, making sure to connect a VMkernel adapter (vmk) to an appropriate physical NIC (vmnic). For information, refer to your virtual switch vendor documentation.
- The Broadcom iSCSI adapter performs data reassembly in hardware, which has a limited buffer space. When you use the Broadcom iSCSI adapter in a congested network or under heavy load, activate flow control to avoid performance degradation.

Flow control manages the rate of data transmission between two nodes to prevent a fast sender from overrunning a slow receiver. For best results, activate flow control at the end points of the I/O path, at the hosts and iSCSI storage systems.

To activate flow control for the host, use the `esxcli system module parameters` command.

- Dependent hardware adapters support IPv4 and IPv6.

View Dependent Hardware iSCSI Adapters

On an ESXi host, view a dependent hardware iSCSI adapter to verify that it is correctly loaded.

If installed, the dependent hardware iSCSI adapter (vmhba#) appears on the list of storage adapters under such category as, for example, Broadcom iSCSI Adapter. If the dependent hardware adapter does not appear on the list of storage adapters, check whether it needs to be licensed. See your vendor documentation.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**.
4. Select the adapter (vmhba#) to view.

The default details for the adapter appear, including the iSCSI name, iSCSI alias, and the status.

Although the dependent iSCSI adapter is enabled by default, to make it functional, you must set up networking for the iSCSI traffic and bind the adapter to the appropriate VMkernel iSCSI port. You then configure discovery addresses and CHAP parameters.

Determine Association Between iSCSI and Network Adapters

On an ESXi host, network connections bind dependent iSCSI and physical network adapters. To create the connections correctly, you must determine the name of the physical NIC with which the dependent hardware iSCSI adapter is associated.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**.
4. Select the iSCSI adapter (vmhba#) and click the **Network Port Binding** tab under adapter details.
5. Click **Add**.

The network adapter (vmnic#) that corresponds to the dependent iSCSI adapter is listed in the Physical Network Adapter column.

If the VMkernel Adapter column is empty, create a VMkernel adapter (vmk#) for the physical network adapter (vmnic#) and then bind them to the associated dependent hardware iSCSI. See [Setting Up Network for iSCSI and iSER with ESXi](#).

Configure the Software iSCSI Adapter with ESXi

With the software-based iSCSI implementation, you can use standard NICs to connect your host to a remote iSCSI target on the IP network. The software iSCSI adapter that is built into ESXi facilitates this connection by communicating with the physical NICs through the network stack.

When you use the software iSCSI adapters, consider the following:

- Designate a separate network adapter for iSCSI. Do not use iSCSI on 100 Mbps or slower adapters.
- Avoid hard coding the name of the software adapter, vmhbaXX, in the scripts. It is possible for the name to change from one ESXi release to another. The change might cause failures of your existing scripts if they use the hardcoded old name. The name change does not affect the behavior of the iSCSI software adapter.

The process of configuring the software iSCSI adapter involves several steps.

Step	Description
Add or Remove the Software iSCSI Adapter	Activate your software iSCSI adapter so that your host can use it to access iSCSI storage.
Modify General Properties for iSCSI or iSER Adapters on ESXi Hosts	If needed, change the default iSCSI name and alias assigned to your adapter.
Configure Port Binding for iSCSI or iSER on ESXi	Configure connections for the traffic between the iSCSI component and the physical network adapters. The process of configuring these connections is called port binding.
Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host	Set up dynamic discovery. With dynamic discovery, each time the initiator contacts a specified iSCSI storage system, it sends the SendTargets request to the system. The iSCSI system responds by supplying a list of available targets to the initiator. In addition to the dynamic discovery method, you can use static discovery and manually enter information for the targets.
Set Up CHAP for iSCSI or iSER Storage Adapter	If your iSCSI environment uses the Challenge Handshake Authentication Protocol (CHAP), configure it for your adapter.
Set Up CHAP for Target	You can also configure different CHAP credentials for each discovery address or static target.

Step	Description
Enable Jumbo Frames for Networking	If your iSCSI environment supports Jumbo Frames, enable them for the adapter.

Add or Remove the Software iSCSI Adapter

You must activate your software iSCSI adapter so that your ESXi host can use it to access iSCSI storage. If you do not need the software iSCSI adapter after activation, you can remove it.

Required privilege: **Host > Configuration > Storage Partition Configuration**

If you plan to remove the software iSCSI adapter, perform these steps:

- Migrate all VMs from the iSCSI datastore.
- Unmount the datastore. See [Unmount Datastores](#).
- Detach the iSCSI device and perform the adapter rescan. See [Performing Planned Storage Device Removal](#).

Note: If you boot from iSCSI using the software iSCSI adapter, the adapter is added and the network configuration is created at the first boot. If you remove the adapter, it is added again each time you boot the host.

You can activate only one software iSCSI adapter.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Add or remove the adapter.

Option	Description
Add the software iSCSI adapter	1. Under Storage, click Storage Adapters, and click Add Software Adapter. 2. From the drop-down menu, select Software iSCSI Adapter and confirm that you want to add the adapter. The software iSCSI adapter (vmhba#) is added and appears on the list of storage adapters. After adding the adapter, the host assigns the default iSCSI name to it. You can now complete the adapter configuration.
Remove the software iSCSI adapter	1. Under Storage, click Storage Adapters, and select the adapter (vmhba#) to remove. 2. Click the Properties tab. 3. From the Adapter Status Actions menu, select Disable and confirm that you want to deactivate the adapter. The adapter will be marked for removal. The adapter will be removed from the host on next host reboot. 4. Reboot the host. After the reboot, the adapter no longer appears on the list of storage adapters. All virtual machines and other data on the storage devices associated with this adapter become inaccessible. You can later activate the adapter.

Configure iSER with ESXi

In addition to traditional iSCSI, ESXi supports the iSCSI Extensions for RDMA (iSER) protocol. When the iSER protocol is enabled, the iSCSI framework on the ESXi host can use the Remote Direct Memory Access (RDMA) transport instead of TCP/IP. You can configure iSER on your ESXi host.

For more information about the iSER protocol, see [Using iSER Protocol with ESXi](#).

The entire setup and configuration process for VMware iSER involves several steps.

Step	Description
Install and View an RDMA Capable Network Adapter	To configure iSER with ESXi, you must first install an RDMA capable network adapter, for example, Mellanox Technologies MT27700 Family ConnectX-4. After you install this type of adapter, the vSphere Client displays its two components, an RDMA adapter and a physical network adapter <code>vmnic#</code> .
Enable the VMware iSER Adapter	To be able to use the RDMA capable adapter for iSCSI, use the <code>esxcli</code> to enable the VMware iSER storage component. The component appears in the vSphere Client as a <code>vmhba#</code> storage adapter under the VMware iSCSI over RDMA (iSER) Adapter category.
Modify General Properties for iSCSI or iSER Adapters on ESXi Hosts	If needed, change the default name and alias assigned to the iSER storage adapter <code>vmhba#</code> .
Configure Port Binding for iSCSI or iSER on ESXi	You must create network connections to bind the iSER storage adapter <code>vmhba#</code> and the RDMA capable network adapter <code>vmnic#</code> . The process of configuring these connections is called port binding. Note: iSER does not support NIC teaming. When configuring port binding, use only one RDMA adapter per vSwitch.
Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host	Set up the dynamic or static discovery for your iSER storage adapter <code>vmhba#</code> . With the dynamic discovery, each time the initiator contacts a specified iSER storage system, it sends the SendTargets request to the system. The iSER system responds by supplying a list of available targets to the initiator. With the static discovery, you manually enter information for the targets.
Set Up CHAP for iSCSI or iSER Storage Adapter	If your environment uses the Challenge Handshake Authentication Protocol (CHAP), configure it for your iSER storage adapter <code>vmhba#</code> .
Set Up CHAP for Target	You can also configure different CHAP credentials for each discovery address or static target.
Enable Jumbo Frames for Networking	If your environment supports Jumbo Frames, enable them for the iSER storage adapter <code>vmhba#</code> .


Install and View an RDMA Capable Network Adapter

ESXi supports RDMA capable network adapters, for example, Mellanox Technologies MT27700 Family ConnectX-4. After you install such adapter on your host, the vSphere Client displays its two components, an RDMA adapter and a physical network adapter.

You can use the vSphere Client to view the RDMA adapter and its corresponding network adapter.

1. In the vSphere Client, navigate to the ESXi host.
2. Under **Networking**, click **RDMA adapters**.

In this example, the RDMA adapter appears on the list as `vmrdma0`. The **Paired Uplink** column displays the network component as the `vmnic1` physical network adapter.



3. To verify the description of the adapter, select the RDMA adapter from the list, and click the **Properties** tab.

You can use the `vmnic#` network component of the adapter for such storage configurations as iSER or NVMe over RDMA. For the iSER configuration steps, see [Configure iSER with ESXi](#). For information about NVMe over RDMA, see [Configuring NVMe over RDMA \(RoCE v2\) on ESXi](#).

Enable the VMware iSER Adapter

To be able to use the RDMA capable adapter for iSCSI, use the `esxcli` to enable the VMware iSER storage component. After you enable the component, it appears in the vSphere Client as a `vmhba#` storage adapter under the VMware iSCSI over RDMA (iSER) Adapter category.

- Make sure that your iSCSI storage supports the iSER protocol.
 - Install the RDMA capable adapter on your ESXi host. For information, see [Install and View an RDMA Capable Network Adapter](#).
 - For RDMA capable adapters that support RDMA over Converged Ethernet (RoCE), determine the RoCE version that the adapter uses.
 - Use the RDMA capable switch.
 - Enable flow control on the ESXi host. To enable flow control for the host, use the `esxcli system module parameters` command.
 - Make sure to configure RDMA switch ports to create lossless connections between the iSER initiator and target.
1. Use the ESXi Shell or vSphere CLI to enable the VMware iSER storage adapter and set its RoCE version.
 - a) Enable the iSER storage adapter.

```
esxcli rdma iser add
```

- b) Verify that the iSER adapter has been added.

```
esxcli iscsi adapter list
```

The output is similar to the following.

```
Adapter Driver State UID Description
-----
vmhba64 iser unbound iscsi.vmhba64 VMware iSCSI over RDMA (iSER) Adapter
```

- c) Specify the RoCE version that iSER uses to connect to the target.

Use the RoCE version of the RDMA capable adapter. The command you enter is similar to the following:

```
esxcli rdma iser params set -a vmhba64 -r 1
```

When the command completes, a message similar to the following appears in the VMkernel log.

```
vmkernel.0:2020-02-18T18:26:15.949Z cpu6:2100717 opID=45abe37e)iser: iser_set_roce: Setting roce type:
1 for vmhba: vmhba64
vmkernel.0:2020-02-18T18:26:15.949Z cpu6:2100717 opID=45abe37e)iser: iser_set_roce: Setting rdma port:
3260 for vmhba: vmhba64
```

If you do not specify the RoCE version, the host defaults to the highest RoCE version the RDMA capable adapter supports.

2. Use the vSphere Client to display the iSER adapter.

- In the vSphere Client, navigate to the ESXi host.
- Click the **Configure** tab.
- Under **Storage**, click **Storage Adapters**, and review the list of adapters.

If you enabled the adapter, it appears as a storage vmhba# adapter on the list under the VMware iSCSI over RDMA (iSER) Adapter category.

Storage Adapters

+ Add Software Adapter Refresh Rescan Storage... Rescan Adapter Remove

Adapter	Type	Status	Identifier	Targets	Devices	Paths
vmhba33	Block SCSI	Unknown	--	1	2	2
Model: VMware iSCSI over RDMA (iSER) Adapter						
vmhba64	iSCSI	Unbound	iser-vmnic9	0	0	0
vmhba65	iSCSI	Unbound	iser-vmnic10	0	0	0
vmhba66	iSCSI	Unbound	iser-vmnic4	0	0	0
vmhba67	iSCSI	Unbound	iser-vmnic5	0	0	0
Model: Wellsburg AHCI Controller						
vmhba1	Block SCSI	Unknown	--	0	0	0
vmhba2	Block SCSI	Unknown	--	1	1	1

Copy All 9 items

Properties Devices Paths Dynamic Discovery Static Discovery Network Port Binding Advanced Options

General

Name: vmhba64 Edit...

Model: VMware iSCSI over RDMA (iSER) Adapter

iSCSI Name: ign.1998-01.com.vmware.prime-fcoe-005.ang.vmx Edit...

iSCSI Alias: iser-vmnic9

Target Discovery: Send Targets, Static Targets

Authentication

Method: None Edit...

3. Select the iSER storage vmhba# to review its properties or perform the following tasks.

Option	Description
Configure port binding for the iSER storage adapter	You must create network connections to bind the iSER storage adapter vmhba# and the RDMA capable network adapter vmnic#. The process of configuring these connections is called port binding.

Option	Description
	For general information about port binding, see Setting Up Network for iSCSI and iSER with ESXi . To configure port binding for iSER, see Configure Port Binding for iSCSI or iSER on ESXi .
Set up dynamic or static discovery for the iSER storage adapter	See Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host .
Configure the Challenge Handshake Authentication Protocol (CHAP) for the iSER storage adapter	See Set Up CHAP for iSCSI or iSER Storage Adapter .

Modify General Properties for iSCSI or iSER Adapters on ESXi Hosts

You can change default name and alias assigned to your iSCSI or iSER storage adapters by the ESXi host. For the independent hardware iSCSI adapters, you can also change the default IP settings.

Required privilege: **Host > Configuration > Storage Partition Configuration**

Important: When you modify any default properties for your adapters, make sure to use correct formats for their names and IP addresses.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
4. Click the **Properties** tab, and click **Edit** in the General panel.
5. Optional: Modify the following general properties.

Option	Description
iSCSI Name	Unique name formed according to iSCSI standards that identifies the iSCSI adapter. If you change the name, make sure that the name you enter is worldwide unique and properly formatted. Otherwise, certain storage devices might not recognize the iSCSI adapter.
iSCSI Alias	A friendly name you use instead of the iSCSI name.

If you change the iSCSI name, it is used for new iSCSI sessions. For existing sessions, the new settings are not used until you log out and log in again.


For other configuration steps you can perform for the iSCSI or iSER storage adapters, see the following topics:

- [Set Up Independent Hardware iSCSI Adapters with ESXi](#)
- [Configure Dependent Hardware iSCSI Adapters with ESXi](#)
- [Configure the Software iSCSI Adapter with ESXi](#)
- [Configure iSER with ESXi](#)

Setting Up Network for iSCSI and iSER with ESXi

Certain types of iSCSI adapters depend on the VMkernel networking. These adapters include the software or dependent hardware iSCSI adapters, and the VMware iSCSI over RDMA (iSER) adapter. If your ESXi environment includes any of these adapters, you must configure connections for the traffic between the iSCSI or iSER component and the physical network adapters.

Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You use 1:1 mapping between each virtual and physical network adapter. You then associate the VMkernel adapter with an appropriate iSCSI or iSER adapter. This process is called port binding.



Follow these rules when configuring the port binding:

- You can connect the software iSCSI adapter with any physical NICs available on your host.
- The dependent iSCSI adapters must be connected only to their own physical NICs.
- You must connect the iSER adapter only to the RDMA-capable network adapter.


For specific considerations on when and how to use network connections with software iSCSI, see the VMware knowledge base article at <http://kb.vmware.com/kb/2038869>.

Multiple Network Adapters in iSCSI or iSER Configuration


If your host has more than one physical network adapter for iSCSI or iSER, you can use the adapters for multipathing.

You can use multiple physical adapters in a single or multiple switch configurations.

In the multiple switch configuration, you designate a separate vSphere switch for each virtual-to-physical adapter pair.

Figure 99: 1:1 Adapter Mapping on Separate vSphere Standard Switches

An alternative is to add all NICs and VMkernel adapters to the single vSphere switch. The number of VMkernel adapters must correspond to the number of physical adapters on the vSphere Standard switch. The single switch configuration is not appropriate for iSER because iSER does not support NIC teaming.

Figure 100: 1:1 Adapter Mapping on a Single vSphere Standard Switch

For that type of configuration, you must override the default network setup and make sure that each VMkernel adapter maps to only one corresponding active physical adapter, as the table indicates.

VMkernel Adapter (vmk#)	Physical Network Adapter (vmnic#)
vmk1 (iSCSI1)	Active Adapters vmnic1 Unused Adapters vmnic2
vmk2 (iSCSI2)	Active Adapters vmnic2 Unused Adapters vmnic1

You can also use distributed switches. For more information about vSphere distributed switches and how to change the default network policy, see the *vSphere Networking* documentation.

The following considerations apply when you use multiple physical adapters:

- Physical network adapters must be on the same subnet as the storage system they connect to.
- (Applies only to iSCSI and not to iSER) If you use separate vSphere switches, you must connect them to different IP subnets. Otherwise, VMkernel adapters might experience connectivity problems and the host fails to discover the LUNs.
- The single switch configuration is not appropriate for iSER because iSER does not support NIC teaming.

Do not use port binding when any of the following conditions exist:

- Array target iSCSI ports are in a different broadcast domain and IP subnet.
- VMkernel adapters used for iSCSI connectivity exist in different broadcast domains, IP subnets, or use different virtual switches.

Note: In iSER configurations, the VMkernel adapters used for iSER connectivity cannot be used for converged traffic. The VMkernel adapters that you created to enable connectivity between the ESXi host with iSER and the iSER target must be used only for iSER traffic.

Best Practices for Configuring Networking with Software iSCSI

When you configure networking with software iSCSI, consider several best practices.

Software iSCSI Port Binding

You can bind the software iSCSI initiator on the ESXi host to a single or multiple VMkernel ports, so that iSCSI traffic flows only through the bound ports. Unbound ports are not used for iSCSI traffic.

When port binding is configured, the iSCSI initiator creates iSCSI sessions from all bound ports to all configured target portals.

See the following examples.

VMkernel Ports	Target Portals	iSCSI Sessions
2 bound VMkernel ports	2 target portals	4 sessions (2 x 2)
4 bound VMkernel ports	1 target portal	4 sessions (4 x 1)
2 bound VMkernel ports	4 target portals	8 sessions (2 x 4)

Note: Make sure that all target portals are reachable from all VMkernel ports when port binding is used. Otherwise, iSCSI sessions might fail to create. As a result, the rescan operation might take longer than expected.

No Port Binding

If you do not use port binding, the ESXi networking layer selects the best VMkernel port based on its routing table. The host uses the port to create an iSCSI session with the target portal. Without the port binding, only one session per each target portal is created.


See the following examples.

VMkernel Ports	Target Portals	iSCSI Sessions
2 unbound VMkernel ports	2 target portals	2 sessions
4 unbound VMkernel ports	1 target portal	1 session
2 unbound VMkernel ports	4 target portals	4 sessions

Software iSCSI Multipathing

Example 1. Multiple paths for an iSCSI target with a single network portal

If your target has only one network portal, you can create multiple paths to the target by adding multiple VMkernel ports on your ESXi host and binding them to the iSCSI initiator.




In this example, all initiator ports and the target portal are configured in the same subnet. The target is reachable through all bound ports. You have four VMkernel ports and one target portal, so total of four paths are created.

Without the port binding, only one path is created.

Example 2. Multiple paths with VMkernel ports in different subnets

You can create multiple paths by configuring multiple ports and target portals on different IP subnets. By keeping initiator and target ports in different subnets, you can force ESXi to create paths through specific ports. In this configuration, you do not use port binding because port binding requires that all initiator and target ports are on the same subnet.



ESXi selects vmk1 when connecting to Port 0 of Controller A and Controller B because all three ports are on the same subnet. Similarly, vmk2 is selected when connecting to Port 1 of Controller A and B. You can use NIC teaming in this configuration.

Total of four paths are created.


Paths	Description
Path 1	vmk1 and Port0 of Controller A
Path 2	vmk1 and Port0 of Controller B
Path 3	vmk2 and Port1 of Controller A
Path 4	vmk2 and Port1 of Controller B

Routing with Software iSCSI

You can use the `esxcli` command to add static routes for your iSCSI traffic. After you configure static routes, initiator and target ports in different subnets can communicate with each other.

Example 1. Using static routes with port binding

In this example, you keep all bound VMkernel ports in one subnet (N1) and configure all target portals in another subnet (N2). You can then add a static route for the target subnet (N2).




Use the following command:

```
# esxcli network ip route ipv4 add -gateway 192.168.1.253 -network 10.115.179.0/24
```

Example 2. Using static routes to create multiple paths

In this configuration, you use static routing when using different subnets. You cannot use the port binding with this configuration.



You configure vmk1 and vmk2 in separate subnets, 192.168.1.0 and 192.168.2.0. Your target portals are also in separate subnets, 10.115.155.0 and 10.115.179.0.

You can add the static route for 10.115.155.0 from vmk1. Make sure that the gateway is reachable from vmk1.

```
# esxcli network ip route ipv4 add -gateway 192.168.1.253 -network 10.115.155.0/24
```

You then add static route for 10.115.179.0 from vmk2. Make sure that the gateway is reachable from vmk2.

```
# esxcli network ip route ipv4 add -gateway 192.168.2.253 -network 10.115.179.0/24
```

When connecting with Port 0 of Controller A, vmk1 is used.

When connecting with Port 0 of Controller B, vmk2 is used.

Example 3. Routing with a separate gateway per vmkernel port

Starting with vSphere 6.5, you can configure a separate gateway per VMkernel port. If you use DHCP to obtain IP configuration for a VMkernel port, gateway information can also be obtained using DHCP.

To see gateway information per VMkernel port, use the following command:

```
# esxcli network ip interface ipv4 address list
```

Name	IPv4 Address	IPv4 Netmask	IPv4 Broadcast	Address Type	Gateway	DHCP DNS
vmk0	10.115.155.122	255.255.252.0	10.115.155.255	DHCP	10.115.155.253	true
vmk1	10.115.179.209	255.255.252.0	10.115.179.255	DHCP	10.115.179.253	true
vmk2	10.115.179.146	255.255.252.0	10.115.179.255	DHCP	10.115.179.253	true

With separate gateways per VMkernel port, you use port binding to reach targets in different subnets.

Managing iSCSI Network

Special considerations apply to network adapters, both physical and VMkernel, that are associated with an iSCSI adapter.

After you create network connections for iSCSI, an iSCSI indicator becomes enabled in the vSphere Client. The indicator shows that a particular virtual or physical network adapter is iSCSI-bound. To avoid disruptions in iSCSI traffic, follow these guidelines and considerations when managing iSCSI-bound virtual and physical network adapters:

- Make sure that the VMkernel network adapters are assigned addresses on the same subnet as the iSCSI storage portal they connect to.
- iSCSI adapters using VMkernel adapters cannot connect to iSCSI ports on different subnets, even if the iSCSI adapters discover those ports.
- When using separate vSphere switches to connect physical network adapters and VMkernel adapters, make sure that the vSphere switches connect to different IP subnets.
- If VMkernel adapters are on the same subnet, they must connect to a single vSwitch.
- If you migrate VMkernel adapters to a different vSphere switch, move associated physical adapters.
- Do not make configuration changes to iSCSI-bound VMkernel adapters or physical network adapters.
- Do not make changes that might break association of VMkernel adapters and physical network adapters. You can break the association if you remove one of the adapters or the vSphere switch that connects them. Or if you change the 1:1 network policy for their connection.

iSCSI Network Troubleshooting

A warning sign indicates non-compliant port group policy for an iSCSI-bound VMkernel adapter.

The VMkernel adapter's port group policy is considered non-compliant in the following cases:

- The VMkernel adapter is not connected to an active physical network adapter.
- The VMkernel adapter is connected to more than one physical network adapter.
- The VMkernel adapter is connected to one or more standby physical adapters.
- The active physical adapter is changed.

Set up the correct network policy for the iSCSI-bound VMkernel adapter. See [Configure Port Binding for iSCSI or iSER on ESXi](#).

Configure Port Binding for iSCSI or iSER on ESXi

The port binding creates connections for the traffic between certain types of iSCSI and iSER adapters and the physical network adapters on your ESXi system.

The following types of adapters require the port binding:

- Software iSCSI adapter
- Dependent hardware iSCSI adapter
- VMware iSCSI over RDMA (iSER) adapter

The following tasks discuss the network configuration with a vSphere Standard switch and a single physical network adapter. If you have multiple network adapters, see [Multiple Network Adapters in iSCSI or iSER Configuration](#).

Note: iSER does not support NIC teaming. When configuring port binding for iSER, use only one RDMA-enabled physical adapter (vmnic#) and one VMkernel adapter (vmk#) per vSwitch.

You can also use the VMware vSphere® Distributed Switch™ and VMware NSX® Virtual Switch™ in the port binding configuration. For information about NSX virtual switches, see the *VMware NSX Data Center for vSphere* documentation.


If you use a vSphere distributed switch with multiple uplink ports, for port binding, create a separate distributed port group per each physical NIC. Then set the team policy so that each distributed port group has only one active uplink port. For detailed information on distributed switches, see the *vSphere Networking* documentation.

Create a Single VMkernel Adapter for iSCSI or iSER

Connect the VMkernel, which runs services for iSCSI storage, to a physical network adapter on your ESXi host. You then use the created VMkernel adapter in the port binding configuration with the iSCSI or iSER adapters.

- If you are creating a VMkernel adapter for dependent hardware iSCSI, you must use the physical network adapter (vmnic#) that corresponds to the iSCSI component. See [Determine Association Between iSCSI and Network Adapters](#).
- With the iSER adapter, make sure to use an appropriate RDMA-capable vmnic#. See [Install and View an RDMA Capable Network Adapter](#).

1. In the vSphere Client, navigate to the ESXi host.
2. Select **Add Networking** from the right-click menu.
3. Select **VMkernel Network Adapter**, and click **Next**.
4. Select **New standard switch** to create a vSphere Standard switch.
5. Click the **Add adapters** icon, and select an appropriate network adapter (vmnic#) to use for iSCSI. Make sure to assign the adapter to Active Adapters.
6. Enter a network label.
A network label is a friendly name that identifies the VMkernel adapter that you are creating, for example, iSCSI or iSER.
7. Specify the IP settings.
8. Review the information and click **Finish**.
You created the virtual VMkernel adapter (vmk#) for a physical network adapter (vmnic#) on your host.
9. Verify your configuration.
 - a) Under **Networking**, select **VMkernel Adapters**, and select the VMkernel adapter (vmk#) from the list.
 - b) Click the **Policies** tab, and verify that the corresponding physical network adapter (vmnic#) appears as an active adapter under **Teaming and failover**.



If your host has one physical network adapter for iSCSI traffic, bind the VMkernel adapter that you created to the iSCSI or iSER vmhba adapter.

If you have multiple network adapters, you can create additional VMkernel adapters and then perform iSCSI binding. The number of virtual adapters must correspond to the number of physical adapters on the host. For information, see [Multiple Network Adapters in iSCSI or iSER Configuration](#).

Bind iSCSI or iSER Adapters to VMkernel Adapters

On the ESXi host, bind an iSCSI or iSER adapter with a VMkernel adapter.

Create a virtual VMkernel adapter for each physical network adapter on your host. If you use multiple VMkernel adapters, set up the correct network policy.

Required privilege: **Host > Configuration > Storage Partition Configuration**

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the appropriate iSCSI or iSER adapter (vmhba#) from the list.
4. Click the **Network Port Binding** tab and click the **Add** icon.
5. Select a VMkernel adapter to bind with the iSCSI or iSER adapter.

Note: Make sure that the network policy for the VMkernel adapter is compliant with the binding requirements.

You can bind the software iSCSI adapter to one or more VMkernel adapters. For a dependent hardware iSCSI adapter or the iSER adapter, only one VMkernel adapter associated with the correct physical NIC is available.

6. Click **OK**.

The network connection appears on the list of network port bindings for the iSCSI or iSER adapter.

The screenshot shows the 'Storage Adapters' configuration window in vSphere. The 'Network Port Binding' tab is selected, showing a table of network port bindings for the selected storage adapter (vmhba33). The table has columns for Adapter, Type, Status, Identifier, Targets, Devices, and Paths. Below the table, there are tabs for Properties, Devices, Paths, Dynamic Discovery, Static Discovery, Network Port Binding, and Advanced Options. The Network Port Binding tab shows a list of port groups and their associated VMkernel adapters, with a 'View Details' icon next to the 'VMkernel-IsER (vSwitch1)' entry.

Adapter	Type	Status	Identifier	Targets	Devices	Paths
vmhba33	Block SCSI	Unknown	--	1	2	2
Model: VMware iSCSI over RDMA (iSER) Adapter						
vmhba64	iSCSI	Unbound	IsER-vmnic9 (vSwitch1) (vSphere VMkernel-IsER)	0	0	0
vmhba65	iSCSI	Unbound	IsER-vmnic10 (vSwitch1) (vSphere VMkernel-IsER)	0	0	0
vmhba66	iSCSI	Online	IsER-vmnic4 (vSwitch1) (vSphere VMkernel-IsER)	0	0	0
vmhba67	iSCSI	Unbound	IsER-vmnic5 (vSwitch1) (vSphere VMkernel-IsER)	0	0	0
Model: Wellsburg AHCI Controller						
vmhba1	Block SCSI	Unknown	--	0	0	0

Port Group	VMkernel Adapter	Port Group Policy	Path Status	Physical Network Adapter
VMkernel-IsER (vSwitch1)	vmk1	Compliant	Not used	vmnic4 (25 Gbit/s, Full)

Review Port Binding Details on the ESXiHost

Review networking details of the VMkernel adapter that is bound to the iSCSI or iSER vmhba adapter.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the appropriate iSCSI or iSER adapter from the list.
4. Click the **Network Port Binding** tab and select the VMkernel adapter from the list.
5. Click the **View Details** icon.
6. Review the VMkernel adapter and physical adapter information by switching between available tabs.

For other configuration steps you can perform for the iSCSI or iSER storage adapters, see the following topics:

- [Configure Dependent Hardware iSCSI Adapters with ESXi](#)
- [Configure the Software iSCSI Adapter with ESXi](#)
- [Configure iSER with ESXi](#)

Using Jumbo Frames with iSCSI and iSER with ESXi

ESXi supports the use of Jumbo Frames with iSCSI and iSER protocols. Jumbo Frames are Ethernet frames with the size that exceeds 1500 Bytes. The maximum transmission unit (MTU) parameter is typically used to measure the size of Jumbo Frames.

When you use Jumbo Frames for iSCSI traffic, the following considerations apply:

- All network components must support Jumbo Frames.
- Check with your vendors to ensure your physical NICs and iSCSI adapters support Jumbo Frames.
- To set up and verify physical network switches for Jumbo Frames, consult your vendor documentation.

The following table explains the level of support that ESXi provides to Jumbo Frames.

Table 212: Support of Jumbo Frames

Type of iSCSI Adapters	Jumbo Frames Support
Software iSCSI	Supported
Dependent Hardware iSCSI	Supported. Check with vendor.
Independent Hardware iSCSI	Supported. Check with vendor.
VMware iSER	Supported. Check with vendor.

Enable Jumbo Frames for Networking

You can enable Jumbo Frames for ESXi storage adapters that use VMkernel networking for their traffic. These adapters include software iSCSI adapters, dependent hardware iSCSI adapters, and VMware iSER adapters.

To enable Jumbo Frames, change the default value of the maximum transmission units (MTU) parameter. You change the MTU parameter on the vSphere switch that you use for iSCSI traffic. For more information, see the *vSphere Networking* documentation.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Networking**, click **Virtual switches**, and select the vSphere switch that you want to modify from the list.
4. Click the **Edit settings** icon.
5. On the Properties page, change the MTU parameter.

This step sets the MTU for all physical NICs on that standard switch. Set the MTU value to the largest MTU size among all NICs connected to the standard switch. ESXi supports the MTU size of up to 9000 Bytes.

Enable Jumbo Frames for Independent Hardware iSCSI

To enable Jumbo Frames for independent hardware iSCSI adapters on your ESXi host, change the default value of the maximum transmission units (MTU) parameter.

Use the Advanced Options settings to change the MTU parameter for the iSCSI HBA.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the independent hardware iSCSI adapter from the list of adapters.
4. Click the **Advanced Options** tab and click **Edit**.
5. Change the value of the MTU parameter.

ESXi supports the MTU size up to 9000 Bytes.

Configure Dynamic or Static Discovery for iSCSI and iSER on ESXi Host

You need to set up target discovery addresses, so that the iSCSI or iSER storage adapter can determine which storage resource on the network is available for access on your ESXi host.

Required privilege: **Host > Configuration > Storage Partition Configuration**

The ESXi system supports these discovery methods:

Dynamic Discovery

Also known as SendTargets discovery. Each time the initiator contacts a specified iSCSI server, the initiator sends the SendTargets request to the server. The server responds by supplying a list of available targets to the initiator. The names and IP addresses of these targets appear on the **Static Discovery** tab. If you remove a static target added by dynamic discovery, the target might be returned to the list the next time a rescan happens, the storage adapter is reset, or the host is rebooted.

Note: With software and dependent hardware iSCSI, ESXi filters target addresses based on the IP family of the iSCSI server address specified. If the address is IPv4, IPv6 addresses that might come in the SendTargets response from the iSCSI server are filtered out. When DNS names are used to specify an iSCSI server, or when the SendTargets response from the iSCSI server has DNS names, ESXi relies on the IP family of the first resolved entry from DNS lookup.

Static Discovery

In addition to the dynamic discovery method, you can use static discovery and manually enter information for the targets. The iSCSI or iSER adapter uses a list of targets that you provide to contact and communicate with the iSCSI servers.

When you set up static or dynamic discovery, you can only add new iSCSI targets. You cannot change any parameters of an existing target. To make changes, remove the existing target and add a new one.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
4. Configure the discovery method.

Discovery Method	Description
Dynamic Discovery	1. Click Dynamic Discovery and click Add .

Discovery Method	Description
	2. Enter the IP address or DNS name of the storage system and click OK. 3. Rescan the iSCSI adapter. After establishing the SendTargets session with the iSCSI system, your host populates the Static Discovery list with all newly discovered targets. Note: A dynamically discovered target remains on the list even after it is removed from the array side.
Static Discovery	1. Click Static Discovery and click Add. 2. Enter the target's information and click OK 3. Rescan the iSCSI adapter.

For other configuration steps you can perform for the iSCSI or iSER storage adapters, see the following topics:

- [Set Up Independent Hardware iSCSI Adapters with ESXi](#)
- [Configure Dependent Hardware iSCSI Adapters with ESXi](#)
- [Configure the Software iSCSI Adapter with ESXi](#)
- [Configure iSER with ESXi](#)

Remove Dynamic or Static iSCSI Targets

Remove iSCSI servers connected to your ESXi host.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the iSCSI adapter to modify from the list.
4. Switch between **Dynamic Discovery** and **Static Discovery**.
5. Select an iSCSI server to remove and click **Remove**.
6. Rescan the iSCSI adapter.

If you are removing the static target that was dynamically discovered, you need to remove it from the storage system before performing the rescan. Otherwise, your host will automatically discover and add the target to the list of static targets when you rescan the adapter.

Configuring CHAP Parameters for iSCSI or iSER Storage Adapters on ESXi Host

The IP networks that the iSCSI technology uses to connect your ESXi host to remote targets do not protect the data they transport. As a result, you must ensure security of the connection. One of the protocols that iSCSI implements is the Challenge Handshake Authentication Protocol (CHAP). The CHAP protocol verifies the legitimacy of ESXi initiators that access targets on the network.

CHAP uses a three-way handshake algorithm to verify the identity of your host and, if applicable, of the iSCSI target when the host and target establish a connection. The verification is based on a predefined private value, or CHAP secret, that the initiator and target share.

ESXi supports CHAP authentication at the adapter level. In this case, all targets receive the same CHAP name and secret from the iSCSI initiator. For software and dependent hardware iSCSI adapters, and for iSER adapters, ESXi also supports per-target CHAP authentication, which allows you to configure different credentials for each target to achieve greater level of security.

Selecting CHAP Authentication Method

ESXi supports unidirectional CHAP for all types of iSCSI and iSER initiators, and bidirectional CHAP for software and dependent hardware iSCSI, and for iSER.

Before configuring CHAP, check whether CHAP is activated at the iSCSI storage system. Also, obtain information about the CHAP authentication method the system supports. If CHAP is activated, configure it for your initiators, making sure that the CHAP authentication credentials match the credentials on the iSCSI storage.

ESXi supports the following CHAP authentication methods:

Unidirectional CHAP

In unidirectional CHAP authentication, the target authenticates the initiator, but the initiator does not authenticate the target.

Bidirectional CHAP

The bidirectional CHAP authentication adds an extra level of security. With this method, the initiator can also authenticate the target. VMware supports this method for software and dependent hardware iSCSI adapters, and for iSER adapters.

For software and dependent hardware iSCSI adapters, and for iSER adapters, you can set unidirectional CHAP and bidirectional CHAP for each adapter or at the target level. Independent hardware iSCSI supports CHAP only at the adapter level.

When you set the CHAP parameters, specify a security level for CHAP.

Note: When you specify the CHAP security level, how the storage array responds depends on the array's CHAP implementation and is vendor-specific. For information on CHAP authentication behavior in different initiator and target configurations, consult the array documentation.

Table 213: CHAP Security Level

CHAP Security Level	Description	Supported Storage Adapters
None	The host does not use CHAP authentication. If authentication is activated, use this option to deactivate it.	Independent hardware iSCSI Software iSCSI Dependent hardware iSCSI iSER
Use unidirectional CHAP if required by target	The host prefers a non-CHAP connection, but can use a CHAP connection if required by the target.	Software iSCSI Dependent hardware iSCSI iSER
Use unidirectional CHAP unless prohibited by target	The host prefers CHAP, but can use non-CHAP connections if the target does not support CHAP.	Independent hardware iSCSI Software iSCSI Dependent hardware iSCSI iSER
Use unidirectional CHAP	The host requires successful CHAP authentication. The connection fails if CHAP negotiation fails.	Independent hardware iSCSI Software iSCSI Dependent hardware iSCSI iSER
Use bidirectional CHAP	The host and the target support bidirectional CHAP.	Software iSCSI Dependent hardware iSCSI iSER

Set Up CHAP for iSCSI or iSER Storage Adapter

When you set up CHAP name and secret at the iSCSI or iSER adapter level, all targets receive the same parameters from the adapter. By default, all discovery addresses or static targets inherit CHAP parameters that you set up at the adapter level.

- Before setting up CHAP parameters for software or dependent hardware iSCSI, determine whether to configure unidirectional or bidirectional CHAP. Independent hardware iSCSI adapters do not support bidirectional CHAP.

- Verify CHAP parameters configured on the storage side. Parameters that you configure must match the ones on the storage side.
- Required privilege: **Host > Configuration > Storage Partition Configuration**

The CHAP name cannot exceed 511 alphanumeric characters and the CHAP secret cannot exceed 255 alphanumeric characters. Some adapters, for example the QLogic adapter, might have lower limits, 255 for the CHAP name and 100 for the CHAP secret.

1. Navigate to the iSCSI or iSER storage adapter.
 - a) In the vSphere Client, navigate to the ESXi host.
 - b) Click the **Configure** tab.
 - c) Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
2. Click the **Properties** tab and click **Edit** in the **Authentication** panel.
3. Specify authentication method.
 - **None**
 - **Use unidirectional CHAP if required by target**
 - **Use unidirectional CHAP unless prohibited by target**
 - **Use unidirectional CHAP**
 - **Use bidirectional CHAP**. To configure bidirectional CHAP, you must select this option.
4. Specify the outgoing CHAP name.

Make sure that the name you specify matches the name configured on the storage side.

 - To set the CHAP name to the iSCSI adapter name, select **Use initiator name**.
 - To set the CHAP name to anything other than the iSCSI initiator name, deselect **Use initiator name** and enter a name in the **Name** text box.
5. Enter an outgoing CHAP secret to be used as part of authentication. Use the same secret that you enter on the storage side.
6. If configuring bidirectional CHAP, specify incoming CHAP credentials.

Make sure to use different secrets for the outgoing and incoming CHAP.
7. Click **OK**.
8. Rescan the iSCSI adapter.

If you change the CHAP parameters, they are used for new iSCSI sessions. For existing sessions, new settings are not used until you log out and log in again.

For other configuration steps you can perform for the iSCSI or iSER storage adapters, see the following topics:

- [Set Up Independent Hardware iSCSI Adapters with ESXi](#)
- [Configure Dependent Hardware iSCSI Adapters with ESXi](#)
- [Configure the Software iSCSI Adapter with ESXi](#)
- [Configure iSER with ESXi](#)

Set Up CHAP for Target

If you use software and dependent hardware iSCSI adapters, or an iSER storage adapter, you can configure different CHAP credentials for each discovery address or static target.

- Before setting up CHAP parameters, determine whether to configure unidirectional or bidirectional CHAP.

- Verify CHAP parameters configured on the storage side. Parameters that you configure must match the ones on the storage side.
- Required privilege: **Host > Configuration > Storage Partition Configuration**

The CHAP name cannot exceed 511 and the CHAP secret 255 alphanumeric characters.

1. Navigate to the iSCSI or iSER storage adapter.
 - a) In the vSphere Client, navigate to the ESXi host.
 - b) Click the **Configure** tab.
 - c) Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
2. Click either **Dynamic Discovery** or **Static Discovery**.
3. From the list of available targets, select a target to configure and click **Authentication**.
4. Deselect **Inherit settings from parent** and specify authentication method.
 - **None**
 - **Use unidirectional CHAP if required by target**
 - **Use unidirectional CHAP unless prohibited by target**
 - **Use unidirectional CHAP**
 - **Use bidirectional CHAP**. To configure bidirectional CHAP, you must select this option.
5. Specify the outgoing CHAP name.

Make sure that the name you specify matches the name configured on the storage side.

 - To set the CHAP name to the iSCSI adapter name, select **Use initiator name**.
 - To set the CHAP name to anything other than the iSCSI initiator name, deselect **Use initiator name** and enter a name in the **Name** text box.
6. Enter an outgoing CHAP secret to be used as part of authentication. Use the same secret that you enter on the storage side.
7. If configuring bidirectional CHAP, specify incoming CHAP credentials.

Make sure to use different secrets for the outgoing and incoming CHAP.
8. Click **OK**.
9. Rescan the storage adapter.

If you change the CHAP parameters, they are used for new iSCSI sessions. For existing sessions, new settings are not used until you log out and login again.

Configuring Advanced Parameters for iSCSI on ESXi Host

You might need to configure additional parameters for iSCSI initiators on your ESXi host. For example, some iSCSI storage systems require ARP (Address Resolution Protocol) redirection to move iSCSI traffic dynamically from one port to another. In this case, you must activate the ARP redirection on your host.

The following table lists advanced iSCSI parameters that you can configure using the vSphere Client. In addition, you can use the vSphere CLI commands to configure some of the advanced parameters. For information, see the *Getting Started with ESXCLI* documentation.

Depending on the type of your adapters, certain parameters might not be available.

Important: Do not change the advanced iSCSI settings unless VMware support or Storage Vendors direct you to change them.

Table 214: Additional Parameters for iSCSI Initiators

Advanced Parameter	Description
Header Digest	Increases data integrity. When the header digest parameter is enabled, the system performs a checksum over each header part of the iSCSI Protocol Data Unit (PDU). The system verifies the data using the CRC32C algorithm.
Data Digest	Increases data integrity. When the data digest parameter is enabled, the system performs a checksum over each PDU data part. The system verifies the data using the CRC32C algorithm. Note: Systems that use the Intel Nehalem processors offload the iSCSI digest calculations for software iSCSI. This offload helps to reduce the impact on performance.
ErrorRecoveryLevel	iSCSI Error Recovery Level (ERL) value that the iSCSI initiator on the host negotiates during a login.
LoginRetryMax	Maximum number of times the ESXi iSCSI initiator attempts to log into a target before ending the attempts.
MaxOutstandingR2T	Defines the R2T (Ready to Transfer) PDUs that can be in transition before an acknowledge PDU is received.
FirstBurstLength	Specifies the maximum amount of unsolicited data an iSCSI initiator can send to the target during the execution of a single SCSI command, in bytes.
MaxBurstLength	Maximum SCSI data payload in a Data-In or a solicited Data-Out iSCSI sequence, in bytes.
MaxRecvDataSegLength	Maximum data segment length, in bytes, that can be received in an iSCSI PDU.
MaxCommands	Maximum SCSI commands that can be queued on the iSCSI adapter.
DefaultTimeToWait	Minimum time in seconds to wait before attempting a logout or an active task reassignment after an unexpected connection termination or reset.
DefaultTimeToRetain	Maximum time in seconds, during which reassigning the active task is still possible after a connection termination or reset.
LoginTimeout	Time in seconds the initiator will wait for the login response to finish.
LogoutTimeout	Time in seconds initiator will wait to get a response for Logout request PDU.
RecoveryTimeout	Specifies the amount of time, in seconds, that can lapse while a session recovery is performed. If the timeout exceeds its limit, the iSCSI initiator ends the session.
No-Op Interval	Specifies the time interval, in seconds, between NOP-Out requests sent from your iSCSI initiator to an iSCSI target. The NOP-Out requests serve as the ping mechanism to verify that a connection between the iSCSI initiator and the iSCSI target is active.
No-Op Timeout	Specifies the amount of time, in seconds, that can lapse before your host receives a NOP-In message. The iSCSI target sends the message in response to the NOP-Out request. When the no-op timeout limit is exceeded, the initiator ends the current session and starts a new one.
ARP Redirect	With this parameter enabled, storage systems can move iSCSI traffic dynamically from one port to another. Storage systems that perform array-based failovers require the ARP parameter.
Delayed ACK	With this parameter enabled, storage systems can delay an acknowledgment of received data packets.

Configure Advanced Parameters for iSCSI on ESXi Host

The advanced iSCSI settings control such parameters as header and data digest, ARP redirection, delayed ACK, and so on.

Required privilege: **Host > Configuration > Storage Partition Configuration**

CAUTION: Do not make any changes to the advanced iSCSI settings unless you are working with the VMware support team or otherwise have thorough information about the values to provide for the settings.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
4. Configure advanced parameters.

Option	Description
At the adapter level	Click the Advanced Options tab and click Edit .
At the target level	1. Click either Dynamic Discovery or Static Discovery. 2. From the list of available targets, select a target to configure and click Advanced.

5. Enter any required values for the advanced parameters you want to modify.

Managing iSCSI Session on ESXi Host

To communicate with each other, iSCSI initiators on the ESXi host and targets establish iSCSI sessions. You can review and manage iSCSI sessions using vSphere CLI.

By default, software iSCSI and dependent hardware iSCSI initiators start one iSCSI session between each initiator port and each target port. If your iSCSI initiator or target has more than one port, your host can have multiple sessions established. The default number of sessions for each target equals the number of ports on the iSCSI adapter times the number of target ports.

Using vSphere CLI, you can display all current sessions to analyze and debug them. To create more paths to storage systems, you can increase the default number of sessions by duplicating existing sessions between the iSCSI adapter and target ports.

You can also establish a session to a specific target port. This capability is useful if your host connects to a single-port storage system that presents only one target port to your initiator. The system then redirects additional sessions to a different target port. Establishing a new session between your iSCSI initiator and another target port creates an additional path to the storage system.

The following considerations apply to iSCSI session management:

- Some storage systems do not support multiple sessions from the same initiator name or endpoint. Attempts to create multiple sessions to such targets can result in an unpredictable behavior of your iSCSI environment.
- Storage vendors can provide automatic session managers. Using the automatic session managers to add or delete sessions, does not guarantee lasting results and can interfere with the storage performance.

Review iSCSI Sessions

Use the vCLI command to display iSCSI sessions between an iSCSI adapter and a storage system.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To list iSCSI sessions, run the following command:

```
esxcli iscsi session list
```

The command takes these options:

Option	Description
-A --adapter=<i>str</i>	The iSCSI adapter name, for example, vmhba34.
-s --isid=<i>str</i>	The iSCSI session identifier.
-n --name=<i>str</i>	The iSCSI target name, for example, iqn.X.

Add iSCSI Sessions

Use the vCLI to add an iSCSI session for a target you specify or to duplicate an existing session. By duplicating sessions, you increase the default number of sessions and create additional paths to storage systems.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To add or duplicate an iSCSI session, run the following command:

```
esxcli iscsi session add
```

The command takes these options:

Option	Description
<code>-A --adapter=</code> <i>str</i>	The iSCSI adapter name, for example, vmhba34. This option is required.
<code>-s --isid=</code> <i>str</i>	The ISID of a session to duplicate. You can find it by listing all sessions.
<code>-n --name=</code> <i>str</i>	The iSCSI target name, for example, iqn.X.

Rescan the iSCSI adapter.

Remove iSCSI Sessions

Use the vCLI command to remove an iSCSI session between an iSCSI adapter and a target.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To remove a session, run the following command:

```
esxcli iscsi session remove
```

The command takes these options:

Option	Description
<code>-A --adapter=</code> <i>str</i>	The iSCSI adapter name, for example, vmhba34. This option is required.
<code>-s --isid=</code> <i>str</i>	The ISID of a session to remove. You can find it by listing all session.
<code>-n --name=</code> <i>str</i>	The iSCSI target name, for example, iqn.X.

Rescan the iSCSI adapter.

Managing ESXi Storage Devices

Manage local and networked storage device that your ESXi host has access to.

For information about storage devices that ESXi supports, see [What Types of Physical Storage Does ESXi Support](#).

Viewing Storage Devices Available to ESXi Hosts

When your ESXi host connects to block-based storage systems, LUNs or storage devices that support ESXi become available to the host.

After the devices get registered with your host, you can display all available local and networked devices and review their information. If you use third-party multipathing plug-ins, the storage devices available through the plug-ins also appear on the list.

Note: If an array supports implicit asymmetric logical unit access (ALUA) and has only standby paths, the registration of the device fails. The device can register with the host after the target activates a standby path and the host detects it as active. The advanced system `/Disk/FailDiskRegistration` parameter controls this behavior of the host.

For each storage adapter, you can display a separate list of storage devices available for this adapter.

Generally, when you review storage devices, you see the following information.

Storage Device Information	Description
Name	Also called Display Name. It is a name that the ESXi host assigns to the device based on the storage type and manufacturer. Generally, you can change this name to a name of your choice. See ESXi Storage Device Names and Identifiers .
Identifier	A universally unique identifier that is intrinsic to the device. See ESXi Storage Device Names and Identifiers .
Operational State	Indicates whether the device is attached or detached. See Detach Storage Devices .
LUN	Logical Unit Number (LUN) within the SCSI target. The LUN number is provided by the storage system. If a target has only one LUN, the LUN number is always zero (0). See Target and Device Representations .
Type	Type of device, for example, disk or CD-ROM.
Drive Type	Information about whether the device is a flash drive or a regular HDD drive. For information about flash drives and NVMe devices, see Using Flash Devices with ESXi Hosts .
Transport	Transportation protocol your host uses to access the device. The protocol depends on the type of storage being used. See What Types of Physical Storage Does ESXi Support .
Capacity	Total capacity of the storage device.
Owner	The plug-in, such as the NMP or a third-party plug-in, that the host uses to manage paths to the storage device. See Using Pluggable Storage Architecture and Path Management with ESXi .
Hardware Acceleration	Information about whether the storage device assists the host with virtual machine management operations. The status can be Supported, Not Supported, or Unknown. See Storage Hardware Acceleration in vSphere .
Sector Format	Indicates whether the device uses a traditional, 512n, or advanced sector format, such as 512e or 4Kn. See Device Sector Formats .
Location	A path to the storage device in the <code>/vmfs/devices/</code> directory.
Partition Format	A partition scheme used by the storage device. It can be of a master boot record (MBR) or GUID partition table (GPT) format. The GPT devices can support datastores greater than 2 TB. See Device Sector Formats .
Partitions	Primary and logical partitions, including a VMFS datastore, if configured. See Working with Datastores in vSphere Storage Environment .
Multipathing Policies	Path Selection Policy and Storage Array Type Policy the host uses to manage paths to storage. See Understanding Multipathing and Failover in the ESXi Environment .
Paths	Paths used to access storage and their status. See Deactivate Storage Paths .

Display Storage Devices for an ESXi Host

Display all storage devices available to a ESXi host. If you use any third-party multipathing plug-ins, the storage devices available through the plug-ins also appear on the list.

The Storage Devices view allows you to list the hosts' storage devices, analyze their information, and modify properties.

1. In the vSphere Client, navigate to the ESXi host.

2. Click the **Configure** tab.

3. Under **Storage**, click **Storage Devices**.

All storage devices available to the host are listed in the Storage Devices table.

4. To view details for a specific device, select the device from the list.

5. Use the icons to perform basic storage management tasks.

Availability of specific icons depends on the device type and configuration.

6. Use the following tabs to access additional information and modify properties for the selected device.

Tab	Description
Properties	View device properties and characteristics. View and modify multipathing policies for the device. See Viewing and Managing Storage Paths on ESXi Hosts .
Paths	Display paths available for the device. Deactivate or enable a selected path. See Deactivate Storage Paths .
Partition Details	Displays information about partitions and their formats.

Display Storage Devices for an Adapter

Display a list of storage devices accessible through a specific storage adapter on an ESXi host.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**.
All storage adapters installed on the host are listed in the Storage Adapters table.
4. Select the adapter from the list and click the **Devices** tab.
Storage devices that the host can access through the adapter are displayed.
5. Use the icons to perform basic storage management tasks.
Availability of specific icons depends on the device type and configuration.

Icon	Description
Refresh	Refresh information about storage adapters, topology, and file systems.
Detach	Detach the selected device from the host.
Attach	Attach the selected device to the host.
Rename	Change the display name of the selected device.
Turn On LED	Turn on the locator LED for the selected devices.
Turn Off LED	Turn off the locator LED for the selected devices.
Mark as Flash Disk	Mark the selected devices as flash disks.
Mark as HDD Disk	Mark the selected devices as HDD disks.
Mark as Local	Mark the selected devices as local for the host.
Mark as Remote	Mark the selected devices as remote for the host.
Erase Partitions	Erase partitions on the selected devices.
Mark as Perennially Reserved	Mark the selected device as perennially reserved.
Unmark as Perennially Reserved	Clear the perennial reservation from the selected device.

Device Sector Formats

ESXi supports storage devices with traditional and advanced sector formats. In storage, a sector is a subdivision of a track on a storage disk or device. Each sector stores a fixed amount of data.

This table introduces different storage device formats that ESXi supports.

Storage Device Format	ESXi Software Emulation	Logical Sector Size	Physical Sector Size	VMFS Datastore
512n	N/A	512	512	VMFS5 and VMFS6 (default)
512e	N/A	512	4096	VMFS5 and VMFS6 (default) Note: Local 512e storage devices do not support VMFS5.
4Kn	512	4096	4096	VMFS6

512-Byte Native Format

ESXi supports traditional 512n storage devices that use a native 512-bytes sector size.

512-Byte Emulation Format

Due to the increasing demand for larger capacities, the storage industry has introduced advanced formats, such as 512-byte emulation, or 512e. 512e is the advanced format in which the physical sector size is 4096 bytes, but the logical sector emulates 512-bytes sector size. Storage devices that use the 512e format can support legacy applications and guest operating systems. These devices serve as an intermediate step to 4Kn sector drives.

4K Native Format with Software Emulation

Another advanced format that ESXi supports is the 4Kn sector technology. In the 4Kn devices, both physical and logical sectors are 4096 bytes (4 KiB) in length. The device does not have an emulation layer, but exposes its 4Kn physical sector size directly to ESXi.

ESXi detects and registers the 4Kn devices and automatically emulates them as 512e. The device is presented to upper layers in ESXi as 512e. But the guest operating systems always see it as a 512n device. You can continue using existing VMs with legacy guest OSes and applications on the host with the 4Kn devices.

When you use 4Kn devices, the following considerations apply:

- ESXi supports only local 4Kn SAS and SATA HDDs.
- ESXi does not support 4Kn SSD and NVMe devices, or 4Kn devices as RDMs.
- ESXi can boot only from a 4Kn device with UEFI.
- You can use the 4Kn device to configure a coredump partition and coredump file.
- Only the NMP plug-in can claim the 4Kn devices. You cannot use the HPP to claim these devices.
- With vSAN, you can use only the 4Kn capacity HDDs for vSAN Hybrid Arrays. For information, see the *Administering VMware vSAN* documentation.
- Due to the software emulation layer, the performance of the 4Kn devices depends on the alignment of the I/Os. For best performance, run workloads that issue mostly 4K aligned I/Os.
- Workloads accessing the emulated 4Kn device directly using scatter-gather I/O (SGIO) must issue I/Os compatible with the 512e disk.

Determine Device Format

To determine whether the device uses the 512n, 512e, or 4Kn format, run the following command.

```
esxcli storage core device capacity list
```

The following sample output shows the format type.

Device	Physical Blocksize	Logical Blocksize	Logical Block Count	Size	For-
mat Type					

-----	-----	-----	-----	-----

naa.5000xxxxxxxx36f	512	512	2344225968	1144641 MiB 512n
naa.5000xxxxxxxx030	4096	512	3516328368	1716957 MiB 4Kn
SWE				
naa.5000xxxxxxxx8df	512	512	2344225968	1144641 MiB 512n
naa.5000xxxxxxxx4f4	4096	512	3516328368	1716957 MiB 4Kn
SWE				

ESXi Storage Device Names and Identifiers

In the ESXi environment, each storage device is identified by several names.

Device Identifiers

Depending on the type of storage, the ESXi host uses different algorithms and conventions to generate an identifier for each storage device.

Storage-provided identifiers

The ESXi host queries a target storage device for the device name. From the returned metadata, the host extracts or generates a unique identifier for the device. The identifier is based on specific storage standards, is unique and persistent across all hosts, and has one of the following formats:

- `naa.xxx`
- `eui.xxx`
- `t10.xxx`

Path-based identifier

When the device does not provide an identifier, the host generates an `mpx.path` name, where *path* represents the first path to the device, for example, `mpx.vmhba1:C0:T1:L3`. This identifier can be used in the same way as the storage-provided identifies.

The `mpx.path` identifier is created for local devices on the assumption that their path names are unique. However, this identifier is not unique or persistent, and can change after every system restart.

Typically, the path to the device has the following format:

`vmhbaAdapter:CChannel:TTarget:LLUN`

- *vmhbaAdapter* is the name of the storage adapter. The name refers to the physical adapter on the host, not to the SCSI controller used by the virtual machines.
- *CChannel* is the storage channel number. Software iSCSI adapters and dependent hardware adapters use the channel number to show multiple paths to the same target.
- *TTarget* is the target number. Target numbering is determined by the host and might change when the mappings of targets visible to the host change. Targets that are shared by different hosts might not have the same target number.
- *LLUN* is the LUN number that shows the position of the LUN within the target. The LUN number is provided by the storage system. If a target has only one LUN, the LUN number is always zero (0).

Legacy identifier

For example, `vmhba1:C0:T3:L1` represents LUN1 on target 3 accessed through the storage adapter `vmhba1` and channel 0.

In addition to the device-provided identifiers or `mpx.path` identifiers, ESXi generates an alternative legacy name for each device. The identifier has the following format:

`vml.number`

The legacy identifier includes a series of digits that are unique to the device. The identifier can be derived in part from the metadata obtained through the SCSI INQUIRY command. For nonlocal devices that do not provide SCSI INQUIRY identifiers, the `vml.number` identifier is used as the only available unique identifier.

Displaying Device Names in the vSphere CLI

You can use the `esxcli storage core device list` command to display all device names in the vSphere CLI. The output is similar to the following example:

```
# esxcli storage core device list
naa.XXX
  Display Name: DGC Fibre Channel Disk(naa.XXX)
  ...
  Other UUIDs: vml.000XXX
mpx.vmhba1:C0:T0:L0
  Display Name: Local VMware Disk (mpx.vmhba1:C0:T0:L0)
  ...
  Other UUIDs: vml.00000000000XYZ
```

NVMe Devices with NGUID Device Identifiers

For NVMe devices, ESXi generates device identifiers based on the information it retrieves from the devices. Generally, the NVMe devices support identifiers in EUI64 or NGUID formats, or use both formats. NGUID is a Namespace Globally Unique Identifier that uses the EUI64 16-byte designator format.

For the devices that support only NGUID format, the host creates two identifiers: `eui.xxx` (*NGUID*) as primary and `t10.xxx_controller_serial_number` as alternative primary.

ID Formats Supported by Device		Device Identifier Generated by Host
EUI64 ID Format	NGUID ID Format	ESXi 8.0
yes	yes	<code>t10.xxx_EUI64</code>
yes	no	<code>t10.xxx_EUI64</code>
no	yes	<code>eui.xxx</code> (<i>NGUID</i>) as primary ID <code>t10.xxx_controller_serial_number</code> as alternative primary ID

Verify Mapping Between Primary and Alternative Device Identifiers

Use the `esxcli storage core device uidmap list` command to verify the device identifiers. The output is similar to the following:

```
esxcli storage core device uidmap list
eui.0000xyz.....
  Primary UID: eui.0000xyz.....
```

```
Alternative Primary UUIDs: t10.0000abc....  
Legacy UUID: vml.0000000000766d68....  
Alternative Legacy UUIDs: vml.000000000080906....
```

Rename Storage Devices

The ESXi host assigns a display name to a storage device based on the storage type and manufacturer. You can change the display name of the device.

You cannot rename certain types of local devices.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**.
4. Select the device to rename and click **Rename**.
5. Change the device name to a friendly name.

Rescan Operations for ESXi Storage

When you perform storage management tasks or make changes in the SAN configuration, you might need to rescan storage available to your ESXi host.

When you perform VMFS datastore management operations, such as creating a VMFS datastore or RDM, adding an extent, and increasing or deleting a VMFS datastore, your host or the vCenter Server automatically rescans and updates your storage. You can deactivate the automatic rescan feature by turning off the Host Rescan Filter. See [Manage Filters for LUNs and vSphere VMFS Datastores](#).

In certain cases, you need to perform a manual rescan. You can rescan all storage available to your host or to all hosts in a folder, cluster, and data center.

If the changes you make are isolated to storage connected through a specific adapter, perform a rescan for this adapter.

Perform the manual rescan each time you make one of the following changes.

- Zone a new disk array on a SAN.
- Create new LUNs on a SAN.
- Change the path masking on a host.
- Reconnect a cable.
- Change CHAP settings (iSCSI only).
- Add or remove discovery or static addresses (iSCSI only).
- Add a single host to the vCenter Server after you have edited or removed from the vCenter Server a datastore shared by the vCenter Server hosts and the single host.

Important: If you rescan when a path is unavailable, the host removes the path from the list of paths to the device. The path reappears on the list as soon as it becomes available and starts working again.

Perform Storage Rescan

When you make changes in your SAN configuration, you might need to rescan storage. You can rescan all storage available to your ESXi host, cluster, or data center. If the changes you make are isolated to storage accessed through a specific host, perform the rescan for only this host.

1. In the vSphere Client object navigator, browse to a host, a cluster, a data center, or a folder that contains hosts.
2. From the right-click menu, select **Storage > Rescan Storage**.
3. Specify extent of rescan.

Option	Description
Scan for New Storage Devices	Rescan all adapters to discover new storage devices. If new devices are discovered, they appear in the device list.
Scan for New VMFS Volumes	Rescan all storage devices to discover new datastores that have been added since the last scan. Any new datastores appear in the datastore list.

Perform Adapter Rescan

When you make changes in your SAN configuration and these changes are isolated to storage accessed through a specific adapter on ESXi host, perform rescan for only this adapter.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter to rescan from the list.
4. Click the **Rescan Adapter** icon.

Change the Number of Scanned Storage Devices

The range of scanned LUN IDs for an ESXi host can be from 0 to 16,383. ESXi ignores LUN IDs greater than 16,383. The configurable `Disk.MaxLUN` parameter controls the range of scanned LUN ID range. The parameter has a default value of 1024.

The `Disk.MaxLUN` parameter also determines how many LUNs the SCSI scan code attempts to discover using individual INQUIRY commands if the SCSI target does not support direct discovery using REPORT_LUNS.

You can modify the `Disk.MaxLUN` parameter depending on your needs. For example, if your environment has a smaller number of storage devices with LUN IDs from 1 through 100, set the value to 101. As a result, you can improve device discovery speed on targets that do not support REPORT_LUNS. Lowering the value can shorten the rescan time and boot time. However, the time to rescan storage devices might also depend on other factors, including the type of the storage system and the load on the storage system.

In other cases, you might need to increase the value if your environment uses LUN IDs that are greater than 1023.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. In the Advanced System Settings table, select **Disk.MaxLUN** and click the **Edit** icon.
5. Change the existing value to the value of your choice, and click **OK**.
The value you enter specifies the LUN ID that is after the last one you want to discover.
For example, to discover LUN IDs from 1 through 100, set **Disk.MaxLUN** to 101.

Identifying ESXi Storage Connectivity Problems

When your ESXi connects to a storage device, it might experience a connectivity problem. Storage connectivity problems might be caused by various reasons. Although ESXi cannot always detect the reason for a device or its paths being unavailable, the host can determine whether the problem is permanent or temporary. In other words, the host can differentiate between a permanent device loss (PDL) state of the device and a transient all paths down (APD) state of storage.

Permanent Device Loss (PDL)

A condition that occurs when a storage device permanently fails or is administratively removed or excluded. It is not expected to become available. When the device becomes permanently unavailable, ESXi receives appropriate sense codes or a login rejection from storage arrays, and is able to recognize that the device is permanently lost.

All Paths Down (APD)

A condition that occurs when a storage device becomes inaccessible to the host and no paths to the device are available. ESXi treats this as a transient condition because typically the problems with the device are temporary and the device is expected to become available again.

Connectivity Problems and vSphere High Availability

When the device enters the PDL or APD state, vSphere High Availability (HA) can detect connectivity problems and provide automated recovery for affected virtual machines on the ESXi host. vSphere HA uses VM Component Protection (VMCP) to protect virtual machines running on the host in the vSphere HA cluster against accessibility failures. For more information about VMCP and how to configure responses for datastores and virtual machines when the APD or PDL condition occurs, see the *vSphere Availability* documentation.

Detecting PDL Conditions

A storage device is considered to be in the permanent device loss (PDL) state when it becomes permanently unavailable to your ESXi host.

Typically, the PDL condition occurs when a device is unintentionally removed, or its unique ID changes, or when the device experiences an unrecoverable hardware error.

When the storage array determines that the device is permanently unavailable, it sends SCSI error sense codes or NVMe error codes to the ESXi host. After receiving these errors, your host recognizes the device as failed and registers the state of the device as PDL. For the device to be considered permanently lost, the sense codes must be received on all its paths.

After registering the PDL state of the device, the host stops attempts to reestablish connectivity or to send commands to the device.

The vSphere Client displays the following information for the device:

- The operational state of the device changes to `Lost Communication`.
- All paths are shown as `Dead`.
- Datastores on the device are not available.

If no open connections to the device exist, or after the last connection closes, the host removes the PDL device and all paths to the device. You can deactivate the automatic removal of paths by setting the advanced host parameter `Disk.AutoremoveOnPDL` to 0.

If the device returns from the PDL condition, the host can discover it, but treats it as a new device. Data consistency for virtual machines on the recovered device is not guaranteed.

Note: When a device fails without sending appropriate SCSI sense codes or NVMe error codes or an iSCSI login rejection, the host cannot detect PDL conditions. In this case, the host continues to treat the device connectivity problems as APD even when the device fails permanently.

Permanent Device Loss and SCSI Sense Codes

The following VMkernel log example of a SCSI sense code indicates that the device is in the PDL state.

```
H:0x0 D:0x2 P:0x0 Valid sense data: 0x5 0x25 0x0 or Logical Unit Not Supported
```

Permanent Device Loss and NVMe Error Codes

The following VMkernel log example of an NVMe error code indicates that the device is in the PDL state.

```
H:0x0 D:0xb P:0x0 or H:0x0 D:0x11a P:0x0
```

Permanent Device Loss and iSCSI

On iSCSI arrays with a single LUN per target, PDL is detected through iSCSI login failure. An iSCSI storage array rejects your host's attempts to start an iSCSI session with a reason `Target Unavailable`. As with the sense codes, this response must be received on all paths for the device to be considered permanently lost.

Permanent Device Loss and Virtual Machines

After registering the PDL state of the device, the host closes all I/O from virtual machines. vSphere HA can detect PDL and restart failed virtual machines.

Performing Planned Storage Device Removal

When a storage device is malfunctioning, you can avoid permanent device loss (PDL) conditions or all paths down (APD) conditions. Perform a planned removal and reconnection of a storage device.

Planned device removal is an intentional disconnection of a storage device. You might also plan to remove a device for such reasons as upgrading your hardware or reconfiguring your storage devices. When you perform an orderly removal and reconnection of a storage device, you complete a number of tasks.

Task	Description
Migrate virtual machines from the device you plan to detach.	<i>vCenter Server and Host Management</i>
Unmount the datastore deployed on the device.	See Unmount Datastores .
Detach the storage device.	See Detach Storage Devices .
For an iSCSI device with a single LUN per target, delete the static target entry from each iSCSI HBA that has a path to the storage device.	See Remove Dynamic or Static iSCSI Targets .
Perform any necessary reconfiguration of the storage device by using the array console.	See your vendor documentation.
Reattach the storage device.	See Attach Storage Devices .
Mount the datastore and restart the virtual machines.	See Mount Datastores .

Detach Storage Devices

Safely detach a storage device from your ESXi host.

- The device does not contain any datastores.

- No virtual machines use the device as an RDM disk.
- The device does not contain a diagnostic partition or a scratch partition.

You might need to detach the device to make it inaccessible to your host, when, for example, you perform a hardware upgrade on the storage side.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**.
4. Select the device to detach and click the **Detach** icon.

The device becomes inaccessible. The operational state of the device changes to Unmounted.

If multiple hosts share the device, detach the device from each host.

Attach Storage Devices

Reattach a storage device that you previously detached from the ESXi host.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**.
4. Select the detached storage device and click the **Attach** icon.

The device becomes accessible.

Recovering from PDL Conditions

An unplanned permanent device loss (PDL) condition occurs when a storage device becomes permanently unavailable without being properly detached from the ESXi host.

The following items in the vSphere Client indicate that the device is in the PDL state:

- The datastore deployed on the device is unavailable.
- Operational state of the device changes to Lost Communication.
- All paths are shown as Dead.
- A warning about the device being permanently inaccessible appears in the VMkernel log file.

To recover from the unplanned PDL condition and remove the unavailable device from the host, perform the following tasks.

Task	Description
Power off and unregister all virtual machines that are running on the datastore affected by the PDL condition.	See <i>vSphere Virtual Machine Administration</i> .
Unmount the datastore.	See Unmount Datastores .
Rescan all ESXi hosts that had access to the device. Note: If the rescan is not successful and the host continues to list the device, some pending I/O or active references to the device might still exist. Check for any items that might still have an active reference to the device or datastore. The items include virtual machines, templates, ISO images, raw device mappings, and so on.	See Perform Storage Rescan .

Handling Transient APD Conditions

A storage device is considered to be in the all paths down (APD) state when it becomes unavailable to your ESXi host for an unspecified time period.

The reasons for an APD state can be, for example, a failed switch or a disconnected storage cable.

In contrast with the permanent device loss (PDL) state, the host treats the APD state as transient and expects the device to be available again.

The host continues to retry issued commands in an attempt to reestablish connectivity with the device. If the host's commands fail the retries for a prolonged period, the host might be at risk of having performance problems. Potentially, the host and its virtual machines might become unresponsive.

To avoid these problems, your host uses a default APD handling feature. When a device enters the APD state, the host turns on a timer. With the timer on, the host continues to retry non-virtual machine commands for a limited time period only.

By default, the APD timeout is set to 140 seconds. This value is typically longer than most devices require to recover from a connection loss. If the device becomes available within this time, the host and its virtual machine continue to run without experiencing any problems.

If the device does not recover and the timeout ends, the host stops its attempts at retries and stops any non-virtual machine I/O. Virtual machine I/O continues retrying. The vSphere Client displays the following information for the device with the expired APD timeout:

- The operational state of the device changes to `Dead or Error`.
- All paths are shown as `Dead`.
- Datastores on the device are dimmed.

Even though the device and datastores are unavailable, virtual machines remain responsive. You can power off the virtual machines or migrate them to a different datastore or host.

If later the device paths become operational, the host can resume I/O to the device and end the special APD treatment.

Deactivate Storage APD Handling

The storage all paths down (APD) handling on your ESXi host is activated by default. When this functionality is activated and a storage device enters the APD state, the host continues to retry nonvirtual machine I/O commands only for a limited time period. After the time period expires, the host stops its retry attempts and terminates any nonvirtual machine I/O. You can deactivate the APD handling feature on your host.

If you deactivate the APD handling, the host will indefinitely continue to retry issued commands in an attempt to reconnect to the APD device. This behavior might cause virtual machines on the host to exceed their internal I/O timeout and become unresponsive or fail. The host might become disconnected from vCenter Server.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. In the Advanced System Settings table, select the **Misc.APDHandlingEnable** parameter and click the `Edit` icon.
5. Change the value to 0.

If you deactivate the APD handling, you can reactivate it and set its value to 1 when a device enters the APD state. The internal APD handling feature turns on immediately and the timer starts with the current timeout value for each device in APD.

Change Timeout Limits for Storage APD

The timeout parameter controls how many seconds the ESXi host must retry the I/O commands to a storage device in an all paths down (APD) state. You can change the default timeout value.

The timeout period begins immediately after the device enters the APD state. After the timeout ends, the host marks the APD device as unreachable. The host stops its attempts to retry any I/O that is not coming from virtual machines. The host continues to retry virtual machine I/O.

By default, the timeout parameter on your host is set to 140 seconds. You can increase the value of the timeout if, for example, storage devices connected to your ESXi host take longer than 140 seconds to recover from a connection loss.

Note: If you change the timeout parameter after the device becomes unavailable, the change does not take effect for that particular APD incident.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. In the Advanced System Settings table, select the **Misc.APDTimeout** parameter and click the **Edit** icon.
5. Change the default value.

You can enter a value between 20 and 99999 seconds.

Verify the Connection Status of a Storage Device on ESXi Host

Use the `esxcli` command to verify the connection status of a particular storage device.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

1. Run the `esxcli storage core device list -d=device_ID` command.
2. Review the connection status in the `Status:` area.
 - `on` - Device is connected.
 - `dead` - Device has entered the APD state. The APD timer starts.
 - `dead timeout` - The APD timeout has expired.
 - `not connected` - Device is in the PDL state.

Use the Locator LED on ESXi Storage Devices

Use the locator LED to identify specific storage devices connected to your ESXi host, so that you can locate them among other devices. You can turn the locator LED on or off.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**.
4. From the list of storage devices, select one or more disks and set the locator LED indicator.
 - Click the **Turn On LED** icon.
 - Click the **Turn Off LED** icon.

Erase ESXi Storage Devices

Certain functionalities, such as vSAN or virtual flash resource, require that your ESXi host uses clean storage devices. You can erase an HDD or flash device and remove all existing data.

- Make sure that the host is in the connected state.

- Verify that the devices you plan to erase are not in use.
 - Required privilege: **Host.Config.Storage**
1. In the vSphere Client, navigate to the ESXi host.
 2. Click the **Configure** tab.
 3. Under **Storage**, click **Storage Devices**.
 4. Select one or more devices and click the **Erase Partitions** icon.
 5. Verify that the partition information you are erasing is not critical.
 6. Click **OK** to confirm your change.

Change Perennial Reservation Settings on ESXiHosts

You can adjust the perennial reservation setting on ESXi storage devices that are used as physical Raw Device Mappings (RDMs) in Windows Server Failover Clustering (WSFC) configurations.

Before marking a device as perennially reserved, make sure the device does not contain a VMFS datastore.

WSFC cluster nodes that are spread over several hosts require physical RDMs. The RDMs are shared among all hosts where cluster nodes run. The host with the active node holds persistent SCSI-3 reservations on all shared RDM devices. When the active node is running and devices are locked, no other host can write to the devices. If another participating host boots while the active node is holding the lock on the devices, the boot might take unusually long time because the host unsuccessfully attempts to contact the locked devices. The same issue might also affect rescan operations.

To prevent this problem, activate perennial reservation for all devices on the ESXi hosts where secondary WSFC nodes with RDMs reside. This setting informs the ESXi host about the permanent SCSI reservation on the devices, so that the host can skip the devices during the boot or storage rescan process.

If you later re-purpose the marked devices as VMFS datastores, remove the reservation to avoid unpredictable datastore behavior.

For information about WSFC clusters, see the *Setup for Windows Server Failover Clustering* documentation.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**.
4. From the list of storage devices, select the device and click one of the following icons.

Option	Description
Mark as Perennially Reserved	Mark the selected device as perennially reserved. Note: Repeat the procedure for each RDM device that is participating in the WSFC cluster.
Unmark as Perennially Reserved	Clear perennial reservation for the device that was previously marked.

The configuration is permanently stored with the ESXi host and persists across restarts.

You can also use the `esxcli` command to mark the devices participating in the WSFC cluster.

1. Mark the devices as perennially reserved.

```
esxcli storage core device setconfig -dnaa.id--perennially-reserved=true
```
2. Verify that the device is perennially reserved.

```
esxcli storage core device list -dnaa.id
```

In the output of the `esxcli` command, search for the entry `Is Perennially Reserved: true`.

3. To remove the perennially reserved flag, run the following command.

```
esxcli storage core device setconfig -dnaa.id--perennially-reserved=false
```

Using Flash Devices with ESXi Hosts

In addition to regular storage hard disk drives (HDDs), ESXi supports flash storage devices.

Unlike regular HDDs that are electromechanical devices containing moving parts, flash devices use semiconductors as their storage medium and have no moving parts. Typically, the flash devices are resilient and provide faster access to data.

To detect flash devices, ESXi uses an inquiry mechanism based on T10 standards. Check with your vendor whether your storage array supports the ESXi mechanism of flash device detection.

After the host detects the flash devices, you can use them for several tasks and functionalities.

If you use NVMe storage, enable the high-performance plug-in (HPP) to improve your storage performance. See [VMware High Performance Plug-In and Path Selection Schemes](#).

For specifics about using NVMe storage with ESXi, see [About VMware NVMe Storage](#).

Table 215: Using Flash Devices with ESXi

Functionality	Description
vSAN	vSAN requires flash devices. For more information, see the <i>Administering VMware vSAN</i> documentation.
VMFS Datastores	Create VMFS datastores on flash devices. Use the datastores for the following purposes: - Store virtual machines. Certain guest operating systems can identify virtual disks stored on these datastores as flash virtual disks. - Allocate datastore space for the ESXi host swap cache. See Configure Host Cache with VMFS Datastore in vSphere Environment.
Virtual Flash Resource (VFFS)	If required by your vendor, set up a virtual flash resource and use it for I/O caching filters. See Filtering Virtual Machine I/O in vSphere .

Flash Devices and Virtual Machines

Guest operating systems can identify virtual disks that reside on flash-based datastores as flash virtual disks.

Guest operating systems can use standard inquiry commands such as SCSI VPD Page (B1h) for SCSI devices and ATA IDENTIFY DEVICE (Word 217) for IDE devices.

For linked clones, native snapshots, and delta-disks, the inquiry commands report the virtual flash status of the base disk.

Operating systems can detect that a virtual disk is a flash disk under the following conditions:

- Detection of flash virtual disks is supported on VMs with virtual hardware version 8 or later.
- Devices backing a shared VMFS datastore must be marked as flash on all hosts.
- If the VMFS datastore includes several device extents, all underlying physical extents must be flash-based.

Best Practices for Flash Devices on ESXi Hosts

Follow these best practices when you use flash devices in vSphere environment.

- Use flash devices approved by the *VMware Compatibility Guide*.

You can monitor certain critical flash device parameters, including Media Wearout Indicator, Temperature, and Reallocated Sector Count, from an ESXi host.

Use the `esxcli` command to monitor flash devices.

```
esxcli storage core device smart get -d=flash device ID
```

Note the number of days passed since the last reboot of your ESXi host. For example, ten days.

Typically, storage vendors provide reliable lifetime estimates for a flash device under ideal conditions. For example, a vendor might guarantee a lifetime of 5 years under the condition of 20 GB writes per day. However, the more realistic life expectancy of the device depends on how many writes per day your ESXi host actually generates. Follow these steps to calculate the lifetime of the flash device.

1. Obtain the total number of blocks written to the flash device since the last reboot.

Run the `esxcli storage core device stats get -d=device ID` command. For example:

```
~ # esxcli storage core device stats get -d t10.xxxxxxxxxxxxxxx
Device: t10.xxxxxxxxxxxxxxx
Successful Commands: xxxxxxxx
Blocks Read: xxxxxxxx
Blocks Written: 629145600
Read Operations: xxxxxxxx
```

The **Blocks Written** item in the output shows the number of blocks written to the device since the last reboot. In this example, the value is 629,145,600. After each reboot, it resets to 0.

2. Calculate the total number of writes and convert to GB.

One block is 512 bytes. To calculate the total number of writes, multiply the Blocks Written value by 512, and convert the resulting value to GB.

In this example, the total number of writes since the last reboot is approximately 322 GB.

3. Estimate the average number of writes per day in GB.

Divide the total number of writes by the number of days since the last reboot.

If the last reboot was ten days ago, you get 32 GB of writes per day. You can average this number over the time period.

4. Estimate lifetime of your device by using the following formula:

vendor provided number of writes per day times *vendor provided life span* divided by *actual average number of writes per day*

For example, if your vendor guarantees a lifetime of 5 years under the condition of 20 GB writes per day, and the actual number of writes per day is 30 GB, the life span of your flash device will be approximately 3.3 years.

Keeping Flash Disks VMFS-Free

If you use the auto-partitioning boot option when installing or auto-deploying ESXi, the auto-partitioning option creates a VMFS datastore on your host's local storage. In certain cases, you need to keep your local storage flash disks unformatted.

By default, auto-partitioning deploys VMFS file systems on any unused local storage disks on your host, including flash disks.

However, a flash disk formatted with VMFS becomes unavailable for such features as virtual flash and vSAN. Both features require an unformatted flash disk and neither can share the disk with any other file system.

To ensure that auto-partitioning does not format the flash disk with VMFS, use the following boot options when you install ESXi or boot the ESXi host for the first time:

- `autoPartition=TRUE`
- `skipPartitioningSsds=TRUE`

If you use Auto Deploy, set these parameters on a reference ESXi host.

1. In the vSphere Client, navigate to the host to use as a reference host and click the **Configure** tab.
2. Click **System** to open the system options, and click **Advanced System Settings**.
3. Set the following items.

Parameter	Value
VMkernel.Boot.autoPartition	True
VMkernel.Boot.skipPartitioningSsds	True

4. Reboot the host.

If flash disks that you plan to use with the virtual flash resource and vSAN already have VMFS datastores, remove the datastores.

Marking Storage Devices on ESXi Hosts

When you configure vSAN or set up a virtual flash resource, your storage environment must include local flash devices. However, ESXi might not recognize certain storage devices as flash devices when their vendors do not support automatic flash device detection. In other cases, certain devices might not be detected as local, and ESXi marks them as remote. When devices are not recognized as the local flash devices, they are excluded from the list of devices offered for vSAN or virtual flash resource. Marking these devices as local flash makes them available for vSAN and virtual flash resource.

Mark Storage Devices as Flash

If ESXi does not recognize its devices as flash, mark them as flash devices.

Verify that the device is not in use.

ESXi does not recognize certain devices as flash when their vendors do not support automatic flash disk detection. The Drive Type column for the devices shows HDD as their type.

CAUTION: Marking the HDD devices as flash might deteriorate the performance of datastores and services that use them. Mark the devices only if you are certain that they are flash devices.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**.
4. From the list of storage devices, select one or several HDD devices and click the **Mark as Flash Disk**  icon.
5. Click **Yes** to save your changes.

The type of the devices changes to flash.

If the flash device that you mark is shared among multiple hosts, make sure that you mark the device from all hosts that share the device.

Mark Storage Devices as Local

ESXi enables you to mark devices as local. This action is useful in cases when ESXi is unable to determine whether certain devices are local.

- Make sure that the device is not shared.
 - Power off virtual machines that reside on the device and unmount an associated datastore.
1. In the vSphere Client, navigate to the ESXi host.
 2. Click the **Configure** tab.
 3. Under **Storage**, click **Storage Devices**.
 4. From the list of storage devices, select one or several remote devices and click the **Mark as Local** icon.
 5. Click **Yes** to save your changes.

About Virtual Flash Resource in vSphere Environment

You can aggregate local flash devices on an ESXi host into a single virtualized caching layer called virtual flash resource. When you set up the virtual flash resource, you create a new file system, Virtual Flash File System (VFFS). VFFS is a derivative of VMFS, which is optimized for flash devices and is used to group the physical flash devices into a single caching resource pool. As a non-persistent resource, it cannot be used to store virtual machines.

After you set up the virtual flash resource, you can use it for I/O caching filters. See [Filtering Virtual Machine I/O in vSphere](#).

Considerations for Virtual Flash Resource

When you configure a virtual flash resource, several considerations apply.

- You can have only one virtual flash resource on a single ESXi host. The virtual flash resource is managed at the host's level.
- You cannot use the virtual flash resource to store virtual machines. Virtual flash resource is a caching layer only.
- You can use only local flash devices for the virtual flash resource.
- You can create the virtual flash resource from mixed flash devices. All device types are treated equally and no distinction is made between SAS, SATA, or PCI express connectivity. When creating the resource from mixed flash devices, make sure to group similar performing devices together to maximize performance.
- You cannot use the same flash devices for the virtual flash resource and vSAN. Each requires its own exclusive and dedicated flash device.

Set Up Virtual Flash Resource

You can set up a virtual flash resource or add capacity to existing virtual flash resource.

To set up a virtual flash resource, you use local flash devices connected to your host or host cluster. To increase the capacity of your virtual flash resource, you can add more devices, up to the maximum number indicated in the *Configuration Maximums* documentation. An individual flash device must be exclusively allocated to the virtual flash resource. No other vSphere functionality, such as vSAN or VMFS, can share the device with the virtual flash resource.

1. In the vSphere Client, navigate to the host.
2. Click the **Configure** tab.
3. Under **Virtual Flash**, click **Virtual Flash Resource Management**.
4. Click one of the following options.

Option	Description
Add Capacity	If you are creating the virtual flash resource on an individual host.
Add Capacity on Cluster	If you are creating the virtual flash resource on a cluster.

5. From the list of available entities, select one or more to use for the virtual flash resource and click **OK**.

If your flash devices do not appear on the list, see [Marking Storage Devices on ESXi Hosts](#).

Option	Description
Local VMware Disk	Select any combination of unclaimed flash devices. ESXi creates the VFFS volume on one of the devices and then extends the volume on the rest of the devices. The system configures the virtual flash resource on the entire VFFS volume. If a VFFS volume exists on your host, you cannot select any unclaimed devices without first selecting the existing VFFS volume.
volume ID - Configure using the existing VFFS volume extents	If you previously created a VFFS volume on one of the host's flash devices using the <code>vmkfstools</code> command, the volume also appears on the list of eligible entities. You can select just this volume for the virtual flash resource. Or combine it with the unclaimed devices. ESXi uses the existing VFFS volume to extend it over other devices.

The virtual flash resource is created. The Device Backing area lists all devices that you use for the virtual flash resource. Use the virtual flash resource for I/O caching filters developed through vSphere APIs for I/O Filtering. See [Using Cache I/O Filters with Flash Storage Devices on ESXi Hosts](#).

You can increase the capacity by adding more flash devices to the virtual flash resource.

Remove Virtual Flash Resource

You might need to remove a virtual flash resource deployed on local flash devices connected to the ESXi host. Removing the virtual flash resource frees the devices for other services.

- Verify that the virtual flash resource is not used for I/O filters.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Virtual Flash**, click **Virtual Flash Resource Management** and click **Remove All**.

After you remove the virtual flash resource and erase the flash device, the device is available for other operations.

Set Alarm for Virtual Flash Use

Set an alarm to indicate when the use of a virtual flash resource on your ESXi host exceeds specified threshold.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. Select the setting to change and click the **Edit** button.

Parameter	Description
VFLASH.ResourceUsageThreshold	The system triggers the Host vFlash resource usage alarm when a virtual flash resource use exceeds the threshold. The default threshold is 80%. You can change the threshold to an appropriate value. The alarm is cleared when the virtual flash resource use drops below the threshold.

5. Click **OK**.

Configure Host Cache with VMFS Datastore in vSphere Environment

Enable your ESXi host to swap to the host cache. You can also change the amount of space allocated for the host cache.

Create a VMFS datastore using flash devices as backing. See [Create a vSphere VMFS Datastore](#).

Your ESXi hosts can use a portion of a flash-backed storage entity as a swap cache shared by all virtual machines.

The host-level cache is made up of files on a low-latency disk that ESXi uses as a write-back cache for virtual machine swap files. All virtual machines running on the host share the cache. Host-level swapping of virtual machine pages makes the best use of potentially limited flash device space.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Host Cache Configuration**.
4. Select the flash datastore in the list and click the **Edit** icon.
5. Allocate appropriate space for host cache.
6. Click **OK**.

About VMware NVMe Storage

Non-Volatile Memory (NVM) storage devices that use persistent memory have become popular in data centers. To connect to local and networked NVM devices, ESXi uses the NVM Express (NVMe) protocol, a standardized protocol designed specifically for high-performance multi-queue communication with NVM devices.

VMware NVMe Concepts

Before you begin working with NVMe storage in the ESXi environment, learn about basic NVMe concepts.

NVM Express (NVMe)

NVMe is a method for connecting and transferring data between a host and a target storage system. NVMe is designed for use with faster storage media equipped with non-volatile memory, such as flash devices. This type of storage can achieve low latency, low CPU usage, and high performance, and generally serves as an alternative to SCSI storage.

NVMe Transports

The NVMe storage can be directly attached to a host using a PCIe interface or indirectly through different fabric transports. VMware

NVMe over Fabrics (NVMe-oF) provides a distance connectivity between a host and a target storage device on a shared storage array.

The following types of transports for NVMe currently exist.

NVMe Transport	ESXi Support
NVMe over PCIe	Local storage.
NVMe over RDMA	Shared NVMe-oF storage. With the RoCE v2 technology.
NVMe over Fibre Channel	Shared NVMe-oF storage.
NVMe over TCP	Shared NVMe-oF storage.

For specific configuration requirements for each transport, see [Requirements for VMware NVMe Storage](#).

NVMe Namespaces

In the NVMe storage array, a namespace is a storage volume backed by some quantity of non-volatile memory. In the context of ESXi, the namespace is analogous to a storage device, or LUN. After your ESXi host discovers the NVMe namespace, a flash device that represents the namespace appears on the list of storage devices in the vSphere Client. You can use the device to create a datastore and store virtual machines.

Targets can present namespaces, equivalent to LUNs in SCSI, to a host in active/active or asymmetric access modes. ESXi can discover and use namespaces presented in either way. ESXi supports *NVMe end to end*, and *without any NVMe to SCSI* emulation.

NVMe Controllers

A controller is associated with one or several NVMe namespaces and provides an access path between the ESXi host and the namespaces in the storage array. To access the controller, the host can use two mechanisms, controller discovery and controller connection. For information, see [Add Controllers for NVMe over Fabrics](#).

Controller Discovery

With this mechanism, the ESXi host first contacts a discovery controller. The discovery controller returns a list of available controllers. After you select a controller for your host to access, all namespaces associated with this controller become available to your host.

Controller Connection

Your ESXi host connects to the controller that you specify. All namespaces associated with this controller become available to your host.

NVMe Subsystem

Generally, an NVMe subsystem is a storage array that might include several NVMe controllers, several namespaces, a non-volatile memory storage medium, and an interface between the controller and non-volatile memory storage medium. The subsystem is identified by a subsystem NVMe Qualified Name (NQN).

VMware High-Performance Plug-in (HPP)

By default, the ESXi host uses the HPP to claim the NVMe-oF targets. When selecting physical paths for I/O requests, the HPP applies an appropriate Path Selection Scheme (PSS). For information about the HPP, see [VMware High Performance](#).


[Plug-In and Path Selection Schemes](#). To change the default path selection mechanism, see [Change the Path Selection Policy](#).

Basic VMware NVMe Architecture and Components

ESXi supports local NVMe over PCIe storage and shared NVMe-oF storage, such as NVMe over Fibre Channel, NVMe over RDMA (RoCE v2), and NVMe over TCP.

VMware NVMe over PCIe


In this configuration, your ESXi host uses a PCIe storage adapter to access one or more local NVMe storage devices. After you install the adapter on the host, the host discovers available NVMe devices, and they appear in the list of storage devices in the vSphere Client.



VMware NVMe over Fibre Channel

This technology maps NVMe onto the Fibre Channel protocol to enable the transfer of data and commands between a host and a target storage device. This transport can use existing Fibre Channel infrastructure upgraded to support NVMe.


To access the NVMe over Fibre Channel storage, install a Fibre Channel storage adapter that supports NVMe on your ESXi host. You do not need to configure the adapter. It automatically connects to an appropriate NVMe subsystem and discovers all shared NVMe storage devices that it can reach. You can later reconfigure the adapter and disconnect its controllers or connect other controllers that were not available during the host boot. For more information, see [Add Controllers for NVMe over Fabrics](#).



NVMe over RDMA (RoCE v2)


This technology uses a remote direct memory access (RDMA) transport between two systems on the network. The transport enables data exchange in the main memory bypassing the operating system or the processor of either system. ESXi supports RDMA over Converged Ethernet v2 (RoCE v2) technology, which enables a remote direct memory access over an Ethernet network.

To access storage, the ESXi host uses an RDMA network adapter installed on your host and a software NVMe over RDMA storage adapter. You must configure both adapters to use them for storage discovery. For more information, see [Configuring NVMe over RDMA \(RoCE v2\) on ESXi](#).



NVMe over TCP

This technology uses Ethernet connections between two systems. To access storage, the ESXi host uses a network adapter installed on your host and a software NVMe over TCP storage adapter. You must configure both adapters to use them for storage discovery. For more information, see [Configuring NVMe over TCP on ESXi](#).



Requirements for VMware NVMe Storage

Learn how to prepare your ESXi environment for NVMe storage. Configuration requirements might change depending on the type of NVMe transport you use. If you use NVMe over RDMA (RoCE v2), in addition to general requirements, you must also configure lossless Ethernet network.

Requirements for NVMe over PCIe

Your ESXi storage environment must include the following components:

- Local NVMe storage devices.
- Compatible ESXi host.
- Hardware NVMe over PCIe adapter. After you install the adapter, your ESXi host detects it and displays in the vSphere Client as a storage adapter (vmhba) with the protocol indicated as PCIe. You do not need to configure the adapter.

Requirements for NVMe over RDMA (RoCE v2)

- NVMe storage array with NVMe over RDMA (RoCE v2) transport support.
- Compatible ESXi host.
- Ethernet switches supporting a lossless network.

- Network adapter that supports RDMA over Converged Ethernet (RoCE v2). To configure the adapter, see [Configuring NVMe over RDMA \(RoCE v2\) on ESXi](#).
- Software NVMe over RDMA adapter. This software component must be enabled on your ESXi host and connected to an appropriate network RDMA adapter. For information, see [Add Software NVMe over RDMA or NVMe over TCP Adapters](#).
- NVMe controller. You must add a controller after you configure the software NVMe over RDMA adapter. See [Add Controllers for NVMe over Fabrics](#).
- Lossless Ethernet. See [Configuring Lossless Ethernet for NVMe over RDMA](#).

Requirements for NVMe over Fibre Channel

- Fibre Channel storage array that supports NVMe. For information, see [Using ESXi with Fibre Channel SAN](#).
- Compatible ESXi host.
- Hardware NVMe adapter. Typically, it is a Fibre Channel HBA that supports NVMe. When you install the adapter, your ESXi host detects it and displays in the vSphere Client as a standard Fibre Channel adapter (vmhba) with the storage protocol indicated as NVMe. You do not need to configure the hardware NVMe adapter to use it.
- NVMe controller. You do not need to configure the controller. After you install the required hardware NVMe adapter, it automatically connects to all targets and controllers that are reachable at the moment. You can later disconnect the controllers or connect other controllers that were not available during the host boot. See [Add Controllers for NVMe over Fabrics](#).

Requirements for NVMe over TCP

- NVMe storage array with NVMe over TCP transport support.
- Compatible ESXi host.
- An Ethernet adapter. To configure the adapter, see [Configuring NVMe over TCP on ESXi](#).
- Software NVMe over TCP adapter. This software component must be enabled on your ESXi host and connected to an appropriate network adapter. For more information, see [Add Software NVMe over RDMA or NVMe over TCP Adapters](#).
- NVMe controller. You must add a controller after you configure the software NVMe over TCP adapter. See [Add Controllers for NVMe over Fabrics](#).

VMware NVMe over Fabrics Shared Storage Support

In the ESXi environment, the NVMe storage devices appear similar to SCSI storage devices, and can be used as shared storage. Follow these rules when using the NVMe-oF storage.

- Do not mix transport types to access the same namespace.
- Make sure that active paths are presented to the host. The namespaces cannot be registered until the active path is discovered.

Shared Storage Functionality	SCSI over Fabric Storage	NVMe over Fabric Storage
RDM	Supported	Not supported
Core dump	Supported	Supported
SCSI-2 reservations	Supported	Not supported
Clustered VMDK	Supported	Supported
Shared VMDK with multi-writer flag	Supported	Supported In vSphere 7.0 Update 1 and later. For more information, see the Knowledge Base article .
Virtual Volumes	Supported	Supported

Shared Storage Functionality	SCSI over Fabric Storage	NVMe over Fabric Storage
		In vSphere 8.0 and later. For more information, see NVMe and Virtual Volumes in vSphere .
Hardware acceleration with VAAI plug-ins	Supported	Not supported
Default MPP	NMP	HPP (NVMe-oF targets cannot be claimed by NMP)

Configuring Lossless Ethernet for NVMe over RDMA

NVMe over RDMA in ESXi requires lossless Ethernet network.

To establish lossless networks, you can select one of the available QoS settings.

Enable Global Pause Flow Control

In this network configuration, ensure global pause flow control is enabled on the network switch ports. Also, ensure that RDMA capable NICs in the host auto-negotiate to the correct flow control automatically.

To check flow control, run the following command:

```
#esxcli network nic get -n vmnicX
  Pause RX: true
  Pause TX: true
```

If the above command options are not set to true, run the following command.

```
#esxcli network nic pauseParams set -r true -t true -n vmnicX
```

Enable Priority Flow Control

For RoCE traffic to be lossless, you must configure the PFC priority value to 3 in the physical switch and hosts. You can configure the PFC in the ESXi host in two ways:

- Automatic Configuration. Apply DCB PFC configuration automatically on the host RNIC, if the RNIC driver supports DCB and DCBx.

You can verify the current DCB settings by running the following command:

```
#esxcli network nic dcb status get -n vmnicX
```

- Manual configuration. In some cases, the RNIC drivers provide a method to manually configure the DCB PFC using driver specific parameters. To use this method, see vendor specific driver documentation. For example, in Mellanox ConnectX-4/5 driver, you can set the PFC priority value to three by running the following command and then rebooting the host.

```
#esxcli system module parameters set -m nmlx5_core -p "pfctx=0x08 pfcrx=0x08"
```

Enable DSCP based PFC

DSCP based PFC is another way to configure lossless network. In physical switches and hosts, you must set the DSCP value to 26. To use this option, see vendor specific driver documentation. For example, in Mellanox ConnectX-4/5 driver, you can set the DSCP tag value to 26 by running the following commands.

1. Enable PFC and DSCP trust mode.

```
#esxcli system module parameters set -m nmlx5_core -p "pfctx=0x08 pfcrx=0x08 trust_state=2"
```

2. Set DSCP value to 26.

```
#esxcli system module parameters set -m nmlx5_rdma -p "dscp_force=26"
```

3. Verify parameters to check to confirm if settings are correct and set.

```
esxcli system module parameters list -m nmlx5_core | grep 'trust_state\|pfcrx\|pfctx'
```

4. Reboot the host.

Configuring NVMe over RDMA (RoCE v2) on ESXi

The adapter configuration process on the ESXi host involves setting up VMkernel binding for an RDMA network adapter, and then adding a software NVMe over RDMA adapter. You can then add an NVMe controller.

The entire configuration process includes these actions.

Action	Description
View RDMA Network Adapters	On your ESXi host, install a network adapter that supports RDMA (RoCE v2). For example, Mellanox Technologies MT27700 Family ConnectX-4. After you install the network adapter, use the vSphere Client to review the RDMA adapter and a physical network adapter.
Configure VMkernel Binding for the RDMA Adapter	Port binding for NVMe over RDMA involves creating a switch and connecting the physical network adapter and the VMkernel adapter to the switch. Through this connection, the RDMA adapter becomes bound to the VMkernel adapter. In the configuration, you can use a vSphere standard switch or a vSphere distributed switch.
Add the Software NVMe over RDMA Adapter	Use the vSphere Client to activate the software storage adapters for NVMe over RDMA.
Add Controllers for NVMe over Fabrics	Use the vSphere Client to add an NVMe controller. After you add the controller, the NVMe namespaces associated with the controller become available to your ESXi host. The NVMe storage devices that represent the namespaces in the ESXi environment appear on the storage devices list.

The following video walks you through the steps of configuring NVMe over RDMA adapters.

View RDMA Network Adapters

After you install a network adapter that supports RDMA (RoCE v2) on your ESXi host, use the vSphere Client to review the RDMA adapter and a physical network adapter.


1. On your ESXi host, install an adapter that supports RDMA (RoCE v2), for example, Mellanox Technologies MT27700 Family ConnectX-4.

The host discovers the adapter and the vSphere Client displays its two components, an RDMA adapter and a physical network adapter.

2. In the vSphere Client, verify that the RDMA adapter is discovered by your host.

- a) Navigate to the host.
- b) Click the **Configure** tab.
- c) Under **Networking**, click **RDMA adapters**.

In this example, the RDMA adapter appears on the list as `vmrdma0`. The **Paired Uplink** column displays the network component as the `vmnic1` physical network adapter.




- d) To verify the description of the adapter, select the RDMA adapter from the list, and click the **Properties** tab.

Configure VMkernel Binding for the RDMA Adapter

Port binding for NVMe over RDMA involves creating a switch and connecting the physical network adapter and the VMkernel adapter to the switch. Through this connection, the RDMA adapter becomes bound to the VMkernel adapter. In the configuration, you can use a vSphere standard switch or a vSphere distributed switch.


The following diagram displays the port binding for the NVMe over RDMA adapter.



For more information about creating switches, see *Create a vSphere Standard Switch* or *Create a vSphere Distributed Switch* in the *vSphere Networking* documentation.

Example of Network Topology with NVMe over RDMA

In this example, two vSphere standard switches and two uplinks (RDMA capable NICs) provide high availability. They connect to two controller pairs in two subnets.

HA with Multiple vSwitches and Multiple Uplinks (RNICs)

Configure VMkernel Binding with a vSphere Standard Switch

You can configure VMkernel port binding for the RDMA adapter using a vSphere standard switch and one uplink per switch. Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You use 1:1 mapping between each virtual and physical network adapter.

1. Create a vSphere standard switch with a VMkernel adapter and the network component.

- a) In the vSphere Client, select your host and click the **Networks** tab.
- b) Click **Actions > Add Networking**.
- c) Select **VMkernel Network Adapter** and click **NEXT**.
- d) Select **New standard switch** and click **NEXT**.
- e) Under **Assigned adapters**, click **+**.

The list of available physical adapters is displayed.

- f) Select the required physical adapter `vmnic`, and click **OK**.


Note: Ensure to select the physical network adapter that corresponds to the RDMA adapter. To see the association between the RDMA adapter `vmrdma`, and the physical network adapter `vmnic`, see [View RDMA Network Adapters](#).

- g) Under **VMkernel port settings**, enter the required values.
If you are using VLAN for the storage path, enter the VLAN ID.
- h) In the **IP settings** list, enter the VMkernel IPv4 settings.
- i) Under Available services, select **NVMe over RDMA**.


2. Verify that your switch is correctly configured.

- a) On the **Configure** tab, select **Virtual switches** under **Networking**.
- b) Expand the switch and verify its configuration.

The illustration shows that the physical network adapter and the VMkernel adapter are connected to the vSphere standard switch. Through this connection, the RDMA adapter is bound to the VMkernel adapter.



3. Verify the configuration of the VMkernel binding for the RDMA adapter.
 - a) Under **Networking** list, click **RDMA adapters**, and select the RDMA adapter from the list.
 - b) Click the **VMkernel adapters binding** tab and verify that the associated VMkernel adapter appears on the page.
 In this example, the `vmrdma0` RDMA adapter is paired to the `vmnic1` network adapter and is connected to the `vmk1` VMkernel adapter.



Configure VMkernel Binding with a vSphere Standard Switch and NIC Teaming

You can configure VMkernel port binding for the RDMA adapter using a vSphere standard switch with the NIC teaming configuration. You can use NIC teaming to achieve network redundancy. You can configure two or more network adapters (NICs) as a team for high availability and load balancing.

1. Create a vSphere standard switch with a VMkernel adapter and the network component with the NIC teaming configuration.
 - a) In the vSphere Client, select your host and click **Networks** tab.
 - b) Click **Actions > Add Networking**.
 - c) Select **VMkernel Network Adapter** and click **NEXT**.
 - d) Select **New standard switch** and click **NEXT**.
 - e) Under **Assigned adapters**, click **+**.
 A list of available physical adapters is displayed.
 - f) Select the required physical adapter `vmnic`, and add it under **Active adapters**.
 - g) Select another physical adapter `vmnic`, and add it under **Unused adapters**.
 - h) Under **VMkernel port settings**, enter the required values.
 If you are using VLAN for the storage path, enter the VLAN ID.
 - i) In the **IP settings** list, specify VMkernel IPv4 settings.
 - j) Under Available services, select **NVMe over RDMA**.
- Repeat step 1 to configure an existing standard switch.

2. Configure your switch for NIC teaming configuration.
 - a) Click the **Configure** tab, and select **Virtual switches** under **Networking**.
 - b) Select the appropriate VMkernel adapter.
 - c) From the right-click menu, click **Edit Settings**.
 - d) Select **Teaming and Failover**.
 - e) Under **Active adapters**, move the required physical adapter `vmnic`.
 - f) Under **Standby adapters > Failover order**, move the other physical adapters.
 - g) Set appropriate load balancing and other properties.
 - h) Repeat the steps to configure additional VMkernel adapters.
3. Repeat steps 1 and 2 to add and configure additional set of teamed `rnic`s. To verify if the adapter is configured, click the **Configure** tab and select **VMkernel adapters**

Configure VMkernel Binding with a vSphere Distributed Switch

You can configure VMkernel port binding for the RDMA adapter using a vSphere distributed switch and one uplink per switch. Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You use 1:1 mapping between each virtual and physical network adapter.

1. Create a vSphere distributed switch with a VMkernel adapter and the network component.
 - a) In the vSphere Client, select **Datacenter**, and click the **Networks** tab.
 - b) Click **Actions**, and select **Distributed Switch > New Distributed Switch**.
 - c) Select a name for the switch.
Ensure that the location of the data center is present within your host, and click **Next**.
 - d) Select a compatible ESXi version, and click **Next**.
 - e) Enter the required number of uplinks, and click **Finish**.
2. Add one or more hosts to your distributed virtual switch.
 - a) In the vSphere Client, select **Datacenter**, and click **Distributed Switches**.
A list of available DSwitches appear.
 - b) Right-click the DSwitch, and select **Add and Manage Hosts** from the menu.
 - c) Select **Add hosts**, and click **Next**.
 - d) Select your host, and click **Next**.
 - e) Select **Assign uplink**.
 - f) Enter the relevant uplink to assign the `vmnic`.
 - g) Assign a VMkernel adapter, and click **Next**.
 - h) In the vSphere Client, select the DSwitch, and click the **Ports** tab.
You can view the uplinks created for your switch here.
3. Create distributed port groups for the NVMe over RDMA storage path.
 - a) In the vSphere Client, select the required DSwitch.
 - b) Click **Actions** and select **Distributed Port Group > New Distributed Port Group**.
 - c) Under **Configure Settings**, enter the general properties of the port group.
If you have configured a specific VLAN, add it in the VLAN ID.

Note: Network connectivity issues might occur if you do not configure VLAN properly.

4. Configure the VMkernel adapters.
 - a) In the vSphere Client, expand the **DSwitch** list, and select the distributed port group.
 - b) Click **Actions > Add VMkernel Adapters**.
 - c) In the **Select Member Hosts** dialog box, select your host and click **OK**.
 - d) In the **Configure VMkernel Adapter** dialog box, ensure that the MTU matches to the Switch MTU.
 - e) Under **Available services**, select **NVMe over RDMA** for appropriate tagging.
 - f) Click **Finish**.
 - g) Repeat step b and step c to add multiple RDMA capable NICs.
5. Set NIC teaming policies for the distributed port groups.
 - a) In the **Distributed Port Group**, click **Actions > Edit Settings**.
 - b) Click **Teaming and Failover**, and verify the active uplinks.
 - c) Assign one uplink as **Active** for the port group, and the other uplink as **Unused**.

Repeat step c for each of the port groups created.

After you complete the configuration, click **Configure**, and verify whether the physical adapter tab on your host lists the DVSwitch for the NICs selected.

Add Software NVMe over RDMA or NVMe over TCP Adapters

ESXi supports NVMe over RDMA and NVMe over TCP software adapters. Use the vSphere Client to add the software storage adapters for NVMe over RDMA or NVMe over TCP.

- On your ESXi host, install an adapter that supports the following types of storage.
 - NVMe over RDMA adapter. For example, Mellanox Technologies MT27700 Family ConnectX-4.
 - NVMe over TCP adapter. For example, i40en.
 - Configure the VMkernel binding for your adapters.
 - For NVMe over RDMA, see [Configure VMkernel Binding for the RDMA Adapter](#).
 - For NVMe over TCP, see [Configure VMkernel Binding for the NVMe over TCP Adapter](#).
1. In the vSphere Client, navigate to the ESXi host.
 2. Click the **Configure** tab.
 3. Under **Storage**, click **Storage Adapters**, and click the **Add Software Adapter** icon.
 4. Select the adapter type as required.
 - **NVMe over RDMA adapter**
 - **NVMe over TCP adapter**
 5. Depending on your selection in Step 4, select an appropriate RDMA adapter or TCP network adapter (`vmnic`) from the drop-down menu.

Note: If you get an error message that prevents you from creating the software adapter, make sure that the VMkernel binding for the adapter is configured correctly. For more information, see [Configure VMkernel Binding for the RDMA Adapter](#) and [Configure VMkernel Binding for the NVMe over TCP Adapter](#).

The software NVMe over RDMA and NVMe over TCP adapters appear in the list as `vmhba` storage adapters. You can remove the adapters if you need to free the underlying RDMA and TCP network adapter for other purposes. See [Remove Software NVMe Adapters from the ESXi Host](#).

Add Controllers for NVMe over Fabrics

Use the vSphere Client to add an NVMe controller. After you add the controller, the NVMe namespaces associated with the controller become available to your ESXi host. The NVMe storage devices that represent the namespaces in the ESXi environment appear on the storage devices list.

- Make sure that your ESXi host has appropriate adapters for your type of storage. See [Requirements for VMware NVMe Storage](#).
- If you use NVMe over RDMA (RoCE v2) storage, perform these tasks:
 - [Configure VMkernel Binding for the RDMA Adapter](#)
 - [Add the Software NVMe over RDMA Adapter](#)
- If you use NVMe over TCP storage, perform these tasks:
 - [Configure VMkernel Binding for the NVMe over TCP Adapter](#)
 - [Add the Software NVMe over TCP Adapter](#)

Note: With NVMe over Fibre Channel, after you install the required adapter, it automatically connects to all targets that are reachable at the moment. You can later reconfigure the adapter and disconnect its controllers or connect other controllers that were not available during the host boot.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
4. Click the **Controllers** tab, and click **Add Controller**.
5. On the **Add controller** dialog box, select one of the following discovery methods.

Option	Description
Automatically	The automatic option indicates that your host can discover controllers automatically and accept connection to any available controller. 1. Specify the following parameters for to discover controllers. – For NVMe over RDMA (RoCE v2), the IP address and transport port number. – For NVMe over TCP, the IP address, transport port number, and the digest parameter. 2. Click Discover Controllers. 3. From the list of controllers, select the controller to use.
Manually	With the manual method, you manually enter controller details. The host requests a connection to a specific controller using the parameters you specify: • Subsystem NQN • Target port identification. – For NVMe over RDMA (RoCE v2), the IP address and transport port number (optional). – For NVMe over TCP, the IP address, transport port number (optional), and the digest parameter (optional). – For NVMe over Fibre Channel, the WorldWideNodeName and WorldWidePortName. • Admin queue size. An optional parameter that specifies the size of the admin queue of the controller. A default value is 16. • Keepalive timeout. An optional parameter to specify in seconds the keep alive timeout between the adapter and the controller. A default timeout value is 60 seconds. Note: IO Queue Size and IO Queue Number are optional parameters that can be set only through <code>esxcli</code>.

The controller appears on the list of controllers. Your host can now discover the NVMe namespaces that are associated with the controller. The NVMe storage devices that represent the namespaces in the ESXi environment appear on the storage devices list in the vSphere Client.

Configuring NVMe over TCP on ESXi

The adapter configuration process on the ESXi host involves setting up VMkernel binding for a TCP network adapter, and then adding a software adapter for NVMe over TCP. After that, you can add an NVMe controller.


The entire configuration process includes these actions.

Action	Description
On your ESXi host, install an adapter that supports NVMe over TCP technology	For example, i40en.
Configure VMkernel Binding for the NVMe over TCP Adapter	VMkernel binding for NVMe over TCP involves creating a virtual switch and connecting the physical network adapter and the VMkernel adapter to the virtual switch. Through this connection, the TCP adapter becomes bound to the VMkernel adapter. In the configuration, you can use a vSphere standard switch or a vSphere distributed switch.
Add the Software NVMe over TCP Adapter	Use the vSphere Client to enable the software storage adapters for NVMe over TCP.
Add Controllers for NVMe over Fabrics	Use the vSphere Client to add an NVMe controller. After you add the controller, the NVMe namespaces associated with the controller become available to your ESXi host. The NVMe storage devices that represent the namespaces in the ESXi environment appear on the storage devices list.

Configure VMkernel Binding for the NVMe over TCP Adapter

Port binding for NVMe over TCP involves creating a virtual switch and connecting the physical network adapter and the VMkernel adapter to the virtual switch. Through this connection, the TCP adapter becomes bound to the VMkernel adapter. In the configuration, you can use a vSphere standard switch or a vSphere distributed switch.


The following diagram displays the port binding for the NVMe over TCP adapter.



For more information about creating switches, see *Create a vSphere Standard Switch* or *Create a vSphere Distributed Switch* in the *vSphere Networking* documentation.

Example of Network Topology with NVMe over TCP

In this example, two vSphere standard switches and two network adapters (vmnic) on the host provide high availability. They connect to two external switches.



Configure VMkernel Binding for the TCP Adapter with a vSphere Standard Switch

You can configure VMkernel binding for the TCP adapter using a vSphere standard switch and one uplink per switch. Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You use 1:1 mapping between each virtual and physical network adapter.

1. Create a vSphere standard switch with a VMkernel adapter and the network component.

- In the vSphere Client, select your host and click the **Networks** tab.
- Click **Actions > Add Networking**.
- Select **VMkernel Network Adapter** and click **NEXT**.
- Select **New standard switch** and click **NEXT**.
- Under **Assigned adapters**, click **+**.

The list of available physical adapters is displayed.

- Select the required physical adapter `vmnic`, and click **OK**.

Note: Make sure to select the physical network adapter that corresponds to the TCP/IP adapter.

- Under **VMkernel port settings**, enter the required values.

If you are using VLAN for the storage path, enter the VLAN ID.

- In the **IP settings** list, enter the VMkernel IPv4 settings.

- Under **Available services**, select **NVMe over TCP** for appropriate tagging.

2. Verify that your switch is correctly configured.

- On the **Configure** tab, select **Virtual switches** under **Networking**.
- Expand the switch and verify its configuration.

The illustration shows that the physical network adapter and the VMkernel adapter are connected to the vSphere standard switch. Through this connection, the TCP adapter is bound to the VMkernel adapter.



3. Set NIC teaming policies for vSphere standard switch.

Note: The NVMe over TCP adapter does not support such NIC teaming features as failover and load balancing. Instead, it relies on Storage Multipathing for these functionalities. However, if you must configure NIC teaming for other network workloads on the uplink serving the NVMe over TCP adapter, follow these steps.

- a) Click the **Configure** tab, and select **Virtual switches** under **Networking**.
- b) Select the appropriate VMkernel adapter.
- c) From the right-click menu, click **Edit Settings**.
- d) Select **Teaming and Failover**.
- e) Under **Active adapters**, move the required physical adapter `vmnic`.
- f) Under **Standby adapters > Failover order**, move the other physical adapters.
- g) Set appropriate load balancing and other properties.
- h) Repeat the steps to configure additional VMkernel adapters.

To verify if the adapter is configured, click the **Configure** tab and select **VMkernel adapters**.

Configure VMkernel Binding for the TCP Adapter with a vSphere Distributed Switch

You can configure VMkernel port binding for the TCP adapter using a vSphere distributed switch and one uplink per switch. Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You use 1:1 mapping between each virtual and physical network adapter.

1. Create a vSphere distributed switch with a VMkernel adapter and the network component.

- a) In the vSphere Client, select **Datacenter**, and click the **Networks** tab.
- b) Click **Actions**, and select **Distributed Switch > New Distributed Switch**.
- c) Select a name for the switch.
Ensure that the location of the data center is present within your host, and click **Next**.
- d) Select a compatible ESXi version, and click **Next**.
- e) Enter the required number of uplinks, and click **Finish**.

2. Add one or more hosts to your distributed virtual switch.

- a) In the vSphere Client, select **Datacenter**, and click **Distributed Switches**.
A list of available DSwitches appear.
- b) Right-click the DSwitch, and select **Add and Manage Hosts** from the menu.
- c) Select **Add hosts**, and click **Next**.
- d) Select your host, and click **Next**.
- e) Select **Assign uplink**.
- f) Enter the relevant uplink to assign the `vmnic`.
- g) Assign a VMkernel adapter, and click **Next**.
- h) In the vSphere Client, select the DSwitch, and click the **Ports** tab.

You can view the uplinks created for your switch here.

3. Create distributed port groups for the NVMe over TCP storage path.

- a) In the vSphere Client, select the required DSwitch.
- b) Click **Actions** and select **Distributed Port Group > New Distributed Port Group**.
- c) Under **Configure Settings**, enter the general properties of the port group.

If you have configured a specific VLAN, add it in the VLAN ID.

Note: Network connectivity issues might occur if you do not configure VLAN properly.

4. Configure the VMkernel adapters.
 - a) In the vSphere Client, expand the **DSwitch** list, and select the distributed port group.
 - b) Click **Actions > Add VMkernel Adapters**.
 - c) In the **Select Member Hosts** dialog box, select your host and click **OK**.
 - d) In the **Configure VMkernel Adapter** dialog box, ensure that the MTU matches to the Switch MTU.
 - e) Click **Finish**.
 - f) Repeat step b and step c to add multiple TCP capable NICs.

5. Set NIC teaming policies for the distributed port groups.

Note: The NVMe over TCP adapter does not support such NIC teaming features as failover and load balancing. Instead, it relies on Storage Multipathing for these functionalities. However, if you must configure NIC teaming for other network workloads on the uplink serving the NVMe over TCP adapter, follow these steps.

- a) In the **Distributed Port Group**, click **Actions > Edit Settings**.
- b) Click **Teaming and Failover**, and verify the active uplinks.
- c) Assign one uplink as **Active** for the port group, and the other uplink as **Unused**.

Repeat step c for each of the port groups created.

After you complete the configuration, click **Configure**, and verify whether the physical adapter tab on your host lists the DVSwitch for the NICs selected.

Add Software NVMe over RDMA or NVMe over TCP Adapters

ESXi supports NVMe over RDMA and NVMe over TCP software adapters. Use the vSphere Client to add the software storage adapters for NVMe over RDMA or NVMe over TCP.

- On your ESXi host, install an adapter that supports the following types of storage.
 - NVMe over RDMA adapter. For example, Mellanox Technologies MT27700 Family ConnectX-4.
 - NVMe over TCP adapter. For example, i40en.
 - Configure the VMkernel binding for your adapters.
 - For NVMe over RDMA, see [Configure VMkernel Binding for the RDMA Adapter](#).
 - For NVMe over TCP, see [Configure VMkernel Binding for the NVMe over TCP Adapter](#).
1. In the vSphere Client, navigate to the ESXi host.
 2. Click the **Configure** tab.
 3. Under **Storage**, click **Storage Adapters**, and click the **Add Software Adapter** icon.
 4. Select the adapter type as required.
 - **NVMe over RDMA adapter**
 - **NVMe over TCP adapter**
 5. Depending on your selection in Step 4, select an appropriate RDMA adapter or TCP network adapter (`vmnic`) from the drop-down menu.

Note: If you get an error message that prevents you from creating the software adapter, make sure that the VMkernel binding for the adapter is configured correctly. For more information, see [Configure VMkernel Binding for the RDMA Adapter](#) and [Configure VMkernel Binding for the NVMe over TCP Adapter](#).

The software NVMe over RDMA and NVMe over TCP adapters appear in the list as `vmhba` storage adapters. You can remove the adapters if you need to free the underlying RDMA and TCP network adapter for other purposes. See [Remove Software NVMe Adapters from the ESXi Host](#).

Add Controllers for NVMe over Fabrics

Use the vSphere Client to add an NVMe controller. After you add the controller, the NVMe namespaces associated with the controller become available to your ESXi host. The NVMe storage devices that represent the namespaces in the ESXi environment appear on the storage devices list.

- Make sure that your ESXi host has appropriate adapters for your type of storage. See [Requirements for VMware NVMe Storage](#).
- If you use NVMe over RDMA (RoCE v2) storage, perform these tasks:
 - [Configure VMkernel Binding for the RDMA Adapter](#)
 - [Add the Software NVMe over RDMA Adapter](#)
- If you use NVMe over TCP storage, perform these tasks:
 - [Configure VMkernel Binding for the NVMe over TCP Adapter](#)
 - [Add the Software NVMe over TCP Adapter](#)

Note: With NVMe over Fibre Channel, after you install the required adapter, it automatically connects to all targets that are reachable at the moment. You can later reconfigure the adapter and disconnect its controllers or connect other controllers that were not available during the host boot.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to configure.
4. Click the **Controllers** tab, and click **Add Controller**.
5. On the **Add controller** dialog box, select one of the following discovery methods.

Option	Description
Automatically	The automatic option indicates that your host can discover controllers automatically and accept connection to any available controller. 1. Specify the following parameters for to discover controllers. – For NVMe over RDMA (RoCE v2), the IP address and transport port number. – For NVMe over TCP, the IP address, transport port number, and the digest parameter. 2. Click Discover Controllers. 3. From the list of controllers, select the controller to use.
Manually	With the manual method, you manually enter controller details. The host requests a connection to a specific controller using the parameters you specify: • Subsystem NQN • Target port identification. – For NVMe over RDMA (RoCE v2), the IP address and transport port number (optional). – For NVMe over TCP, the IP address, transport port number (optional), and the digest parameter (optional). – For NVMe over Fibre Channel, the WorldWideNodeName and WorldWidePortName. • Admin queue size. An optional parameter that specifies the size of the admin queue of the controller. A default value is 16. • Keepalive timeout. An optional parameter to specify in seconds the keep alive timeout between the adapter and the controller. A default timeout value is 60 seconds. Note: IO Queue Size and IO Queue Number are optional parameters that can be set only through <code>esxcli</code>.

The controller appears on the list of controllers. Your host can now discover the NVMe namespaces that are associated with the controller. The NVMe storage devices that represent the namespaces in the ESXi environment appear on the storage devices list in the vSphere Client.

Remove Software NVMe Adapters from the ESXi Host

Use the vSphere Client to remove software NVMe over RDMA or NVMe over TCP adapters from the ESXi host. You can remove the adapter if you need to free the underlying RDMA adapter or ethernet adapter for other purposes.

You cannot remove the NVMe over PCIe and NVMe over Fibre Channel adapters.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters**, and select the adapter (vmhba#) to remove.
4. Remove the NVMe controller connected to the adapter.
 - a) Click the **Controllers** tab.
 - b) Select the controller and click **Remove**.
The NVMe controller is disconnected and disappears from the list.
5. Click the **Remove** icon (Remove the host's storage adapter) to remove the NVMe over RDMA adapter or the NVMe over TCP adapter.

Working with Datastores in vSphere Storage Environment

In the vSphere environment, datastores are logical containers, analogous to file systems, that hide specifics of physical storage and provide a uniform model for storing virtual machine files. You can also use datastores for storing ISO images, virtual machine templates, and floppy images. vCenter Server and ESXi support VMFS, NFS, vSAN, and Virtual Volumes datastores.

vCenter Server and ESXi support the following types of datastores.

Table 216: Types of Datastores

Datastore Type	Description
VMFS (version 5 and 6)	Datastores that you deploy on block storage devices use the vSphere Virtual Machine File System (VMFS) format. VMFS is a special high-performance file system format that is optimized for storing virtual machines. See vSphere VMFS Datastore Concepts and Operations .
NFS (version 3 and 4.1)	An NFS client built into ESXi uses the Network File System (NFS) protocol over TCP/IP to access a designated NFS volume. The volume is located on a NAS server. The ESXi host mounts the volume as an NFS datastore, and uses it for storage needs. ESXi supports versions 3 and 4.1 of the NFS protocol. See NFS Datastore Concepts and Operations in vSphere Environment .
vSAN	vSAN aggregates all local capacity devices available on the hosts into a single datastore shared by all hosts in the vSAN cluster. See the <i>Administering VMware vSAN</i> documentation.
vVol	Virtual Volumes datastore represents a storage container in vCenter Server and vSphere Client. See Working with VMware vSphere Virtual Volumes .

Depending on your storage type, some of the following tasks are available for the datastores.

- Create datastores. You can use the vSphere Client to create certain types of datastores.
- Perform administrative operations on the datastores. Several operations, such as renaming a datastore, are available for all types of datastores. Others apply to specific types of datastores.
- Organize the datastores. For example, you can group them into folders according to business practices. After you group the datastores, you can assign the same permissions and alarms on the datastores in the group at one time.
- Add the datastores to datastore clusters. A datastore cluster is a collection of datastores with shared resources and a shared management interface. When you create the datastore cluster, you can use Storage DRS to manage storage resources. For information about datastore clusters, see the *vSphere Resource Management* documentation.

vSphere VMFS Datastore Concepts and Operations

To store virtual disks, ESXi uses datastores. The datastores are logical containers that hide specifics of physical storage from virtual machines and provide a uniform model for storing the virtual machine files. The datastores that you deploy on block storage devices use the native vSphere Virtual Machine File System (VMFS) format. It is a special high-performance file system format that is optimized for storing virtual machines.

Use the vSphere Client to set up the VMFS datastore in advance on the block-based storage device that your ESXi host discovers.

After creating a datastore, you can perform basic administration operations, such as renaming, unmounting, or removing the datastore. You can also perform more advanced tasks, for example, expand the datastore, manage duplicate datastores, or collect datastore diagnostic information.

The workflow that you use when managing the VMFS datastores generally includes the following actions.

Action	Description
1. Create a vSphere VMFS Datastore	Set up the VMFS datastore on the block-based storage device that your ESXi host discovers.
2. Perform Basic Administration Operations for Datastores	Perform basic administration operations, such as renaming, unmounting, or removing the datastore.
3. Perform Advanced VMFS Datastore Tasks	Perform more advanced tasks, for example, expand the datastore, manage duplicate datastores, or collect datastore diagnostic information.
4. Browse the Datastore	Use the datastore file browser in your vSphere environment to manage contents of your datastores. You can upload files or folders, download files from the datastore, or move or copy folders or files to a new location. You can also use the datastore browser to inflate the thin virtual disk.

Versions of VMFS Datastores

Several versions of the VMFS file system have been released since its introduction. Currently, ESXi supports VMFS5 and VMFS6.

For all supported VMFS version, ESXi offers complete read and write support. On the supported VMFS datastores, you can create and power on virtual machines.

Table 217: Host Access to VMFS Versions

VMFS	ESXi
VMFS6	Read and write

VMFS	ESXi
VMFS5	Read and write

The following table compares major characteristics of VMFS5 and VMFS6. For additional information, see *Configuration Maximums*.

Table 218: Comparing VMFS5 and VMFS6

Features and Functionalities	VMFS5	VMFS6
Access for ESXi hosts version 6.5 and later	Yes	Yes
Access for ESXi hosts version 6.0 and earlier	Yes	No
Datastores per host	512	512
512n storage devices	Yes	Yes (default)
512e storage devices	Yes. Not supported on local 512e devices.	Yes (default)
4Kn storage devices	No	Yes
Automatic space reclamation	No	Yes
Manual space reclamation through the esxcli command. See Manually Reclaim Accumulated Storage Space on VMFS5 .	Yes	Yes
Space reclamation from guest OS	Limited	Yes
GPT storage device partitioning	Yes	Yes
MBR storage device partitioning	Yes For a VMFS5 datastore that has been previously upgraded from VMFS3.	No
Storage devices greater than 2 TB for each VMFS extent	Yes	Yes
Support for virtual machines with large capacity virtual disks, or disks greater than 2 TB	Yes	Yes
Support of small files of 1 KB	Yes	Yes
Default use of ATS-only locking mechanisms on storage devices that support ATS. See vSphere VMFS Locking Mechanisms .	Yes	Yes
Block size	Standard 1 MB	Standard 1 MB
Default snapshots	VMFSsparse for virtual disks smaller than 2 TB. SEsparse for virtual disks larger than 2 TB.	SEsparse
Virtual disk emulation type	512n	512n
vMotion	Yes	Yes
Storage vMotion across different datastore types	Yes	Yes
High Availability and Fault Tolerance	Yes	Yes
DRS and Storage DRS	Yes	Yes
RDM	Yes	Yes

When you work with VMFS datastores, consider the following:

- **Datastore Extents.** A spanned VMFS datastore must use only homogeneous storage devices, either 512n, 512e, or 4Kn. The spanned datastore cannot extend over devices of different formats.
- **Block Size.** The block size on a VMFS datastore defines the maximum file size and the amount of space a file occupies. VMFS5 and VMFS6 datastores support the block size of 1 MB.
- **Storage vMotion.** Storage vMotion supports migration across VMFS, vSAN, and Virtual Volumes datastores. vCenter Server performs compatibility checks to validate Storage vMotion across different types of datastores.
- **Storage DRS.** VMFS5 and VMFS6 can coexist in the same datastore cluster. However, all datastores in the cluster must use homogeneous storage devices. Do not mix devices of different formats within the same datastore cluster.
- **Device Partition Formats.** Any new VMFS5 or VMFS6 datastore uses GUID partition table (GPT) to format the storage device. The GPT format enables you to create datastores larger than 2 TB. If your VMFS5 datastore has been previously upgraded from VMFS3, it continues to use the master boot record (MBR) partition format, which is characteristic for VMFS3. Conversion to GPT happens only after you expand the datastore to a size larger than 2 TB.

VMFS Datastores as Repositories

ESXi can format SCSI-based storage devices as VMFS datastores. VMFS datastores primarily serve as repositories for virtual machines.

Note: Always have only one VMFS datastore for each LUN.


You can store multiple virtual machines on the same VMFS datastore. Each virtual machine, encapsulated in a set of files, occupies a separate single directory. For the operating system inside the virtual machine, VMFS preserves the internal file system semantics, which ensures correct application behavior and data integrity for applications running in virtual machines.

When you run multiple virtual machines, VMFS provides specific locking mechanisms for the virtual machine files. As a result, the virtual machines can operate safely in a SAN environment where multiple ESXi hosts share the same VMFS datastore.

In addition to the virtual machines, the VMFS datastores can store other files, such as the virtual machine templates and ISO images.

Sharing a VMFS Datastore Across Hosts

As a cluster file system, VMFS lets multiple ESXi hosts access the same VMFS datastore concurrently.

Figure 101: Sharing a VMFS Datastore Across Hosts

For information on the maximum number of hosts that can connect to a single VMFS datastore, see the *Configuration Maximums* document.

To ensure that multiple hosts do not access the same virtual machine at the same time, VMFS provides on-disk locking. For more details about the ATS-only and ATS+SCSI locking mechanisms that VMFS uses, see [vSphere VMFS Locking Mechanisms](#). You can also learn how to display information about locking mechanisms, how to change from ATS+SCSI to ATS-only, or how to downgrade to ATS+SCSI.

Sharing the VMFS volume across multiple hosts offers several advantages, for example, the following:

- You can use VMware Distributed Resource Scheduling (DRS) and VMware High Availability (HA). You can distribute virtual machines across different physical servers. That means you run a mix of virtual machines on each server, so that not all experience high demand in the same area at the same time. If a server fails, you can restart virtual machines on another physical server. If the failure occurs, the on-disk lock for each virtual machine is released. For more information about VMware DRS, see the *vSphere Resource Management* documentation. For information about VMware HA, see the *vSphere Availability* documentation.
- You can use vMotion to migrate running virtual machines from one physical server to another. For information about migrating virtual machines, see the *vCenter Server and Host Management* documentation.

To create a shared datastore, mount the datastore on those ESXi hosts that require the datastore access. See [Mount Datastores](#).

VMFS Metadata Updates

A VMFS datastore holds virtual machine files, directories, symbolic links, RDM descriptor files, and so on. The datastore also maintains a consistent view of all the mapping information for these objects. This mapping information is called metadata.

Metadata is updated each time you perform datastore or virtual machine management operations. Examples of operations requiring metadata updates include the following:

- Creating, growing, or locking a virtual machine file
- Changing attributes of a file
- Powering a virtual machine on or off
- Creating or deleting a VMFS datastore
- Expanding a VMFS datastore

- Creating a template
- Deploying a virtual machine from a template
- Migrating a virtual machine with vMotion

When metadata changes are made in a shared storage environment, VMFS uses special locking mechanisms to protect its data and prevent multiple hosts from concurrently writing to the metadata.

Snapshot Formats on VMFS

When you take a snapshot, the state of the virtual disk is preserved, which prevents the guest operating system from writing to it. A delta or child disk is created. The delta represents the difference between the current state of the VM disk and the state that existed when you took the previous snapshot. On the VMFS datastore, the delta disk is a sparse disk.

Sparse disks use the copy-on-write mechanism, in which the virtual disk contains no data, until the data is copied there by a write operation. This optimization saves storage space.

Depending on the type of your datastore, delta disks use different sparse formats.

Snapshot Formats	VMFS5	VMFS6
VMFSsparse	For virtual disks smaller than 2 TB.	N/A
SEsparse	For virtual disks larger than 2 TB.	For all disks.

VMFSsparse

VMFS5 uses the VMFSsparse format for virtual disks smaller than 2 TB.

VMFSsparse is implemented on top of VMFS. The VMFSsparse layer processes I/Os issued to a snapshot VM. Technically, VMFSsparse is a redo-log that starts empty, immediately after a VM snapshot is taken. The redo-log expands to the size of its base vmdk, when the entire vmdk is rewritten with new data after the VM snapshotting. This redo-log is a file in the VMFS datastore. Upon snapshot creation, the base vmdk attached to the VM is changed to the newly created sparse vmdk.

SEsparse

SEsparse is a default format for all delta disks on the VMFS6 datastores. On VMFS5, SEsparse is used for virtual disks of the size 2 TB and larger.

SEsparse is a format similar to VMFSsparse with some enhancements. This format is space efficient and supports the space reclamation technique. With space reclamation, blocks that the guest OS deletes are marked. The system sends commands to the SEsparse layer in the hypervisor to unmap those blocks. The unmapping helps to reclaim space allocated by SEsparse once the guest operating system has deleted that data. For more information about space reclamation, see [Storage Space Reclamation in vSphere](#).

Snapshot Migration

You can migrate VMs with snapshots across different datastores. The following considerations apply:

- If you migrate a VM with the VMFSsparse snapshot to VMFS6, the snapshot format changes to SEsparse.
- When a VM with a vmdk of the size smaller than 2 TB is migrated to VMFS5, the snapshot format changes to VMFSsparse.
- You cannot mix VMFSsparse redo-logs with SEsparse redo-logs in the same hierarchy.

Upgrading VMFS Datastores

You cannot upgrade a VMFS5 datastore to VMFS6. If you have a VMFS5 datastore in your environment, create a VMFS6 datastore and migrate virtual machines from the VMFS5 datastore to VMFS6.

Creating vSphere Datastores

You use the New Datastore wizard to create your datastores in the vSphere environment. Depending on the type of your storage and storage needs, you can create a VMFS, NFS, or Virtual Volumes datastore.

A vSAN datastore is automatically created when you enable vSAN. For information, see the *Administering VMware vSAN* documentation.

You can also use the New Datastore wizard to manage VMFS datastore copies. See [vSphere VMFS Datastore Copies and Datastore Resignaturing](#).

To learn more about the vSphere datastores, see the following:

- [vSphere VMFS Datastore Concepts and Operations](#)
- [NFS Datastore Concepts and Operations in vSphere Environment](#)
- [Working with VMware vSphere Virtual Volumes](#)

Create a vSphere VMFS Datastore

VMFS datastores serve as repositories for virtual machines. You can set up VMFS datastores on any SCSI or NVMe based storage devices that the host discovers, including Fibre Channel, iSCSI, and local storage devices.

- Install and configure any adapters that your storage requires.
 - For Fibre Channel, see [Setting Up ESXi Fibre Channel Environment](#).
 - For iSCSI, see [Configuring iSCSI and iSER Adapters and Storage with ESXi](#).
 - For NVMe, see [About VMware NVMe Storage](#).
- To discover newly added storage devices, perform a rescan. See [Rescan Operations for ESXi Storage](#).
- Verify that storage devices you are planning to use for your datastores are available. See [Viewing Storage Devices Available to ESXi Hosts](#).

1. In the vSphere Client object navigator, browse to a host, a cluster, or a data center.
2. From the right-click menu, select **Storage > New Datastore**.
3. Select VMFS as the datastore type.
4. Enter the datastore name and if necessary, select the placement location for the datastore.
The system enforces an 80 character limit for the datastore name.
5. Select the device to use for your datastore.

Important: The device you select must not have any values displayed in the Snapshot Volume column. If a value is present, the device contains a copy of an existing VMFS datastore. For information on managing datastore copies, see [vSphere VMFS Datastore Copies and Datastore Resignaturing](#).

6. Specify the datastore version.

Option	Description
VMFS6	Default format on all hosts that support VMFS6. The ESXi hosts of version 6.0 or earlier cannot recognize the VMFS6 datastore.
VMFS5	VMFS5 datastore supports access by the ESXi hosts of version 6.7 or earlier.

7. Define configuration details for the datastore.

Note: The required minimum size for a VMFS6 datastore is 2 GB.

a) Specify partition configuration.

Option	Description
Use all available partitions	Dedicates the entire disk to a single VMFS datastore. If you select this option, all file systems and data currently stored on this device are destroyed.
Use free space	Deploys a VMFS datastore in the remaining free space of the disk.

b) If the space allocated for the datastore is excessive for your purposes, adjust the capacity values in the Datastore Size field.

By default, the entire free space on the storage device is allocated.

c) For VMFS6, specify the block size and define space reclamation parameters. See [Space Reclamation on vSphere VMFS Datastores](#).

8. In the Ready to Complete page, review the datastore configuration information and click **Finish**.

The datastore on the SCSI or NVMe based storage device is created. It is available to all hosts that have access to the device.

After you create the VMFS datastore, you can perform the following tasks:

- Change the capacity of the datastore. See [Increase Capacity of vSphere VMFS Datastore](#).
- Edit space reclamation settings. See [Configure Fixed Reclamation for the Existing VMFS6 Datastore](#).
- Enable shared vmdk support. See [Configure Settings for Clustered Virtual Disks on the vSphere VMFS6 Datastore](#).

Create an NFS Datastore in vSphere Environment

You can use the **New Datastore** wizard to mount an NFS volume.

- Set up NFS storage environment.
- If you plan to use Kerberos authentication with the NFS 4.1 datastore, make sure to configure the ESXi hosts for Kerberos authentication.
- If you plan to isolate the NFS 3 traffic to a specific VMkernel adapter that you connect to the datastore, create and configure the VMkernel adapter on default TCP/IP stack. For more information, see [Configure VMkernel Binding for NFS 3 Datastores](#).

1. In the vSphere Client object navigator, browse to a host, a cluster, or a data center.

2. From the right-click menu, select **Storage > New Datastore**.

3. Select NFS as the datastore type and specify an NFS version.

- NFS 3
- NFS 4.1

Important: If multiple hosts access the same datastore, you must use the same protocol on all hosts.

4. Enter the datastore parameters.

Option	Description
Datastore name	The system enforces an 80 character limit for the datastore name.
Folder	The mount point folder name
Server	The server name or IP address. You can use IPv6 or IPv4 formats.

Option	Description
	With NFS 4.1, you can add multiple IP addresses or server names if the NFS server supports trunking. The ESXi host uses these values to achieve multipathing to the NFS server mount point.

5. Select **Mount NFS read only** if the volume is exported as read-only by the NFS server.
6. To connect the NFS 3 datastore to a specific VMkernel adapter, select **Bind to vmknics**. For more information, see [Configure VMkernel Binding for NFS 3 Datastores](#).

Note: Only NFS 3 datastores support connections to VMkernel adapters.

7. To use Kerberos security with NFS 4.1, enable Kerberos and select an appropriate Kerberos model.

Option	Description
Use Kerberos for authentication only (krb5)	Supports identity verification
Use Kerberos for authentication and data integrity (krb5i)	In addition to identity verification, provides data integrity services. These services help to protect the NFS traffic from tampering by checking data packets for any potential modifications.

If you do not enable Kerberos, the datastore uses the default AUTH_SYS security.

8. If you create a datastore at the data center or cluster level, select hosts that mount the datastore.
9. Review the configuration options and click **Finish**.

Create a Virtual Volumes Datastore in vSphere Environment

You use the **New Datastore** wizard to create a Virtual Volumes datastore.

For more information about the Virtual Volumes datastore, see [Virtual Volumes Datastores](#).

1. In the vSphere Client object navigator, browse to a host, a cluster, or a data center.
2. From the right-click menu, select **Storage > New Datastore**.
3. Select **vVol** as the datastore type.
4. Enter the datastore name and select a backing storage container from the list of storage containers.
Make sure to use the name that does not duplicate another datastore name in your data center environment.

If you mount the same Virtual Volumes datastore to several hosts, the name of the datastore must be consistent across all hosts.

5. Select the hosts that require access to the datastore.
6. Review the configuration options and click **Finish**.

New Datastore

- Type
- Name and container selection
- Hosts accessibility
- Ready to complete**

Ready to complete

Review your selections before finishing the wizard

▼ Name and container selection

Datastore name	Datastore
Datastore type	vVol
Storage container name:	coke2
Storage container UUID:	vvol:c0b1db- 
Storage array(s):	xVP SCSI Array 2
Storage provider(s):	xvp-vasa4, xvp-vasa3

▼ Hosts accessibility

Hosts	 10.161.16 
	 10.161.16 
	 10.161.18 
	 10.161.17 

After you create the Virtual Volumes datastore, you can perform the following operations:

- Rename the datastore, browse datastore files, unmount the datastore, and so on. See [Basic Administrative Operations for vSphere Datastores](#).
- Note:** You cannot add the Virtual Volumes datastore to a datastore cluster.
- Provision virtual machines on the Virtual Volumes datastore. See [Provision VMs on vSphere Virtual Volumes Datastores](#) and the *vSphere Virtual Machine Administration* documentation.
- Review and manage protocol endpoints. See [Review and Manage Protocol Endpoints in vSphere](#).
- Manage or upgrade storage providers. See [Manage Storage Providers for vSphere Virtual Volumes](#).

Basic Administrative Operations for vSphere Datastores

After creating datastores, you can perform several administrative operations on the datastores, such as renaming, unmounting, or removing the datastore. Certain operations, for example, renaming, are available for all types of datastores. Others, such as deleting the datastore, apply only to VMFS datastores.

Change Datastore Name

Use the vSphere Client to change the name of an existing datastore. You can rename the datastore that has virtual machines running on it without any negative impact.

Note: If the host is managed by vCenter Server, you cannot rename the datastore by directly accessing the host from the VMware Host Client. You must rename the datastore from vCenter Server.

1. In the vSphere Client, navigate to the datastore.
2. Right-click the datastore to rename, and select **Rename**.
3. Enter a new datastore name.

The system enforces an 80 character limit for the datastore name.

The new name appears on all hosts that have access to the datastore.

Unmount Datastores

When you unmount a datastore, it remains intact, but can no longer be seen from the hosts that you specify. The datastore continues to appear on other hosts, where it remains mounted.

When appropriate, before unmounting datastores, make sure that the following prerequisites are met:

- No virtual machines reside on the datastore.
- Storage DRS does not manage the datastore.
- Storage I/O Control is deactivated for this datastore.

Do not perform any configuration operations that might result in I/O to the datastore while the unmounting is in progress.

Note: Make sure that the datastore is not used by vSphere HA Heartbeating. vSphere HA Heartbeating does not prevent you from unmounting the datastore. However, if the datastore is used for heartbeating, unmounting it might cause the host to fail and restart any active virtual machine.

1. In the vSphere Client, navigate to the datastore.
2. Right-click the datastore and select **Unmount Datastore**.
3. If the datastore is shared, select the hosts from which to unmount the datastore.
4. Confirm that you want to unmount the datastore.

After you unmount a VMFS datastore from all hosts, the datastore is marked as inactive. If you unmount an NFS or a Virtual Volumes datastore from all hosts, the datastore disappears from the inventory. You can mount the unmounted VMFS datastore. To mount the NFS or Virtual Volumes datastore that has been removed from the inventory, use the New Datastore wizard.

If you unmounted the VMFS datastore as a part of a storage removal procedure, you can now detach the storage device that is backing the datastore. See [Detach Storage Devices](#).

Mount Datastores

You can mount a datastore you previously unmounted. You can also mount a datastore on additional hosts, so that it becomes a shared datastore.

A VMFS datastore that has been unmounted from all hosts remains in inventory, but is marked as inaccessible. See [Unmount Datastores](#).

You can use this task to mount the VMFS datastore to a specified host or multiple hosts.

If you have unmounted an NFS or a Virtual Volumes datastore from all hosts, the datastore disappears from the inventory. To mount the NFS or Virtual Volumes datastore that has been removed from the inventory, use the New Datastore wizard.

A datastore of any type that is unmounted from some hosts while being mounted on others, is shown as active in the inventory.

1. In the vSphere Client, navigate to the datastore.
2. Right-click the datastore to mount and select one of the following options:

- **Mount Datastore**
- **Mount Datastore on Additional Hosts**

You can also select the VMkernel binding option for the required datastore. For more information, see step g under [Configure VMkernel Binding for NFS 3 Datastores](#).

Whether you see one or another option depends on the type of datastore you use.

3. Select the hosts that should access the datastore and click **OK**.
4. To list all hosts that share the datastore, navigate to the datastore, and click the **Hosts** tab.

Remove VMFS Datastores

You can delete any type of VMFS datastore, including copies that you have mounted without resignaturing. When you delete a datastore, it is destroyed and disappears from all hosts that have access to the datastore.

- Remove or migrate all virtual machines from the datastore.
- Unmount the datastore from all hosts.
- Disable Storage DRS for the datastore. For more information, see [Activate and Deactivate Storage DRS](#).
- Disable Storage I/O Control for the datastore. For more information, see [Enable Storage I/O Control](#).
- Make sure that the datastore is not used for vSphere HA heartbeating. For more information, see [Configure Heartbeat Datastores](#).

Note: The delete operation for the datastore permanently deletes all files associated with virtual machines on the datastore. Although you can delete the datastore without unmounting, it is preferable that you unmount the datastore first.

1. In the vSphere Client, navigate to the datastore.
2. Right-click the datastore to remove, and select **Delete Datastore**.
3. Confirm that you want to remove the datastore.

Managing vSphere VMFS Datastores

After you create a VMFS datastore in your vSphere environment, you can perform such tasks as expanding the datastore, managing duplicate datastores, or collecting datastore diagnostic information.

For general information about VMFS datastores, see [vSphere VMFS Datastore Concepts and Operations](#).

To create a VMFS datastore and perform basic operations, see these topics:

- [Create a vSphere VMFS Datastore](#)

- [Basic Administrative Operations for vSphere Datastores](#)

vSphere VMFS Locking Mechanisms

In a shared storage environment, when multiple ESXi hosts access the same VMFS datastore, locking mechanisms are used. The locking mechanisms prevent multiple hosts from concurrently writing to the metadata and ensure that no data corruption occurs. VMFS supports ATS-only and ATS+SCSI locking. Learn how to display information about locking mechanisms, how to change from ATS+SCSI to ATS-only, or how to downgrade to ATS+SCSI.

To learn more about VMFS as a cluster file system, see [vSphere VMFS Datastore Concepts and Operations](#).

Depending on its configuration and the type of underlying storage, a VMFS datastore can use different types of locking mechanisms. It can exclusively use the atomic test and set locking mechanism (ATS-only), or use a combination of ATS and SCSI reservations (ATS+SCSI).

What is ATS-Only Mechanism?

For storage devices that support T10 standard-based VAAI specifications, VMFS provides ATS locking, also called hardware assisted locking. The ATS algorithm supports discrete locking per disk sector. All newly formatted VMFS5 and VMFS6 datastores use the ATS-only mechanism if the underlying storage supports it, and never use SCSI reservations.

When you create a multi-extent datastore where ATS is used, vCenter Server filters out non-ATS devices. This filtering allows you to use only those devices that support the ATS primitive.

In certain cases, you might need to change the default locking mechanism for a VMFS5 or VMFS6 datastore. For information, see [Change Locking Mechanism to ATS+SCSI](#).

Note: If you are running a VMware vSAN environment or have ATS-only VMFS volumes, do not deactivate ATS. Deactivating ATS might result in an outage as no locking mechanism is available. For more information, see a [VMware Knowledge Base article](#).

What is ATS+SCSI Mechanism?

A VMFS datastore that supports the ATS+SCSI mechanism is configured to use ATS and attempts to use it when possible. If ATS fails, the VMFS datastore reverts to SCSI reservations. In contrast with the ATS locking, the SCSI reservations lock an entire storage device while an operation that requires metadata protection is performed. After the operation completes, VMFS releases the reservation, and other operations can continue.

Datastores that use the ATS+SCSI mechanism include VMFS5 datastores that were upgraded from VMFS3. In addition, new VMFS5 or VMFS6 datastores on storage devices that do not support ATS use the ATS+SCSI mechanism.

If the VMFS datastore reverts to SCSI reservations, you might notice performance degradation caused by excessive SCSI reservations.

Display VMFS Locking Information

Use the `esxcli` command to obtain information about the locking mechanism that a VMFS datastore uses.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To display information related to VMFS locking mechanisms, run the following command:

```
esxcli storage vmfs lockmode list
```

The table lists items that the output of the command might include.

Table 219: VMFS Locking Information

Fields	Values	Descriptions
Locking Modes		Indicates the locking configuration of the datastore.
	ATS	The datastore is configured to use the ATS-only locking mode.
	ATS+SCSI	The datastore is configured to use the ATS mode. If ATS fails or is not supported, the datastore can revert to SCSI.
	ATS upgrade pending	The datastore is in the process of an online upgrade to the ATS-only mode.
	ATS downgrade pending	The datastore is in the process of an online downgrade to the ATS+SCSI mode.
ATS Compatible		Indicates whether the datastore can be or cannot be configured for the ATS-only mode.
ATS Upgrade Modes		Indicates the type of upgrade that the datastore supports.
	None	The datastore is not ATS-only compatible.
	Online	The datastore can be used during its upgrade to the ATS-only mode.
	Offline	The datastore cannot be used during its upgrade to the ATS-only mode.
ATS Incompatibility Reason		If the datastore is not compatible with ATS-only, the item indicates the reason for the incompatibility.

Upgrading Legacy VMFS Datastores to ATS-Only Locking Mechanism

If your VMFS datastore uses the ATS+SCSI locking mechanism, you can change to ATS-only locking.

Typically, VMFS5 datastores that were previously upgraded from VMFS3 continue using the ATS+SCSI locking mechanism. If the datastores are deployed on ATS-enabled hardware, they are eligible for an upgrade to ATS-only locking. Depending on your vSphere environment, you can use one of the following upgrade modes:

- The online upgrade to the ATS-only mechanism is available for most single-extent VMFS5 datastores. While you perform the online upgrade on one of the hosts, other hosts can continue using the datastore.
- The offline upgrade to ATS-only must be used for VMFS5 datastores that span multiple physical extents. Datastores composed of multiple extents are not eligible for the online upgrade. These datastores require that no hosts actively use the datastores at the time of the upgrade request.

Prepare for an Upgrade to ATS-Only Locking

You must perform several steps to prepare your environment for an online or offline upgrade to ATS-only locking.

1. Upgrade all hosts that access the VMFS5 datastore to the newest version of vSphere.
2. Determine whether the datastore is eligible for an upgrade of its current locking mechanism by running the `esxcli storage vmfs lockmode list` command.

The following sample output indicates that the datastore is eligible for an upgrade. It also shows the current locking mechanism and the upgrade mode available for the datastore.

```
Locking Mode   ATS Compatible   ATS Upgrade Modes
-----
ATS+SCSI      true            Online or Offline
```

3. Depending on the upgrade mode available for the datastore, perform one of the following actions:

Upgrade Mode	Action
Online	Verify that all hosts have consistent storage connectivity to the VMFS datastore.
Offline	Verify that no hosts are actively using the datastore.

Upgrade Locking Mechanism to the ATS-Only Type

If a VMFS datastore is ATS-only compatible, you can upgrade its locking mechanism from ATS+SCSI to ATS-only.

If you plan to complete the upgrade of the locking mechanism by putting the datastore into maintenance mode, deactivate Storage DRS. This prerequisite applies only to an online upgrade.

Most datastores that do not span multiple extents are eligible for an online upgrade. While you perform the online upgrade on one of the ESXi hosts, other hosts can continue using the datastore. The online upgrade completes only after all hosts have closed the datastore.

1. Perform an upgrade of the locking mechanism by running the following command:

```
esxcli storage vmfs lockmode set -a|--ats -l|--volume-label=VMFS_label -u|--volume-uuid=VMFS_UUID.
```

2. For an online upgrade, perform additional steps.
 - a) Close the datastore on all hosts that have access to the datastore, so that the hosts can recognize the change. You can use one of the following methods:
 - Unmount and mount the datastore.
 - Put the datastore into maintenance mode and exit maintenance mode.
 - b) Verify that the Locking Mode status for the datastore changed to ATS-only by running:


```
esxcli storage vmfs lockmode list
```
 - c) If the Locking Mode displays any other status, for example ATS UPGRADE PENDING, check which host has not yet processed the upgrade by running:


```
esxcli storage vmfs host list
```

Change Locking Mechanism to ATS+SCSI

When you create a VMFS5 datastore on a device that supports the atomic test and set (ATS) locking, the datastore uses the ATS-only locking mechanism. In certain circumstances, you might need to downgrade the ATS-only locking to ATS+SCSI.

You might need to switch to the ATS+SCSI locking mechanism when, for example, your storage device is downgraded. Or when firmware updates fail and the device no longer supports ATS.

The downgrade process is similar to the ATS-only upgrade. As with the upgrade, depending on your storage configuration, you can perform the downgrade in online or offline mode.

Note: If you are running a VMware vSAN environment or have ATS-only VMFS volumes, do not deactivate ATS. Deactivating ATS might result in an outage as no locking mechanism is available. For more information, see a [VMware Knowledge Base article](#).

1. Change the locking mechanism to ATS+SCSI by running the following command:

```
esxcli storage vmfs lockmode set -s|--scsi -l|--volume-label= VMFS label -u|--volume-uuid= VMFS UUID.
```

2. For an online mode, close the datastore on all hosts that have access to the datastore, so that the hosts can recognize the change.

vSphere VMFS Datastore Copies and Datastore Resignaturing

When a storage device contains a VMFS datastore copy, ESXi detects the copy. You can mount the datastore copy with the original datastore UUID or change the UUID. The process of changing the UUID is called datastore resignaturing.

Each VMFS datastore created on a storage device has a unique signature, also called UUID, that is stored in the file system superblock. When the storage device is replicated or its snapshot is taken on the array side, the resulting device copy is identical, byte-for-byte, with the original device. For example, if the original storage device contains a VMFS datastore with UUIDX, the copy appears to contain a datastore copy with the same UUIDX.

In addition to LUN snapshots and replications, certain device operations, such as LUN ID changes, might produce a copy of the original datastore.

After the ESXi host detects the VMFS datastore copy, you can mount the datastore copy with its original UUID or change the UUID using the datastore resignaturing.

Whether you select resignaturing or mounting without resignaturing depends on how the LUNs are masked in the storage environment. If your hosts can see both copies of the LUN, then resignaturing is the optimal method.

Keeping Existing Datastore Signature

If you do not need to resignature a VMFS datastore copy, you can mount it without changing its signature.

You can keep the signature if, for example, you maintain synchronized copies of virtual machines at a secondary site as part of a disaster recovery plan. In the event of a disaster at the primary site, you mount the datastore copy and power on the virtual machines at the secondary site.

Resignaturing a VMFS Datastore Copy

Use datastore resignaturing if you want to retain the data stored on the VMFS datastore copy.

When resignaturing a VMFS copy, ESXi assigns a new signature (UUID) to the copy, and mounts the copy as a datastore distinct from the original. All references to the original signature in virtual machine configuration files are updated.

When you perform datastore resignaturing, consider the following points:

- Datastore resignaturing is irreversible.
- After resignaturing, the storage device replica that contained the VMFS copy is no longer treated as a replica.
- A spanned datastore can be resignatured only if all its extents are online.
- The resignaturing process is fault tolerant. If the process is interrupted, you can resume it later.
- You can mount the new VMFS datastore without a risk of its UUID conflicting with UUIDs of any other datastore from the hierarchy of device snapshots.

Resignature a VMFS Datastore Copy

Use datastore resignaturing if you want to retain the data stored on the VMFS datastore copy.

- Unmount the datastore copy.
- Perform a storage rescan on your host.

When resignaturing a VMFS copy, ESXi assigns a new signature (UUID) to the copy, and mounts the copy as a datastore distinct from the original. All references to the original signature from virtual machine configuration files are updated.

When you perform datastore resignaturing, consider the following points:

- Datastore resignaturing is irreversible.
 - After resignaturing, the storage device replica that contained the VMFS copy is no longer treated as a replica.
 - A spanned datastore can be resignatured only if all its extents are online.
 - The resignaturing process is fault tolerant. If the process is interrupted, you can resume it later.
 - You can mount the new VMFS datastore without a risk of its UUID conflicting with UUIDs of any other datastore from the hierarchy of device snapshots.
1. In the vSphere Client object navigator, browse to a host, a cluster, or a data center.
 2. Click the **New Datastore** icon.
 3. Enter the datastore name and if necessary, select the placement location for the datastore.
 4. Select VMFS as the datastore type.
 5. From the list of storage devices, select the device that has a specific value displayed in the Snapshot Volume column. The value present in the Snapshot Volume column indicates that the device is a copy that contains a copy of an existing VMFS datastore.
 6. Under **Mount Options**, select **Assign a New Signature** and click **Next**.
 7. Review the datastore configuration information and click **Finish**.

Increase Capacity of vSphere VMFS Datastore

You can increase the capacity of a VMFS datastore on the ESXi host. Additional capacity might be required when you add virtual machines to the datastore, or when the virtual machines running on the datastore require more space.

You can increase the datastore capacity if the host storage meets one of the following conditions:

- The backing device for the existing datastore has enough free space.
- You added new storage devices to the host.

If a shared datastore has powered on virtual machines and becomes 100% full, you can increase the datastore capacity. You can perform this action only from the host where the powered on virtual machines are registered.

Depending on your storage configuration, you can use one of the following methods to increase the datastore capacity. You do not need to power off virtual machines when using either method of increasing the datastore capacity.

Expand an Existing Datastore

Increase the size of an expandable datastore. The datastore is considered expandable when the backing storage device has free space immediately after the datastore extent.

Add an Extent

Increase the capacity of an existing VMFS datastore by adding new storage devices to the datastore. The datastore can span over multiple storage devices, yet appear as a single volume. The spanned VMFS datastore can use any or all its extents at any time. It does not need to fill up a particular extent before using the next one.

Note: Datastores that support only the hardware assisted locking, also called the atomic test and set (ATS) mechanism, cannot span

over non-ATS devices. For more information, see [vSphere VMFS Locking Mechanisms](#).

1. In the vSphere Client, navigate to the datastore.
2. Select **Increase Datastore Capacity** from the datastore right-click menu.
3. Select a device from the list of storage devices.

Your selection depends on whether an expandable storage device is available.

Option	Description
To expand an existing datastore extent	Select the device for which the Expandable column reads YES.
To add an extent	Select the device for which the Expandable column reads NO.

4. Review the **Partition Layout** to see the available configurations.
5. Select a configuration option from the bottom panel.

Depending on the current layout of the disk and on your previous selections, the menu items you see might vary.

Menu Item	Description
Use free space to expand the datastore	Expands an existing extent to a required capacity.
Use free space	Deploys an extent in the remaining free space of the disk. This menu item is available only when you are adding an extent.
Use all available partitions	Dedicates the entire disk to a single extent. This menu item is available only when you are adding an extent and when the disk you are formatting is not blank. The disk is reformatted, and the datastores and any data that it contains are erased.

6. Set the capacity for the extent.
The minimum extent size is 1.3 GB. By default, the entire free space on the storage device is available.
7. Click **Next**.
8. Review the proposed layout and the new configuration of your datastore, and click **Finish**.

Note: It is recommended to perform a cluster level rescan at vCenter Server after you perform any form of expansion done at the host level. If you do not perform a cluster level rescan from vCenter Server, it may not detect any change. This causes a disparity in the current space/usage for the datastore.

If the device where you increased the capacity is exposed to multiple hosts in multiple clusters, then the cluster level rescan must be triggered on all the clusters where the device is visible.

Configure Settings for Clustered Virtual Disks on the vSphere VMFS6 Datastore

If you plan to use a virtual disk in Windows Server Failover Clustering (WSFC) configurations, your VMFS6 datastore must support clustered virtual disks. Use the vSphere Client to configure the clustered disk support.

Follow these guidelines when you use a datastore for clustered virtual disks:

- The storage array supports the ATS, Write Exclusive – All Registrant (WEAR) SCSI-3 reservation type. For this type of configurations, ESXi supports only Fibre Channel arrays.
- Storage devices that back the datastore are claimed by the NMP. ESXi does not support third-party plug-ins (MPPs) in the clustered virtual disk configurations.
- The datastore is of VMFS6 type. You cannot expand or span the datastore.
- The VMFS6 datastore is deployed on the storage device with the attribute `Clustered VMDK Supported` set to Yes.

- Virtual disks you use for clustering are in the Thick Provision Eager Zeroed format.

For information on the use of clustered virtual disks in VM clusters, see the *Setup for Windows Server Failover Clustering* documentation.

- In the vSphere Client, navigate to the datastore.
- Click the **Configure** tab, and click **General**.
- Under **Datastore Capabilities**, click one of the following options next to the **Clustered VMDK** item.

Option	Description
Enable	To activate support for clustered virtual disks on the datastore. After you activate the support, you can place the clustered virtual disks on this VMFS datastore.
Disable	To deactivate the support. Before deactivating, make sure to power off all virtual machines with the clustered virtual disks.

- Confirm your configuration.

Manage Filters for LUNs and vSphere VMFS Datastores

When you perform VMFS datastore management operations, vCenter Server uses default storage protection filters. The filters help you to avoid storage corruption by retrieving only the storage devices that can be used for a particular operation. Unsuitable devices are not displayed for selection. You can turn off the filters to view all devices.

Before you change the device filters, consult with the VMware support team.

These filters are available by default.

Filter Name	Description
config.vpxd.filter.vmfsFilter (VMFS Filter)	Filters out storage devices, or LUNs, that are already used by a VMFS datastore on any host managed by vCenter Server. The LUNs do not show up as candidates to be formatted with another VMFS datastore or to be used as an RDM.
config.vpxd.filter.rdmFilter (RDM Filter)	Filters out LUNs that are already referenced by an RDM on any host managed by vCenter Server. The LUNs do not show up as candidates to be formatted with VMFS or to be used by a different RDM. For your virtual machines to access the same LUN, the virtual machines must share the same RDM mapping file. For information about this type of configuration, see the <i>vSphere Resource Management</i> documentation.
config.vpxd.filter.sameHostsAndTransportsFilter (Same Hosts and Transports Filter)	Filters out LUNs ineligible for use as VMFS datastore extents because of host or storage type incompatibility. Prevents you from adding the following LUNs as extents: - LUNs not exposed to all hosts that share the original VMFS datastore. - LUNs that use a storage type different from the one the original VMFS datastore uses. For example, you cannot add a Fibre Channel extent to a VMFS datastore on a local storage device.
config.vpxd.filter.hostRescanFilter (Host Rescan Filter)	Automatically rescans and updates VMFS datastores after you perform datastore management operations. The filter helps provide a consistent view of all VMFS datastores on all hosts managed by vCenter Server.

Filter Name	Description
	Note: If you present a new LUN to a host or a cluster, the hosts automatically perform a rescan no matter whether you have the Host Rescan Filter on or off.

1. Browse to the vCenter Server instance.
2. Click the **Configure** tab.
3. Under **Settings**, click **Advanced Settings**, and click **EDIT SETTINGS**.
4. Specify the filter to turn off.

In the **Name** and **Value** text boxes at the bottom of the page, enter appropriate information.

Name	Value
config.vpxd.filter.vmfsFilter	False
config.vpxd.filter.rdmFilter	False
config.vpxd.filter.sameHostsAndTransportsFilter	False
config.vpxd.filter.hostRescanFilter	False
	Note: If you turn off this filter, your hosts continue to perform a rescan each time you present a new LUN to a host or a cluster.

5. Click **ADD**, and click **SAVE** to save your changes.
You are not required to restart the vCenter Server system.

Set Up Dynamic Disk Mirroring for Virtual Machines on vSphere VMFS

Typically, in the ESXi environment, you cannot use LUN manager software on virtual machines to mirror virtual disks. However, if your Microsoft Windows virtual machines support dynamic disks, you can mirror virtual disks across two SAN LUNs. Mirroring helps you to protect the virtual machines from an unplanned storage device loss.

- Use a Windows virtual machine that supports dynamic disks.
- Required privilege: **Virtual machine > Configuration > Settings**

1. Create a virtual machine with two virtual disks.
Place the disks on different datastores.
2. Log in to your virtual machine and configure the disks as dynamic mirrored disks.
See Microsoft documentation.
3. After the disks synchronize, power off the virtual machine.
4. Change virtual machine settings to enable the dynamic disk mirroring.
 - a) Right-click the virtual machine and select **Edit Settings**.
 - b) Click the **VM Options** tab and expand the **Advanced** menu.
 - c) Click **Edit Configuration** next to Configuration Parameters.
 - d) Click **Add Configuration Params** and add the following parameters:
Replace # with the boot device SCSI bus controller ID number. By default, it is set to 0.

Name	Value
scsi#.returnNoConnectDuringAPD	True

Name	Value
scsi#.returnBusyOnNoConnectStatus	False

- e) For ESXi version 6.7 or later, include an additional parameter for each virtual disk participating in the software RAID-1 configuration.

The parameter prevents guest OS I/O failures when a storage device fails.

Replace # with the boot device SCSI bus controller ID number. By default, it is set to 0.

Name	Value
scsi#:1.passthruTransientErrors	True
scsi#:2.passthruTransientErrors	True

- f) Click **OK**.

Setting Up ESXi Core Dump on a VMFS Datastore

During a host failure, ESXi must be able to save diagnostic information to a preconfigured location for diagnostic and technical support purposes.

Typically, a partition to collect diagnostic information, also called a core dump, is created on a local storage device during ESXi installation. You can also configure an ESXi Dump Collector and keep core dumps on a network server for use during debugging. For information on setting up the ESXi Dump Collector, see [Configure an ESXi Dump Collector](#) in the *VMware ESXi Installation and Setup* documentation.

Another option is to use a file on a VMFS datastore to collect the diagnostic information.

Set Up a File as Core Dump Location

If the size of your available core dump partition is insufficient, you can configure ESXi to use a file on a VMFS datastore for diagnostic information.

- Make sure that the amount of available space in the datastore for a core dump file is twice as large as the recommended file size. For more information, see the Knowledge Base article [2058416](#).
- Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Note: VMFS datastores on software iSCSI do not support core dump files.

- Create a VMFS datastore core dump file by running the following command:

```
esxcli system coredump file add
```

The command takes the following options, but they are not required and can be omitted:

Option	Description
<code>--datastore -d datastore_UUID or datastore_name</code>	Specify the datastore for the dump file. If not provided, the system selects a datastore of sufficient size.
<code>--file -f file_name</code>	Specify the filename of the dump file. If not provided, the system creates a unique name for the file.
<code>--size -s file_size_MB</code>	Set the size in MB of the dump file. If not provided, the system creates a file of the size appropriate for the memory installed in the host.

2. Verify that the file has been created:

```
esxcli system coredump file list
```

You can see the output similar to the following:

```
Path                                     Active Configured Size
-----
/vmfs/volumes/52b021c3-.../vmkdump/test.dumpfile false  false    104857600
```

3. Activate the core dump file for the host:

```
esxcli system coredump file set
```

The command takes the following options:

Option	Description
--enable -e	Activate or deactivate the dump file. This option cannot be specified when unconfiguring the dump file.
--path -p	The path of the core dump file to use. The file must be pre-allocated.
--smart -s	This flag can be used only with --enable -e=true . It causes the file to be selected using the smart selection algorithm. For example, <code>esxcli system coredump file set --smart --enable true</code>
--unconfigure -u	Unconfigure the current VMFS dump file.

4. Verify that the core dump file is active and configured:

```
esxcli system coredump file list
```

The output similar to the following indicates that the core dump file is active and configured:

```
Path                                     Active Configured Size
-----
/vmfs/volumes/52b021c3-.../vmkdump/test.dumpfile True    True    104857600
```

For information about other commands you can use to manage the core dump files, see the *ESXCLI Reference* documentation.

Deactivate and Delete a Core Dump File

Deactivate a configured core dump file and, if needed, remove it from the VMFS datastore.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

You can temporarily deactivate the core dump file. If you do not plan to use the deactivated file, you can remove it from the VMFS datastore. To remove the file that has not been deactivated, you can use the `esxcli system coredump file remove` command with the `--force | -F` parameter.

1. List the core dump files:

```
esxcli system coredump file list
```

2. Deactivate the core dump file by running the following command:

```
esxcli system coredump file set --unconfigure | -u
```


3. Remove the file from the VMFS datastore:

```
esxcli system coredump file remove --file | -f file_name
```

The command takes the following options:

Option	Description
<code>--file -f</code>	Enter the name of the dump file to be removed. If you do not enter the name, the command removes the default configured core dump file.
<code>--force -F</code>	Deactivate and unconfigure the dump file being removed. This option is required if the file has not been previously deactivated and is active.

The core dump file is removed from the VMFS datastore.

Checking Metadata Consistency with vSphere On-disk Metadata Analyzer

Use the vSphere On-disk Metadata Analyzer (VOMA) to check metadata consistency and identify and fix incidents of metadata corruption on the VMFS datastores or logical volumes that back the VMFS datastores.

You can check metadata consistency when you experience problems with a VMFS datastore or a virtual flash resource. For example, perform a metadata check if one of the following occurs:

- You experience storage outages.
- After you rebuild RAID or perform a disk replacement.
- You see metadata errors in the `vmkernel.log` file similar to the following:

```
cpu1:268057)WARNING: HBX: 599: Volume 50fd60a3-3aae1ae2-3347-0017a4770402 ("<Datastore_name>") may be
damaged on disk. Corrupt heartbeat detected at offset 3305472: [HB state 0 offset 6052837899185946624 gen
15439450 stampUS 5 $
```

- You are unable to access files on a VMFS.
- You see corruption being reported for a datastore in events tabs of vCenter Server.

To check metadata consistency, run VOMA from the CLI of an ESXi host. VOMA can be used to check and fix minor inconsistency issues for a VMFS datastore or logical volumes that back the VMFS datastore.

VOMA can check and fix the following items.

Table 220: VOMA Functions

VOMA Functions	Description
Metadata check and fix	Examples of metadata check and fix include, but are not limited to, the following: - Validation of VMFS volume header for basic metadata consistency. - Checking consistency of VMFS resource files (system file). - Checking the pathname and connectivity of all files.
Directory validation	VOMA can detect and correct the following errors: - Directory hash block corruption. - Alloc map corruption. - Link blocks corruptions. - Directory entry block corruptions. Based on the nature of the corruption, VOMA can either fix only the corrupted entries or fully reconstruct the hash block, alloc map blocks and link blocks.
Lost and found files	During a filesystem check, VOMA can find files that are not referenced anywhere in the filesystem. These orphaned files are valid and complete, but do not have a name or directory entry on the system.

VOMA Functions	Description
	If VOMA encounters orphaned files during scanning, it creates a directory named <code>lost+found</code> at the root of the volume to store the orphaned files. The names of the files use the <code>Filesequence-number</code> format.

Command options that the VOMA tool takes include the following.

Table 221: VOMA Command Options

Command Option	Description
<code>-m --module</code>	The modules to run include the following:
	<code>vmfs</code> If you do not specify the name of the module, this option is used by default. You can check the VMFS file systems and the file systems that back virtual flash resources. If you specify this module, minimal checks are performed for LVM as well.
	<code>lvm</code> Check logical volumes that back the VMFS datastores.
	<code>ptbl</code> Check and validate VMFS partitions, such as MBR or GPT. If no partition exists, determine whether partitions should exist.
<code>-f --func</code>	Functions to be performed include the following:
	<code>query</code> List functions supported by module.
	<code>check</code> Check for errors.
	<code>fix</code> Check and fix errors for PTBL module.
	<code>advfix</code> Check and fix for VMFS and LVM modules.
	<code>dump</code> Collect metadata dump.
<code>-A --skipAffinityChk</code>	Skip affinity related check/fix for VMFS6.
<code>-d --device</code>	Indicate the device or disk to inspect. Make sure to provide the absolute path to the device partition backing the VMFS datastore. If the datastore spans multiple devices, provide the UUID of the head extent. For example, <code>voma -m vmfs -f check -d /vmfs/devices/disks/naa.xxxx:x</code> If you use the <code>-x --extractDump</code> command, enter multiple device paths, with a partition qualifier, separated with a comma. The number of device paths you enter equals the number of spanned devices.
<code>-b --blockSize</code>	Indicate the disk block size.
<code>-s --logfile</code>	Specify the path to log file to output the results.
<code>-x --extractDump</code>	Extract the collected dump using VOMA.
<code>-D --dumpfile</code>	Indicate the dump file to save the collected metadata dump.
<code>-v --version</code>	Display the version of VOMA.
<code>-h --help</code>	Display the help message for the VOMA command.
<code>-Z --file</code>	Indicate that you run VOMA on extracted device files.
<code>-p --patchfiledir</code>	Indicate the file to store the corrupted VMFS/LVM metadata.
<code>-r --restorefile</code>	Restore previous state using patch file generated by advfix mode.
<code>-l --live</code>	Perform dump operation on an active datastore.

Use the following as examples:

```
voma -m vmfs -f check -d /vmfs/devices/disks/naa.xxxx:x
voma -m vmfs -f dump -d /vmfs/devices/disks/naa.xxxx:x -D dumpfilename
voma -r patchfilename -d /vmfs/devices/disks/naa.xxxx:x
```

Use VOMA to Check Metadata Consistency

The task demonstrates how to use VOMA to check VMFS metadata consistency. VOMA can be used to check and fix minor inconsistency issues for a VMFS datastore or a virtual flash resource. Run VOMA from the CLI of an ESXi host.

Power off any virtual machines that are running or migrate them to a different datastore.

1. Obtain the name and partition number of the device that backs the VMFS datastore that you want to check.

```
#esxcli storage vmfs extent list
```

The Device Name and Partition columns in the output identify the device. For example:

```
Volume Name      .... Device Name      Partition
1TB_VMFS6        .... naa.xxxx          3
```

2. Check for VMFS errors.

Provide the absolute path to the device partition that backs the VMFS datastore, and provide a partition number with the device name. For example:

```
# voma -m vmfs -f check -d /vmfs/devices/disks/naa.xxxx:x
```

The output lists possible errors. For example, the following output indicates that the heartbeat address is invalid.

```
XXXXXXXXXXXXXXXXXXXXXXXXX
Phase 2: Checking VMFS heartbeat region
ON-DISK ERROR: Invalid HB address
Phase 3: Checking all file descriptors.
Phase 4: Checking pathname and connectivity.
Phase 5: Checking resource reference counts.

Total Errors Found:      1
```

Configuring Pointer Block Cache for vSphere VMFS Datastores

Pointer blocks, also called indirection blocks, are file system resources that contain addresses to VMFS file blocks. When you open a `vmdk` file on an ESXi host, pointer blocks related to that file are stored in the pointer block cache. The size of the pointer block cache is a configurable parameter.

The pointer block cache is a host-wide cache that is independent from VMFS. The cache is shared across all datastores that are accessed from the same ESXi host.

The size of the pointer block cache is controlled by `/VMFS3/MinAddressableSpaceTB` and `/VMFS3/MaxAddressableSpaceTB`. You can configure the minimum and maximum sizes on each ESXi host.

Parameter	Description
<code>/VMFS3/MinAddressableSpaceTB</code>	The minimum value is minimum amount of memory that the system guarantees to the pointer block cache. For example, 1 TB of open file space requires approximately 4 MB of memory. Default value is 10 TB.

Parameter	Description
/VMFS3/MaxAddressableSpaceTB	The parameter defines the maximum limit of pointer blocks that can be cached in memory. Default value is 32 TB. Maximum value is 128 TB. Typically, the default value of the /VMFS3/MaxAddressableSpaceTB parameter is adequate. However, as the size of the open vmdk files increases, the number of pointer blocks related to those files also increases. If the increase causes any performance degradation, you can adjust the parameter to its maximum value to provide more space for the pointer block cache. Base the maximum size of the pointer block cache on the working set, or the active pointer blocks required.
Pointer Block Eviction	The /VMFS3/MaxAddressableSpaceTB parameter also controls the growth of the pointer block cache. When the size of the pointer block cache approaches the configured maximum size, a pointer block eviction process starts. The mechanism leaves active pointer blocks, but removes non-active or less active blocks from the cache, so that space can be reused.

To change the values for the pointer block cache, use the **Advanced System Settings** dialog box of the vSphere Client or the `esxcli system settings advanced set -o` command.

You can use the `esxcli storage vmfs pbcache` command to obtain information about the size of the pointer block cache and other statistics. This information assists you in adjusting minimum and maximum sizes of the pointer block cache, so that you can get maximum performance.

Obtain Information for VMFS Pointer Block Cache

You can get information about VMFS pointer block cache use. This information helps you understand how much space the pointer block cache consumes. You can also determine whether you must adjust the minimum and maximum sizes of the pointer block cache.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To obtain or reset the pointer block cache statistics, use the following command:

```
esxcli storage vmfs pbcache
```

Option	Description
get	Get VMFS pointer block cache statistics.
reset	Reset the VMFS pointer block cache statistics.

Getting Statistics for Pointer Block Cache

```
#esxcli storage vmfs pbcache get
Cache Capacity Miss Ratio: 0 %
Cache Size: 0 MiB
Cache Size Max: 132 MiB
Cache Usage: 0 %
Cache Working Set: 0 TiB
Cache Working Set Max: 32 TiB
Vmfs Heap Overhead: 0 KiB
Vmfs Heap Size: 23 MiB
Vmfs Heap Size Max: 256 MiB
```

Change the Size of the Pointer Block Cache

You can adjust the minimum and maximum sizes of the pointer block cache.

CAUTION: Changing advanced options is considered unsupported. Typically, the default settings produce the optimum result. Change the advanced options only when you get specific instructions from VMware technical support or a knowledge base article.

1. Browse to the host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. In Advanced System Settings, select the appropriate item.

Option	Description
VMFS3.MinAddressableSpaceTB	Minimum size of all open files that VMFS cache guarantees to support.
VMFS3.MaxAddressableSpaceTB	Maximum size of all open files that VMFS cache supports before eviction starts.

5. Click the **Edit** button and change the value.
6. Click **OK**.

Use the `esxcli` Command to Change the Pointer Block Cache

You can also use the `esxcli system settings advanced set -o` to modify the size of the pointer block cache. The following example describes how to set the size to its maximum value of 128 TB.

1. To change the value of `/VMFS3/MaxAddressableSpaceTB` to 128 TB, enter the following command:

```
# esxcli system settings advanced set -i 128 -o /VMFS3/MaxAddressableSpaceTB
```

2. To confirm that the value is set correctly, enter this command:

```
# esxcli system settings advanced list -o /VMFS3/MaxAddressableSpaceTB
```

NFS Datastore Concepts and Operations in vSphere Environment

An NFS client built into ESXi uses the Network File System (NFS) protocol over TCP/IP to access a designated NFS volume that is located on a NAS server. The ESXi host can mount the volume and use it for its storage needs. vSphere supports versions 3 and 4.1 of the NFS protocol.

Typically, the NFS volume or directory is created by a storage administrator and is exported from the NFS server. You do not need to format the NFS volume with a local file system, such as VMFS. Instead, you mount the volume directly on the ESXi hosts and use it to store and boot virtual machines in the same way that you use the VMFS datastores.

In addition to storing virtual disks on NFS datastores, you can use NFS as a central repository for ISO images, virtual machine templates, and so on. If you use the datastore for the ISO images, you can connect the CD-ROM device of the virtual machine to an ISO file on the datastore. You then can install a guest operating system from the ISO file.

After you set up NFS storage environment, see the following topics to create and manage NFS datastores:

- [Create an NFS Datastore in vSphere Environment](#)
- [Basic Administrative Operations for vSphere Datastores](#)

NFS Protocols and ESXi

ESXi supports NFS protocols version 3 and 4.1. To support both versions, ESXi uses two different NFS clients.

Comparing Versions of NFS Clients

The following table lists capabilities that the NFS version 3 and 4.1 support.

Characteristics	NFS version 3	NFS version 4.1
Security mechanisms	AUTH_SYS	AUTH_SYS and Kerberos (krb5 and krb5i)
Encryption algorithms with Kerberos	N/A	AES256-CTS-HMAC-SHA1-96 and AES128-CTS-HMAC-SHA1-96
Multipathing	Not supported	Supported through the session trunking
Locking mechanisms	Propriety client-side locking	Server-side locking
Hardware acceleration	Supported	Supported
Thick virtual disks	Supported	Supported
IPv6	Supported	Supported for AUTH_SYS and Kerberos
ISO images presented as CD-ROMs to virtual machines	Supported	Supported
Virtual machine snapshots	Supported	Supported
Virtual machines with virtual disks greater than 2 TB	Supported	Supported

NFS Protocols and vSphere Solutions

The following table lists major vSphere solutions that NFS versions support.

vSphere Features	NFS version 3	NFS version 4.1
vMotion and Storage vMotion	Yes	Yes
High Availability (HA)	Yes	Yes
Fault Tolerance (FT)	Yes	Yes
Distributed Resource Scheduler (DRS)	Yes	Yes
Host Profiles	Yes	Yes
Storage DRS	Yes	Initial placement and load balancing based on space utilization.
Storage I/O Control	Yes	No
Site Recovery Manager	Yes	Site Recovery Manager does not support NFS 4.1 datastores for array-based replication and Virtual Volumes replication. You can use Site Recovery Manager with NFS v 4.1 datastores for vSphere Replication.
Virtual Volumes	Yes	Yes
vSphere Replication	Yes	Yes
VMware Aria Operations	Yes	Yes

NFS 4.1 and Fault Tolerance

Virtual machines on NFS 4.1 support the new Fault Tolerance mechanism that was introduced in vSphere 6.0. The mechanism can accommodate symmetric multiprocessor (SMP) virtual machines with up to four vCPUs.

NFS 4.1 VMs do not support the legacy Fault Tolerance mechanism.

NFS Upgrades

When you upgrade ESXi from a version earlier than 6.5, existing NFS 4.1 datastores automatically begin supporting functionalities that were not available in the previous ESXi release. These functionalities include Virtual Volumes, hardware acceleration, and so on.

ESXi does not support automatic datastore conversions from NFS version 3 to NFS 4.1.

If you want to upgrade your NFS 3 datastore, the following options are available:

- Create the NFS 4.1 datastore, and then use Storage vMotion to migrate virtual machines from the old datastore to the new one.
- Use conversion methods provided by your NFS storage server. For more information, contact your storage vendor.
- Unmount the NFS 3 datastore, and then mount as NFS 4.1 datastore.

CAUTION: If you use this option, make sure to unmount the datastore from all hosts that have access to the datastore. The datastore can never be mounted by using both protocols at the same time.

Guidelines and Requirements for NFS Storage with ESXi

When you use NFS storage with ESXi, follow specific guidelines related to NFS server configuration, networking, NFS datastores, and so on.

NFS Server Configuration

When you configure NFS servers to work with ESXi, follow recommendation of your storage vendor. In addition to these general recommendations, use specific guidelines that apply to NFS in vSphere environment.

The guidelines include the following items.

- Make sure that the NAS servers you use are listed in the *VMware HCL*. Use the correct version for the server firmware.
- Ensure that the NFS volume is exported using NFS over TCP.
- Make sure that the NAS server exports a particular share as either NFS 3 or NFS 4.1. The NAS server must not provide both protocol versions for the same share. The NAS server must enforce this policy because ESXi does not prevent mounting the same share through different NFS versions.
- NFS 3 and non-Kerberos (AUTH_SYS) NFS 4.1 do not support the delegate user functionality that enables access to NFS volumes using nonroot credentials. If you use NFS 3 or non-Kerberos NFS 4.1, ensure that each host has root access to the volume. Different storage vendors have different methods of enabling this functionality, but typically the NAS servers use the `no_root_squash` option. If the NAS server does not grant root access, you can still mount the NFS datastore on the host. However, you cannot create any virtual machines on the datastore.
- If the underlying NFS volume is read-only, make sure that the volume is exported as a read-only share by the NFS server. Or mount the volume as a read-only datastore on the ESXi host. Otherwise, the host considers the datastore to be read-write and might not open the files.

NFS Networking

An ESXi host uses TCP/IP network connection to access a remote NAS server. Certain guidelines and best practices exist for configuring the networking when you use NFS storage.

For more information, see the *vSphere Networking* documentation.

- For network connectivity, use a standard network adapter in your ESXi host.

- ESXi supports Layer 2 and Layer 3 Network switches. If you use Layer 3 switches, ESXi hosts and NFS storage arrays must be on different subnets and the network switch must handle the routing information.
- Configure a VMkernel port group for NFS storage. You can create the VMkernel port group for IP storage on an existing virtual switch (vSwitch) or on a new vSwitch. The vSwitch can be a vSphere Standard Switch (VSS) or a vSphere Distributed Switch (VDS).
- If you use multiple ports for NFS traffic, make sure that you correctly configure your virtual switches and physical switches.
- NFS 3 and NFS 4.1 support IPv6.
- You can configure NFS storage with multiple connections by using the `nconnect` option. For NFS 4.1, you can create multiple connections per session. For NFS 3, you can mount the datastore with multiple connections. You can set a maximum of 4 connections per NFS datastore by default. However, you can increase it up to 8 by using the advanced NFS option. Ensure that the total number of connections across all mounted NFS datastores does not exceed 256. See [Configure Multiple TCP Connections for NFS](#).
- You can isolate NFS traffic to specific VMkernel adapters. Without binding, if the VMkernel adapter that ESXi uses for NFS traffic fails, the network infrastructure redirects the traffic to an alternative route. As a result, the NFS traffic might unintentionally flow through a random VMkernel adapter. VMkernel port binding for the NFS datastore allows you to bind an NFS volume to a specific VMkernel adapter to connect to an NFS server. See [Configure VMkernel Binding for NFS Datastore](#).

NFS File Locking

File locking mechanisms are used to restrict access to data stored on a server to only one user or process at a time. The locking mechanisms of the two NFS versions are not compatible. NFS 3 uses proprietary locking and NFS 4.1 uses native protocol specified locking.

NFS 3 locking on ESXi does not use the Network Lock Manager (NLM) protocol. Instead, VMware provides its own locking protocol. NFS 3 locks are implemented by creating lock files on the NFS server. Lock files are named `.lck-file_id..`

NFS 4.1 uses share reservations as a locking mechanism.

Because NFS 3 and NFS 4.1 clients do not use the same locking protocol, you cannot use different NFS versions to mount the same datastore on multiple hosts. Accessing the same virtual disks from two incompatible clients might result in incorrect behavior and cause data corruption.

NFS Security

With NFS 3 and NFS 4.1, ESXi supports the AUTH_SYS security. In addition, for NFS 4.1, the Kerberos security mechanism is supported.

NFS 3 supports the AUTH_SYS security mechanism. With this mechanism, storage traffic is transmitted in an unencrypted format across the LAN. Because of this limited security, use NFS storage on trusted networks only and isolate the traffic on separate physical switches. You can also use a private VLAN.

NFS 4.1 supports the Kerberos authentication protocol to secure communications with the NFS server. Nonroot users can access files when Kerberos is used. For more information, see [Using Kerberos for NFS 4.1 with ESXi](#).

In addition to Kerberos, NFS 4.1 supports traditional non-Kerberos mounts with the AUTH_SYS security. In this case, use root access guidelines for NFS version 3.

Note: You cannot use two security mechanisms, AUTH_SYS and Kerberos, for the same NFS 4.1 datastore shared by multiple hosts.

NFS Multipathing

NFS 4.1 supports multipathing as per protocol specifications. For NFS 3 multipathing is not applicable.

NFS 3 uses one TCP connection for I/O. As a result, ESXi supports I/O on only one IP address or hostname for the NFS server, and does not support multiple paths. Depending on your network infrastructure and configuration, you can use the network stack to configure multiple connections to the storage targets. In this case, you must have multiple datastores, each datastore using separate network connections between the host and the storage.

NFS 4.1 provides multipathing for servers that support the session trunking. When the trunking is available, you can use multiple IP addresses to access a single NFS volume. Client ID trunking is not supported.

NFS and Hardware Acceleration

Virtual disks created on NFS datastores are thin-provisioned by default. To be able to create thick-provisioned virtual disks, you must use hardware acceleration that supports the Reserve Space operation.

NFS 3 and NFS 4.1 support hardware acceleration that allows your host to integrate with NAS devices and use several hardware operations that NAS storage provides. For more information, see [vSphere Hardware Acceleration on NAS Devices](#).

NFS Datastores

When you create an NFS datastore, make sure to follow specific guidelines.

The NFS datastore guidelines and best practices include the following items. To create an NFS datastores, see [Create an NFS Datastore in vSphere Environment](#).

- You cannot use different NFS versions to mount the same datastore on different hosts. NFS 3 and NFS 4.1 clients are not compatible and do not use the same locking protocol. As a result, accessing the same virtual disks from two incompatible clients might result in incorrect behavior and cause data corruption.
- NFS 3 and NFS 4.1 datastores can coexist on the same host.
- ESXi cannot automatically upgrade NFS version 3 to version 4.1, but you can use other conversion methods. For information, see [NFS Upgrades](#).
- When you mount the same NFS 3 volume on different hosts, make sure that the server and folder names are identical across the hosts. If the names do not match, the hosts see the same NFS version 3 volume as two different datastores. This error might result in a failure of such features as vMotion. An example of such discrepancy is entering `filer` as the server name on one host and `filer.domain.com` on the other. This guideline does not apply to NFS version 4.1.
- If you use non-ASCII characters to name datastores and virtual machines, make sure that the underlying NFS server offers internationalization support. If the server does not support international characters, use only ASCII characters, or unpredictable failures might occur.

Using Layer 3 Routed Connections to Access NFS Storage with ESXi

When you use Layer 3 (L3) routed connections to access NFS storage, consider certain requirements and restrictions.

Ensure that your environment meets the following requirements:

- Use Cisco's Hot Standby Router Protocol (HSRP) in IP Router. If you are using a non-Cisco router, use Virtual Router Redundancy Protocol (VRRP) instead.
- To prioritize NFS L3 traffic on networks with limited bandwidths, or on networks that experience congestion, use Quality of Service (QoS). See your router documentation for details.
- Follow Routed NFS L3 recommendations offered by storage vendor. Contact your storage vendor for details.

- Deactivate Network I/O Resource Management (NetIORM).
- If you are planning to use systems with top-of-rack switches or switch-dependent I/O device partitioning, contact your system vendor for compatibility and support.

In an L3 environment, the following restrictions apply:

- The environment does not support VMware Site Recovery Manager.
- The environment supports only the NFS protocol. Do not use other storage protocols such as FCoE over the same physical network.
- The NFS traffic in this environment does not support IPv6.
- The NFS traffic in this environment can be routed only over a LAN. Other environments such as WAN are not supported.

Firewall Configurations for NFS Storage ESXi

Learn about firewall rule set files, `nfsClient` and `nfs41client`, that ESXi creates when you mount NFS datastores version 3 or 4.1

For general information about firewall configurations, see [Configuring the ESXi Firewall](#) in the *vSphere Security* documentation.

NFS Client Firewall Behavior

The NFS Client firewall rule set behaves differently than other ESXi firewall rule sets. ESXi configures NFS Client settings when you mount or unmount an NFS datastore. The behavior differs for different versions of NFS.

When you add, mount, or unmount an NFS datastore, the resulting behavior depends on the version of NFS.

NFS v3 Firewall Behavior

When you add or mount an NFS v3 datastore, ESXi checks the state of the NFS Client (`nfsClient`) firewall rule set.

- If the `nfsClient` rule set is deactivated, ESXi activates the rule set and deactivates the Allow All IP Addresses policy by setting the `allowedAll` flag to `FALSE`. The IP address of the NFS server is added to the allowed list of outgoing IP addresses.
- If the `nfsClient` rule set is activated, the state of the rule set and the allowed IP address policy are not changed. The IP address of the NFS server is added to the allowed list of outgoing IP addresses.

Note: If you manually activate the `nfsClient` rule set or manually set the Allow All IP Addresses policy, either before or after you add an NFS v3 datastore to the system, your settings are overridden when the last NFS v3 datastore is unmounted. The `nfsClient` rule set is deactivated when all NFS v3 datastores are unmounted.

When you remove or unmount an NFS v3 datastore, ESXi performs one of the following actions.

- If none of the remaining NFS v3 datastores are mounted from the server of the datastore being unmounted, ESXi removes the server's IP address from the list of outgoing IP addresses.
- If no mounted NFS v3 datastores remain after the unmount operation, ESXi deactivates the `nfsClient` firewall rule set.

NFS v4.1 Firewall Behavior

When you mount the first NFS v4.1 datastore, ESXi activates the `nfs41client` rule set and sets its `allowedAll` flag to `TRUE`. This action opens port 2049 for all IP addresses. Unmounting an NFS v4.1 datastore does not affect the firewall state. That is, the first NFS v4.1 mount opens port 2049 and that port remains activated unless you close it explicitly.

Verify Firewall Ports for NFS Clients

To enable access to NFS storage, ESXi automatically opens firewall ports for the NFS clients when you mount an NFS datastore. For troubleshooting reasons, you might need to verify that the ports are open.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Firewall**, and click **Edit**.
4. Scroll down to an appropriate version of NFS to make sure that the port is open.

Using Kerberos for NFS 4.1 with ESXi

With NFS version 4.1, ESXi supports the Kerberos authentication mechanism. The ESXi implementation of Kerberos for NFS 4.1 provides two security models, krb5 and krb5i, that offer different levels of security.

The RPCSEC_GSS Kerberos mechanism is an authentication service. It allows an NFS 4.1 client installed on ESXi to prove its identity to an NFS server before mounting an NFS share. The Kerberos security uses cryptography to work across an insecure network connection.

The ESXi implementation of Kerberos for NFS 4.1 provides two security models, krb5 and krb5i, that offer different levels of security.

- Kerberos for authentication only (krb5) supports identity verification.
- Kerberos for authentication and data integrity (krb5i), in addition to identity verification, provides data integrity services. These services help to protect the NFS traffic from tampering by checking data packets for any potential modifications.

Kerberos supports cryptographic algorithms that prevent unauthorized users from gaining access to NFS traffic. The NFS 4.1 client on ESXi attempts to use either the AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 algorithm to access a share on the NAS server. Before using your NFS 4.1 datastores, make sure that AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 are enabled on the NAS server.

The following table compares Kerberos security levels that ESXi supports.

Type of Kerberos Security		ESXi Support
Kerberos for authentication only (krb5)	Integrity checksum for RPC header	Yes with AES
	Integrity checksum for RPC data	No
Kerberos for authentication and data integrity (krb5i)	Integrity checksum for RPC header	Yes with AES
	Integrity checksum for RPC data	Yes with AES

When you use Kerberos authentication, the following considerations apply:

- ESXi uses Kerberos with the Active Directory domain.
- As a vSphere administrator, you specify Active Directory credentials to provide access to NFS 4.1 Kerberos datastores for an NFS user. A single set of credentials is used to access all Kerberos datastores mounted on that host.
- When multiple ESXi hosts share the NFS 4.1 datastore, you must use the same Active Directory credentials for all hosts that access the shared datastore. To automate the assignment process, set the user in host profiles and apply the profile to all ESXi hosts.
- You cannot use two security mechanisms, AUTH_SYS and Kerberos, for the same NFS 4.1 datastore shared by multiple hosts.

Configure ESXi Hosts for Kerberos Authentication

If you use NFS 4.1 with Kerberos, you must perform several tasks to set up your hosts for Kerberos authentication.

- Make sure that Microsoft Active Directory (AD) and NFS servers are configured to use Kerberos.
- Enable AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 encryption modes on AD. The NFS 4.1 client does not support the DES-CBC-MD5 encryption mode.
- Make sure that the NFS server exports are configured to grant full access to the Kerberos user.

When multiple ESXi hosts share the NFS 4.1 datastore, you must use the same Active Directory credentials for all hosts that access the shared datastore. You can automate the assignment process by setting the user in host profiles and applying the profile to all ESXi hosts.

Configure DNS for NFS 4.1 with Kerberos

When you use NFS 4.1 with Kerberos, you must change the DNS settings on ESXi hosts. The settings must point to the DNS server that is configured to hand out DNS records for the Kerberos Key Distribution Center (KDC). For example, use the Active Directory server address if AD is used as a DNS server.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Networking**, click **TCP/IP configuration**.
4. Select **Default** and click the **Edit** icon.
5. Manually enter the DNS settings.

Option	Description
Domain	<i>AD Domain Name</i>
Preferred DNS server	<i>AD Server IP</i>
Search domains	<i>AD Domain Name</i>

Configure Network Time Protocol for NFS 4.1 with Kerberos

If you use NFS 4.1 with Kerberos, ESXi hosts, the NFS server, and the Active Domain server need to be time synchronized. Typically, in the setup the Active Domain server is used as the Network Time Protocol (NTP) server.

The following task describes how to synchronize the ESXi host with the NTP server.

The best practice is to use the Active Domain server as the NTP server.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, select **Time Configuration**.
4. Click **Edit** and set up the NTP server.
 - a) Select **Use Network Time Protocol (Enable NTP client)**.
 - b) To synchronize with the NTP server, enter its IP addresses.
 - c) Select **Start NTP Service**.
 - d) Set the NTP Service Startup Policy.
5. Click **OK**.

The host synchronizes with the NTP server.

Enable Kerberos Authentication in Active Directory

If you use NFS 4.1 storage with Kerberos, you must add each ESXi host to an Active Directory domain and enable Kerberos authentication. Kerberos integrates with Active Directory to enable single sign-on and provides an extra layer of security when used across an insecure network connection.

Set up an AD domain and a domain administrator account with the rights to add hosts to the domain.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Authentication Services**.
4. Add the ESXi host to an Active Directory domain.
 - a) In the Authentication Services pane, click **Join Domain**.
 - b) Supply the domain settings, and click **OK**.

The directory services type changes to Active Directory.
5. Configure or edit credentials for an NFS Kerberos user.
 - a) In the NFS Kerberos Credentials pane, click **Edit**.
 - b) Enter a user name and password.

Files stored in all Kerberos datastores are accessed using these credentials.

The state for NFS Kerberos credentials changes to Enabled.

After you configure your host for Kerberos, you can create an NFS 4.1 datastore with Kerberos enabled.

For information about creating an NFS user account to be used by the NFS 4.1 for Kerberos authentication, see VMware PowerCLI on the Broadcom Developer Portal page at <https://developer.broadcom.com/powercli/latest/vmware.vimautomation.storage/commands/new-nfsuser>.

Set Up NFS Storage Environment with ESXi

You must perform several configuration steps before you mount an NFS datastore in vSphere.

- Familiarize yourself with the guidelines in [Guidelines and Requirements for NFS Storage with ESXi](#).
 - For details on configuring NFS storage, consult your storage vendor documentation.
 - If you use Kerberos, make sure that AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 are enabled on the NAS server.
1. On the NFS server, configure an NFS volume and export it to be mounted on the ESXi hosts.
 - a) Note the IP address or the DNS name of the NFS server and the full path, or folder name, for the NFS share.

For NFS 4.1, you can collect multiple IP addresses or DNS names to use the multipathing support that the NFS 4.1 datastore provides.
 - b) If you plan to use Kerberos authentication with NFS 4.1, specify the Kerberos credentials to be used by ESXi for authentication.
 2. On each ESXi host, configure a VMkernel Network port for NFS traffic.

For more information, see the *vSphere Networking* documentation.
 3. If you plan to use Kerberos authentication with the NFS 4.1 datastore, configure the ESXi hosts for Kerberos authentication.

See [Configure ESXi Hosts for Kerberos Authentication](#).

You can now create an NFS datastore on the ESXi hosts. See [Create an NFS Datastore in vSphere Environment](#).

Configuring the NFS Datastore

You can use the `esxcli` command or the vSphere Client to configure the NFS datastore depending on the NFS version and the configuration you want to set.

Understand the esxcli Command Parameters for NFS Datastores

Use the `esxcli` command to configure specific settings on NFS.

Command	Description
NFS 4.1: <code>esxcli storage nfs41 add</code> NFS 3: <code>esxcli storage nfs add</code>	Adds a new NFS or NAS volume and mounts it with the given volume name on the ESXi host .
NFS 4.1: <code>esxcli storage nfs41 param set</code> NFS 3: <code>esxcli storage nfs param set</code>	Sets the NFS volume parameters.

The commands accept the following options:

Option	Description
<code>-H --hosts=<str></code>	NFS 4.1: The hostname or IP address of the server for the NFS 4.1 volume you want to add. You can add multiple NFS volumes separated with a comma. NFS 3: The hostname or IP address of the NAS volume you want to add and mount on the system.
<code>-I --host-vmknics=<str></code>	NFS 4.1: The hostname or IP address of the server for the NFS 4.1 volume you want to add followed by a colon(:) and the VMknics name. You can specify multiple names. Example: <code>--host-vmknics=<host1>:<vmknics1> --host-vmknics=<ip2>:<vmknics2></code> NFS 3: The hostname or IP address of the NAS volume you want to add. You can specify only one name. Example: <code><ip1>:<vmknics1></code>
<code>-p --ispe</code>	If set, this flag will set PE as the mount point.
<code>-r --readonly</code>	If set, this flag will make the mount point be read-only.
<code>-a --sec=<str></code>	Security setting. Acceptable values are: AUTH_SYS, SEC_KRB5, and SEC_KRB5I. This option is applicable only for NFS 4.1 and later.
<code>-s --share=<str></code>	The share name to use for the NFS or NAS mount point on the remote system. This option is required.
<code>-v --volume-name=<str></code>	The volume name to use for the NFS or NAS mount. The volume name must be unique and should not conflict with existing NAS, NFS, VMFS, or other volume names. This option is required.
<code>-c --connections=<long></code>	The number of RPC connections to be established for the NFS datastore. A higher number of connections results in better performance on datastores with higher I/O load.
<code>-q --maxqueuedepth=<long></code>	The maximum queue depth for the NFS volume. This option can be used only with the <code>param set</code> command.

Configure Multiple TCP Connections for NFS

You can configure NFS storage with multiple connections by using the `nconnect` option. For NFS 4.1, you can create multiple connections per session. For NFS 3, you can mount the datastore with multiple connections. Use the `esxcli storage nfs41` command to configure multiple connections.

Set up the NFS storage environment.

You can set a maximum of 4 connections per NFS datastore by default. However, you can increase it up to 8 by using the advanced NFS option. Ensure that the total number of connections across all mounted NFS datastores does not exceed 256.

For NFS 4.1, setting up multiple connections does not impact multipathing. Multiple connections and multiple paths can co-exist.

Note: The commands in the following procedure are examples for NFS 4.1 datastores. For NFS 3 datastores, replace `nfs41` with `nfs`.

1. To specify the number of connections while mounting the NFS datastore, use the following command:

```
esxcli storage nfs41 add -H <host> -v <volume-label> -s <remote_share> -c <number_of_connections>
```

2. To specify the number of connections on an existing NFS datastore, use the following command:

```
esxcli storage nfs41 param set -v <volume-label> -c <number_of_connections>
```

3. (Optional) To increase the number of connections up to 8, use the following command:

```
esxcli storage nfs41 param set -v <volume-label> -c <number_of_connections>
```

Configure VMkernel Binding for NFS Datastore

VMkernel port binding for the NFS datastore on an ESXi host allows you to bind an NFS volume to a specific VMkernel adapter to connect to an NFS server. With NFS datastores, you can isolate NFS traffic to a specific VMkernel adapter. To route the traffic to this adapter, connect the datastore to the adapter.

Note: Without binding, if the VMkernel adapter that ESXi uses for NFS traffic fails, the network infrastructure redirects the traffic to an alternative route. As a result, the NFS traffic might unintentionally flow through a random VMkernel adapter.

Configure VMkernel Binding for NFS 4.1 Datastores

To connect an NFS 4.1 datastore to a specific VMkernel adapter, use the `esxcli storage nfs41` command. You can specify the VMkernel adapters either while mounting the NFS 4.1 datastore or, if the NFS datastore is already mounted with a specific VMkernel adapter, you can modify the existing NFS volume configuration to specify the desired VMkernel adapters.

- Set up the NFS storage environment.
- Create and configure a VMkernel adapter on default TCP/IP stack.

1. To specify VMkernel adapters while mounting a new NFS v4.1 volume on the ESXi host, use the following command:

```
[root@:~] esxcli storage nfs41 add
```

Example: `[root@:~] esxcli storage nfs41 add -I 10.187.114.32:vmk0 -I 10.187.214.27:vmk1 -s /store1 -v store1`

2. To set the VMkernel adapters on an existing NFS 4.1 datastore, use the following command:

```
[root@:~] esxcli storage nfs41 param set
```

Example: `[root@:~] esxcli storage nfs41 param set -I 10.187.114.32:vmk2 -v store1`

3. (Optional) Verify whether the specified VMkernel adapters are set:

```
[root@:~] esxcli storage nfs41 list
```

The output should be similar to the following:

Volume Name	Host(s)	Share	Vmknics	Accessible	Mounted	Connections	Read-Only
Security	isPE Hardware Acceleration						
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----	-----	-----

store1	10.187.114.32,10.187.214.27	/store1	vmk2,vmk1	true	true	1
	false	AUTH_SYS	false	Not Supported		

Configure VMkernel Binding for NFS 3 Datastores

For NFS 3 datastores, you can use the vSphere Client to connect the datastore to the desired VMkernel adapters.

- Set up the NFS storage environment.
 - Create and configure a VMkernel adapter on default TCP/IP stack.
1. Bind an NFS 3 datastore to a VMkernel adapter.
 - a) In the vSphere Client navigator, browse to a host, a cluster, or a data center.
 - b) From the right-click menu, select **Storage > New Datastore**.
 - c) Select **NFS** as the datastore type, and select **NFS 3**.
 - d) Enter the datastore parameters. For more information, see [Create an NFS Datastore in vSphere Environment](#).
 - e) Select **Bind to vmknic**.
 - f) If you are creating the datastore at the data center or cluster level, select hosts that mount the datastore.
 - g) For the datastore at the data center or cluster level, select one of the following options.
 - Select **Vmkcnics on all hosts**. With this option, selectable VMkernel adapters are available across all hosts where you mount your datastore.
 - Select **Vmknic per host**. With this option, you can specify a VMkernel adapter per host.
- Note:** You can perform only one selection at a time.


h) Review the configuration options you made and click **Finish**.

2. Review the hosts connected to the NFS 3 datastore.

a) In the vSphere Client, select the NFS 3 datastore with VMkernel adapter binding.

b) Click the **Configure** tab, and click **Connectivity with Hosts**.

You can view all the hosts connected to the NFS 3 datastore. The **Vmknics Binding** column shows the VMkernel adapters connected to the



datastore

Note: If you delete the VMkernel adapter associated with the NFS 3 datastore, the datastore becomes inaccessible triggering a warning. The warning displays the text in the **Events** screen similar to the following.

```
vmknics vmkX removed, NFSv3 datastore datastore_name configured with the vmknics will be inaccessible.
```

Collecting Statistical Information for NFS Storage on ESXi Host

You can use the `nfsStats` tool in your ESXi host to display statistical information about NFS calls and Remote Procedure Calls (RPC). The command displays statistical information for NFS 3 and NFS 4.1 mounts on the ESXi host.

Generally, the `nfsStats` tool performs the following tasks.

- Gathers NFS statistics to investigate problems when you deploy a new configuration, such as a new NFS server or network, in the NFS environment.
- Provides statistics about the success and failure of NFS operations.
- Publishes latency statistics about the success and failure of NFS operations.
- Troubleshoots NFS performance issues.

The command syntax is `nfsStatsoptions`.

The following command options are available.

Table 222: nfsStats commands

Command Option	Description
No option	Obtain both NFS statistics and RPC statistics for all NFS datastores.
-3	Display only NFS 3 statistics.
-4	Display only NFS 4.1 statistics.
-n	Display only NFS 3 and NFS 4.1 statistics.
-r	Display RPC statistics.
-i <i>interval</i>	Display NFS and RPC statistics at an interval equal to the specified value in seconds. For example, if the value you enter is 10, the statistics is refreshed every 10 seconds.
-j	Display statistics in json format.
-e	Display error statistics for NFS 3, NFS 4.1, and RPC protocols.
-d, --datastore-name <i>DSNAME1</i> , <i>DSNAME2</i> , ...	Display both NFS and RPC statistics for the specified NFS datastores. Use this option in conjunction with the type of the NFS datastore, for example, -3 or -4.

Using Datastore Browser in vSphere Environment

Use the datastore file browser in your vSphere environment to manage contents of your datastores. You can upload files or folders, download files from the datastore, or move or copy folders or files to a new location. You can also use the datastore browser to inflate the thin virtual disk.

Navigate to the Datastore Browser

Navigate to your datastore to browse folders and files that are stored on the datastore and perform administrative tasks on your folders and files.

1. Open the datastore browser.
 - a) Display the datastore in the inventory.
 - b) Right-click the datastore and select **Browse Files**.
2. Explore the contents of the datastore by navigating to existing folders and files.
3. Perform administrative tasks by using the icons and options.

Icons and Options	Descriptions
Upload Files	Upload a file to the datastore.
Upload Folder (Available only in the vSphere Client)	Upload a folder to the datastore.
Download	Download from the datastore.
New Folder	Create a folder on the datastore.
Copy to	Copy selected folders or files to a new location, either on the same datastore or on a different datastore.
Move to	Move selected folders or files to a new location, either on the same datastore or on a different datastore.
Rename to	Rename selected files.
Delete	Delete selected folders or files.
Inflate	Convert a selected thin virtual disk to thick. This option applies only to thin-provisioned disks.

Upload Files or Folders to Datastores

Use the datastore file browser to upload files to datastores on the ESXi host. If you use the vSphere Client, you can also upload folders.

Required privilege: **Datastore > Browse Datastore** and **Host > System Management**

In addition to their traditional use as storage for virtual machines files, datastores can serve to store data or files related to virtual machines. For example, you can upload ISO images of operating systems from a local computer to a datastore on the host. You then use these images to install guest operating systems on the new virtual machines.

Note: You cannot upload files directly to the Virtual Volumes datastores. You must first create a folder on the Virtual Volumes datastore, and then upload the files into the folder. The created folders in Virtual Volumes datastores for block storage have a limited storage capacity space of 4GB. The Virtual Volumes datastore supports direct uploads of folders.

1. Open the datastore browser.
 - a) Display the datastore in the inventory.
 - b) Right-click the datastore and select **Browse Files**.
2. Optional: Create a folder to store the file or folder.
3. Upload the file or folder.

Option	Description
Upload a file	1. Select the target folder and click Upload Files. 2. Locate the item to upload on the local computer and click Open.
Upload a folder (available only in the vSphere Client)	1. Select the datastore or the target folder and click Upload Folders. 2. Locate the item to upload on the local computer and click Ok.

4. Refresh the datastore file browser to see the uploaded files or folders on the list.

Download Files from Datastores

Use the datastore file browser to download files from the datastore available on your ESXi host to your local computer.

- Required privilege: **Datastore > Browse Datastore**
 - Privileges when a VM is encrypted: **Cryptographic operations > Decrypt** and **Cryptographic operations > Direct Access**
1. Open the datastore browser.
 - a) Display the datastore in the inventory.
 - b) Right-click the datastore and select **Browse Files**.
 2. Navigate to the file to download and click **Download**.
 3. Follow the prompts to save the file to your local computer.

Rename Datastore Files

Use the datastore browser to rename files.

Required privilege: **Datastore > Browse Datastore**

1. Open the datastore browser.
 - a) Display the datastore in the inventory.
 - b) Right-click the datastore and select **Browse Files**.
2. Browse to a file you want to rename.
3. Select the file and click **Rename to**.
4. Specify the new name and click **OK**.

Move or Copy Datastore Folders or Files

Use the datastore browser to move or copy folders or files to a new location, either on the same datastore or on a different datastore.

Required privilege: **Datastore > Browse Datastore**

Note: Virtual disk files are moved or copied without format conversion. If you move a virtual disk to a datastore that belongs to a host different from the source host, you might need to convert the virtual disk. Otherwise, you might not be able to use the disk.

You cannot copy VM files across vCenter Servers.

1. Open the datastore browser.
 - a) Display the datastore in the inventory.
 - b) Right-click the datastore and select **Browse Files**.
2. Browse to an object you want to move or copy, either a folder or a file.
3. Select the object and click **Move to** or **Copy to**.
4. Specify the destination location.
5. Optional: Select **Overwrite files and folders with matching names at the destination**.
6. Click **OK**.

Inflate Thin Virtual Disks

If you created a virtual disk in the thin format, you can change the format to thick.

- Make sure that the datastore where the virtual machine resides has enough space.
- Make sure that the virtual disk is thin.
- Remove snapshots.
- Power off your virtual machine.

You use the datastore browser to inflate the thin virtual disk.

1. In the vSphere Client, navigate to the folder of the virtual disk you want to inflate.
 - a) Navigate to the virtual machine.
 - b) Click the **Datastores** tab.

The datastore that stores the virtual machine files is listed.
 - c) Right-click the datastore and select **Browse Files**.

The datastore browser displays contents of the datastore.
2. Expand the virtual machine folder and browse to the virtual disk file that you want to convert.

The file has the `.vmdk` extension and is marked with the virtual disk () icon.

3. Select the virtual disk file and click **Inflate**.

Note: The option might not be available if the virtual disk is thick or when the virtual machine is running.

The inflated virtual disk occupies the entire datastore space originally provisioned to it.

Understanding Multipathing and Failover in the ESXi Environment

To maintain a constant connection between a host and its storage, ESXi supports multipathing. With multipathing, you can use more than one physical path that transfers data between the host and an external storage device.

If a failure of any element in the SAN network, such as an adapter, switch, or cable, occurs, ESXi can switch to another viable physical path. This process of path switching to avoid failed components is known as path failover.

In addition to path failover, multipathing provides load balancing. Load balancing is the process of distributing I/O loads across multiple physical paths. Load balancing reduces or removes potential bottlenecks.


Note: Virtual machine I/O might be delayed for up to 60 seconds while path failover takes place. With these delays, the SAN can stabilize its configuration after topology changes. In general, the I/O delays might be longer on active-passive arrays and shorter on active-active arrays.

Failover with Fibre Channel

To support multipathing, your host typically has two or more HBAs available. This configuration supplements the SAN multipathing configuration. Generally, the SAN multipathing provides one or more switches in the SAN fabric and one or more storage processors on the storage array device itself.

In the following illustration, multiple physical paths connect each server with the storage device. For example, if HBA1 or the link between HBA1 and the FC switch fails, HBA2 takes over and provides the connection. The process of one HBA taking over for another is called HBA failover.

Figure 102: Multipathing and Failover with Fibre Channel



Similarly, if SP1 fails or the links between SP1 and the switches breaks, SP2 takes over. SP2 provides the connection between the switch and the storage device. This process is called SP failover. VMware ESXi supports both HBA and SP failovers.

Host-Based Failover with iSCSI

When setting up your ESXi host for multipathing and failover, you can use multiple iSCSI HBAs or combine multiple NICs with the software iSCSI adapter.


For information on different types of iSCSI adapters, see [iSCSI Initiators](#).

When you use multipathing, specific considerations apply.

- ESXi does not support multipathing when you combine an independent hardware adapter with software iSCSI or dependent iSCSI adapters in the same host.
- Multipathing between software and dependent adapters within the same host is supported.
- On different hosts, you can mix both dependent and independent adapters.

The following illustration shows multipathing setups possible with different types of iSCSI initiators.

Figure 103: Host-Based Path Failover



Hardware iSCSI and Failover

With hardware iSCSI, the host typically has two or more hardware iSCSI adapters. The host uses the adapters to reach the storage system through one or more switches. Alternatively, the setup might include one adapter and two storage processors, so that the adapter can use different paths to reach the storage system.

On the illustration, Host1 has two hardware iSCSI adapters, HBA1 and HBA2, that provide two physical paths to the storage system. Multipathing plug-ins on your host, whether the VMkernel NMP or any third-party MPPs, have access to the paths by default. The plug-ins can monitor health of each physical path. If, for example, HBA1 or the link between HBA1 and the network fails, the multipathing plug-ins can switch the path over to HBA2.

Software iSCSI and Failover

With software iSCSI, as shown on Host 2 of the illustration, you can use multiple NICs that provide failover and load balancing capabilities for iSCSI connections.

Multipathing plug-ins do not have direct access to physical NICs on your host. As a result, for this setup, you first must connect each physical NIC to a separate VMkernel port. You then associate all VMkernel ports with the software iSCSI initiator using a port binding technique. Each VMkernel port connected to a separate NIC becomes a different path that the iSCSI storage stack and its storage-aware multipathing plug-ins can use.

For information about configuring multipathing for software iSCSI, see [Setting Up Network for iSCSI and iSER with ESXi](#).

Array-Based Failover with iSCSI

Some iSCSI storage systems manage path use of their ports automatically and transparently to ESXi.


When using one of these storage systems, your host does not see multiple ports on the storage and cannot choose the storage port it connects to. These systems have a single virtual port address that your host uses to initially communicate. During this initial communication, the storage system can redirect the host to communicate with another port on the storage system. The iSCSI initiators in the host obey this reconnection request and connect with a different port on the system. The storage system uses this technique to spread the load across available ports.

If the ESXi host loses connection to one of these ports, it automatically attempts to reconnect with the virtual port of the storage system, and should be redirected to an active, usable port. This reconnection and redirection happens quickly and generally does not disrupt running virtual machines. These storage systems can also request that iSCSI initiators reconnect to the system, to change which storage port they are connected to. This allows the most effective use of the multiple ports.

The Port Redirection illustration shows an example of port redirection. The host attempts to connect to the 10.0.0.1 virtual port. The storage system redirects this request to 10.0.0.2. The host connects with 10.0.0.2 and uses this port for I/O communication.

Note: The storage system does not always redirect connections. The port at 10.0.0.1 could be used for traffic, also.


Figure 104: Port Redirection



If the port on the storage system that is acting as the virtual port becomes unavailable, the storage system reassigns the address of the virtual port to another port on the system. Port Reassignment shows an example of this type of port

reassignment. In this case, the virtual port 10.0.0.1 becomes unavailable and the storage system reassigns the virtual port IP address to a different port. The second port responds to both addresses.

Figure 105: Port Reassignment



With this form of array-based failover, you can have multiple paths to the storage only if you use multiple ports on the ESXi host. These paths are active-active. For additional information, see [Managing iSCSI Session on ESXi Host](#).

Path Failover and Virtual Machines

When a path fails, storage I/O might pause for 30-60 seconds until your host determines that the link is unavailable and performs the failover. If you attempt to display the host, its storage devices, or its adapters, the operation might appear to stall. Virtual machines with their disks installed on the SAN can appear unresponsive. After the failover, I/O resumes normally and the virtual machines continue to run.

A Windows virtual machine might interrupt the I/O and eventually fail when failovers take too long. To avoid the failure, set the disk timeout value for the Windows virtual machine to at least 60 seconds.

To avoid disruptions during a path failover, increase the standard disk timeout value on a Windows guest operating system.

This procedure explains how to change the timeout value by using the Windows registry.

Tip: Before you begin: Back up the Windows registry.

1. Select **Start > Run**.
2. Type `regedit.exe`, and click **OK**.
3. In the left-panel hierarchy view, double-click **HKEY_LOCAL_MACHINE > System > CurrentControlSet > Services > Disk**.
4. Double-click **TimeOutValue**.
5. Set the value data to 0x3c (hexadecimal) or 60 (decimal) and click **OK**.
After you make this change, Windows waits at least 60 seconds for delayed disk operations to finish before it generates errors.
6. Reboot guest OS for the change to take effect.

Using Pluggable Storage Architecture and Path Management with ESXi

To manage multipathing, ESXi uses a special VMkernel layer, Pluggable Storage Architecture (PSA). The PSA is an open and modular framework that coordinates various software modules responsible for multipathing operations, VMware NMP, VMware HPP, and third-party MPPs.


VMware provides generic native multipathing modules, called VMware NMP and VMware HPP. In addition, the PSA offers a collection of VMkernel APIs that third-party developers can use. The software developers can create their own load balancing and failover modules for a particular storage array. These third-party multipathing modules (MPPs) can be installed on the ESXi host and run in addition to the VMware native modules, or as their replacement.

When coordinating the VMware native modules and any installed third-party MPPs, the PSA performs the following tasks:

- Loads and unloads multipathing plug-ins.
- Hides virtual machine specifics from a particular plug-in.
- Routes I/O requests for a specific logical device to the MPP managing that device.
- Handles I/O queueing to the logical devices.
- Implements logical device bandwidth sharing between virtual machines.
- Handles I/O queueing to the physical storage HBAs.
- Handles physical path discovery and removal.
- Provides logical device and physical path I/O statistics.

As the Pluggable Storage Architecture illustration shows, multiple third-party MPPs can run in parallel with the VMware NMP or HPP. When installed, the third-party MPPs can replace the behavior of the native modules. The MPPs can take control of the path failover and the load-balancing operations for the specified storage devices.

Figure 106: Pluggable Storage Architecture



Multipathing Concepts and Acronyms

Multipathing Concept	Acronym	Definition
Pluggable Storage Architecture	PSA	The PSA is an open and modular framework that coordinates various software modules responsible for multipathing operations. These modules include generic multipathing modules that VMware provides, NMP and HPP, and third-party MPPs.
Native Multipathing Plug-in	NMP	The NMP is the VMkernel multipathing module that ESXi provides by default and is used for SCSI storage devices. The NMP associates physical paths with a specific storage device and provides a default path selection algorithm based on the array type. The NMP is extensible and manages additional submodules, called Path Selection Policies (PSPs) and Storage Array Type Policies (SATPs). PSPs and SATPs can be provided by VMware, or by a third party. For more information, see VMware Native Multipathing Plug-In .

Multipathing Concept	Acronym	Definition
Path Selection Plug-in	PSP	The PSPs are submodules of the VMware NMP. PSPs are responsible for selecting a physical path for SCSI I/O requests. For more information, see VMware Path Selection Plug-Ins and Policies .
Storage Array Type Plug-in	SATP	SATPs are responsible for array-specific operations for a SCSI storage array. The SATP can determine the state of a particular array-specific path, perform a path activation, and detect any path errors. For more information, see VMware SATPs .
Multipathing Plug-in	MPP (third-party)	The PSA offers a collection of VMkernel APIs that third parties can use to create their own multipathing plug-ins (MPPs). The modules provide specific load balancing and failover functionalities for a particular storage array. The MPPs can be installed on the ESXi host. They can run in addition to the VMware native modules, or as their replacement.
Native High-Performance Plug-in	HPP	The HPP replaces the NMP for high-speed devices, such as NVMe. The HPP can improve the performance of ultra-fast flash devices that are installed locally on your ESXi host, and is the default plug-in that claims NVMe-oF targets. For information, see VMware High Performance Plug-In and Path Selection Schemes .
Path Selection Scheme	PSS	To support multipathing, the HPP uses the Path Selection Schemes (PSS). A particular PSS is responsible for selecting physical paths for I/O requests and handles multipathing for NVMe storage devices. For information, see VMware High Performance Plug-In and Path Selection Schemes .
Claim Rules		The PSA uses claim rules to determine which plug-in owns the paths to a particular storage device. See Using Claim Rules to Control ESXi Multipathing Modules .

VMware Native Multipathing Plug-In

By default, ESXi provides an extensible multipathing module called Native Multipathing Plug-In (NMP).

Generally, the VMware NMP supports all storage arrays listed on the VMware storage HCL and provides a default path selection algorithm based on the array type. The NMP associates a set of physical paths with a specific storage device, or LUN.

For additional multipathing operations, the NMP uses submodules, called SATPs and PSPs. The NMP delegates to the SATP the specific details of handling path failover for the device. The PSP handles path selection for the device.

Typically, the NMP performs the following operations:

- Manages physical path claiming and unclaiming.
- Registers and de-registers logical devices.
- Associates physical paths with logical devices.
- Supports path failure detection and remediation.
- Processes I/O requests to logical devices:
 - Selects an optimal physical path for the request.
 - Performs actions necessary to handle path failures and I/O command retries.
- Supports management tasks, such as reset of logical devices.

ESXi automatically installs an appropriate SATP for an array you use. You do not need to obtain or download any SATPs.

VMware NMP Flow of I/O

When a virtual machine issues an I/O request to a storage device managed by the NMP, the following process takes place.

1. The NMP calls the PSP assigned to this storage device.
2. The PSP selects an appropriate physical path on which to issue the I/O.
3. The NMP issues the I/O request on the path selected by the PSP.
4. If the I/O operation is successful, the NMP reports its completion.
5. If the I/O operation reports an error, the NMP calls the appropriate SATP.
6. The SATP interprets the I/O command errors and, when appropriate, activates the inactive paths.
7. The PSP is called to select a new path on which to issue the I/O.

Display Multipathing Modules

Use the `esxcli` command to list all multipathing modules loaded into the system. Multipathing modules manage physical paths that connect your host with storage. The modules include VMware native NMP and HPP, and any third-party MPPs.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To list multipathing modules, run the following command:

```
esxcli storage core plugin list --plugin-class=MP
```

This command typically shows the NMP and, if loaded, the HPP and the MASK_PATH module. If any third-party MPPs have been loaded, they are listed as well.

Plugin name	Plugin class
-----	-----
NMP	MP

For more information about the command, see the *ESXCLI Concepts and Examples* and *ESXCLI Reference* documentation.

Display NMP Storage Devices

Use the `esxcli` command to list all storage devices controlled by the VMware NMP and display SATP and PSP information associated with each device.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To list all storage devices, run the following command:

```
esxcli storage nmp device list
```

Use the `--device | -d=device_ID` parameter to filter the output of this command to show a single device.

Displaying NMP Storage Devices

```
# esxcli storage nmp device list
mpx.vmhbal:C0:T2:L0
  Device Display Name: Local VMware Disk (mpx.vmhbal:C0:T2:L0)
  Storage Array Type: VMW_SATP_LOCAL
  Storage Array Type Device Config: SATP VMW_SATP_LOCAL does not support device configuration.
  Path Selection Policy: VMW_PSP_FIXED
  Path Selection Policy Device Config: {preferred=vmhbal:C0:T2:L0;current=vmhbal:C0:T2:L0}
  Path Selection Policy Device Custom Config:
  Working Paths: vmhbal:C0:T2:L0
```

```
Is USB: false
```

```
.....
```

```
eui.6238666462643332
```

```
Device Display Name: SCST_BIO iSCSI Disk (eui.6238666462643332)
```

```
Storage Array Type: VMW_SATP_DEFAULT_AA
```

```
Storage Array Type Device Config: {action_OnRetryErrors=off}
```

```
Path Selection Policy: VMW_PSP_FIXED
```

```
Path Selection Policy Device Config: {preferred=vmhba65:C0:T0:L0;current=vmhba65:C0:T0:L0}
```

```
Path Selection Policy Device Custom Config:
```

```
Working Paths: vmhba65:C0:T0:L0
```

```
Is USB: false
```

For more information about the command, see the *ESXCLI Concepts and Examples* and *ESXCLI Reference* documentation.

VMware Path Selection Plug-Ins and Policies

VMware Path Selection Plug-ins (PSPs) are responsible for selecting a physical path for I/O requests.

The plug-ins are submodules of the VMware NMP. The NMP assigns a default PSP for each logical device based on the device type. You can override the default PSP. For more information, see [Change the Path Selection Policy](#).

Each PSP enables and enforces a corresponding path selection policy.

VMW_PSP_MRU - Most Recently Used (VMware)

The Most Recently Used (VMware) policy is enforced by VMW_PSP_MRU. It selects the first working path discovered at system boot time. When the path becomes unavailable, the host selects an alternative path. The host does not revert to the original path when that path becomes available. The Most Recently Used policy does not use the preferred path setting. This policy is default for most active-passive storage devices.

The VMW_PSP_MRU supports path ranking. To set ranks to individual paths, use the `esxcli storage nmp psp generic pathconfig set` command. For details, see the VMware knowledge base article at <http://kb.vmware.com/kb/2003468> and the *ESXCLI Reference* documentation.

VMW_PSP_FIXED - Fixed (VMware)

This Fixed (VMware) policy is implemented by VMW_PSP_FIXED. The policy uses the designated preferred path. If the preferred path is not assigned, the policy selects the first working path discovered at system boot time. If the preferred path becomes unavailable, the host selects an alternative available path. The host returns to the previously defined preferred path when it becomes available again.

VMW_PSP_RR - Round Robin (VMware)

Fixed is the default policy for most active-active storage devices. VMW_PSP_RR enables the Round Robin (VMware) policy. Round Robin is the default policy for many arrays. It uses an automatic path selection algorithm rotating through the configured paths. Both active-active and active-passive arrays use the policy to implement load balancing across paths for different LUNs. With active-passive arrays, the policy uses active paths. With active-active arrays, the policy uses available paths.

The latency mechanism that can be activated for the policy makes it more adaptive. To achieve better load balancing results, the mechanism dynamically selects an optimal path by considering the following path characteristics:

- I/O bandwidth

- Path latency

To change the default parameters for the adaptive latency Round Robin policy or to deactivate the latency mechanism, see the [Change Default Parameters for Latency Round Robin](#).

To set other configurable parameters for VMW_PSP_RR, use the `esxcli storage nmp psp roundrobin` command. For details, see the *ESXCLI Reference* documentation.

VMware SATPs

Storage Array Type Plug-ins (SATPs) are responsible for array-specific operations. The SATPs are submodules of the VMware NMP.

ESXi offers an SATP for every type of array that VMware supports. ESXi also provides default SATPs that support non-specific active-active, active-passive, ALUA, and local devices.

Each SATP accommodates special characteristics of a certain class of storage arrays. The SATP can perform the array-specific operations required to detect path state and to activate an inactive path. As a result, the NMP module itself can work with multiple storage arrays without having to be aware of the storage device specifics.

Generally, the NMP determines which SATP to use for a specific storage device and associates the SATP with the physical paths for that storage device. The SATP implements the tasks that include the following:

- Monitors the health of each physical path.
- Reports changes in the state of each physical path.
- Performs array-specific actions necessary for storage fail-over. For example, for active-passive devices, it can activate passive paths.

ESXi includes several generic SATP modules for storage arrays.

VMW_SATP_LOCAL

SATP for local direct-attached devices.

As of vSphere 6.5 Update 2 release, VMW_SATP_LOCAL provides multipathing support for the local devices, except the devices in 4K native format. To claim multiple paths to the local devices, you are no longer required to use other SATPs as you were in the earlier vSphere releases.

VMW_SATP_LOCAL supports the VMW_PSP_MRU and VMW_PSP_FIXED path selection plug-ins, but does not support VMW_PSP_RR.

VMW_SATP_DEFAULT_AA

Generic SATP for active-active arrays.

VMW_SATP_DEFAULT_AP

Generic SATP for active-passive arrays.

VMW_SATP_ALUA

SATP for ALUA-compliant arrays.

For more information, see the *VMware Compatibility Guide* and the *ESXCLI Reference* documentation.

Display SATPs for the Host

Use the `esxcli` command to list VMware NMP SATPs loaded into the system. Display information about the SATPs.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To list VMware SATPs, run the following command:

```
esxcli storage nmp satp list
```

For each SATP, the output displays information that shows the type of storage array or system the SATP supports. The output also shows the default PSP for any LUNs that use this SATP. Placeholder (plugin not loaded) in the Description column indicates that the SATP is not loaded.

Displaying SATPs for the Host

```
# esxcli storage nmp satp list
Name                Default PSP      Description
VMW_SATP_MSA        VMW_PSP_MRU     Placeholder (plugin not loaded)
VMW_SATP_ALUA        VMW_PSP_MRU     Placeholder (plugin not loaded)
VMW_SATP_DEFAULT_AP VMW_PSP_MRU     Placeholder (plugin not loaded)
VMW_SATP_SVC         VMW_PSP_FIXED   Placeholder (plugin not loaded)
VMW_SATP_EQL         VMW_PSP_FIXED   Placeholder (plugin not loaded)
VMW_SATP_INV         VMW_PSP_FIXED   Placeholder (plugin not loaded)
VMW_SATP_EVA         VMW_PSP_FIXED   Placeholder (plugin not loaded)
VMW_SATP_ALUA_CX     VMW_PSP_RR      Placeholder (plugin not loaded)
VMW_SATP_SYMM        VMW_PSP_RR      Placeholder (plugin not loaded)
VMW_SATP_CX          VMW_PSP_MRU     Placeholder (plugin not loaded)
VMW_SATP_LSI         VMW_PSP_MRU     Placeholder (plugin not loaded)
VMW_SATP_DEFAULT_AA  VMW_PSP_FIXED   Supports non-specific active/active arrays
VMW_SATP_LOCAL       VMW_PSP_FIXED   Supports direct attached devices
```

For more information about the command, see the *ESXCLI Concepts and Examples* and *ESXCLI Reference* documentation.

VMware High Performance Plug-In and Path Selection Schemes

VMware provides the High-Performance Plug-in (HPP) to improve the performance of storage devices on your ESXi host.

The HPP replaces the NMP for high-speed devices, such as NVMe. The HPP is the default plug-in that claims NVMe-oF targets. ESXi supports end to end NVMe without emulation and SCSI to NVMe emulation. The HPP supports only active/active and implicit ALUA targets.

HPP Support	vSphere 8.0 Update 3
Storage devices	Local NVMe and SCSI Shared NVMe-oF (active/active and implicit ALUA targets only)
Multipathing	Yes
Second-level plug-ins	No
SCSI-3 persistent reservations	No
4Kn devices with software emulation	Yes

Path Selection Schemes

To support multipathing, the HPP uses the Path Selection Schemes (PSS) when selecting physical paths for I/O requests.

You can use the vSphere Client or the `esxcli` command to change the default path selection mechanism.

For information about configuring the path mechanisms in the vSphere Client, see [Change the Path Selection Policy](#). To configure with the `esxcli` command, see [ESXi esxcli HPP Commands](#).

ESXi supports the following path selection mechanisms.

FIXED

With this scheme, a designated preferred path is used for I/O requests. If the preferred path is not assigned, the host selects the

first working path discovered at the boot time. If the preferred path becomes unavailable, the host selects an alternative available path. The host returns to the previously defined preferred path when it becomes available again.

When you configure **FIXED** as a path selection mechanism, select the preferred path.

LB-RR (Load Balance - Round Robin)

This is the default scheme for the devices claimed by HPP. After transferring a specified number of bytes or I/Os on a current path, the scheme selects the path using the round robin algorithm.

To configure the **LB-RR** path selection mechanism, specify the following properties:

- **IOPS** indicates the I/O count on the path to be used as criteria to switch a path for the device.
- **Bytes** indicates the byte count on the path to be used as criteria to switch a path for the device.

LB-IOPS (Load Balance - IOPs)

After transferring a specified number of I/Os on a current path, default is 1000, the system selects an optimal path that has the least number of outstanding I/Os.

When configuring this mechanism, specify the **IOPS** parameter to indicate the I/O count on the path to be used as criteria to switch a path for the device.

LB-BYTES (Load Balance - Bytes)

After transferring a specified number of bytes on a current path, default is 10 MB, the system selects an optimal path that has the least number of outstanding bytes.

To configure this mechanism, use the **Bytes** parameter to indicate the byte count on the path to be used as criteria to switch a path for the device.

Load Balance - Latency (LB-Latency)

To achieve better load balancing results, the mechanism dynamically selects an optimal path by considering the following path characteristics:

- **Latency evaluation time** parameter indicates at what time interval, in milliseconds, the latency of paths must be evaluated.
- **Sampling I/Os per path** parameter controls how many sample I/Os must be issued on each path to calculate latency of the path.

HPP Best Practices

To achieve the fastest throughput from a high-speed storage device, follow these recommendations.

- Use the vSphere version that supports the HPP.
- Use the HPP for local NVMe and SCSI devices, and NVMe-oF devices.
- If you use NVMe over Fibre Channel devices, follow general recommendations for Fibre Channel storage. See [Using ESXi with Fibre Channel SAN](#).
- If you use NVMe-oF, do not mix transport types to access the same namespace.
- When using NVMe-oF namespaces, make sure that active paths are presented to the host. The namespaces cannot be registered until the active path is discovered.
- When you configure your VMs, you can use VMware Paravirtual controllers or add NVMe controllers. Both types have their advantages and disadvantages. To check which works best for your environment, see [SCSI, SATA, and](#)

[NVMe Storage Controller Conditions, Limitations, and Compatibility](#) in the *vSphere Virtual Machine Administration* documentation.

- Set the latency sensitive threshold.
- If a single VM drives a significant share of the device's I/O workload, consider spreading the I/O across multiple virtual disks. Attach the disks to separate virtual controllers in the VM.

Otherwise, I/O throughput might be limited due to saturation of the CPU core responsible for processing I/Os on a particular virtual storage controller.

For information about device identifiers for NVMe devices that support only NGUID ID format, see [NVMe Devices with NGUID Device Identifiers](#).

Enable the High-Performance Plug-In and the Path Selection Schemes

The high-performance plug-in (HPP) is the default plug-in that claims local NVMe and SCSI devices, and NVMe-oF targets. If necessary, you can replace it with NMP. In vSphere version 7.0 Update 1 and earlier, NMP remains the default plug-in for local NVMe and SCSI devices, but you can replace it with HPP.

Set up your VMware NVMe storage environment. For more information, see [About VMware NVMe Storage](#).

Use the `esxcli storage core claimrule add` command to enable the HPP or NMP on your ESXi host.

To run the `esxcli storage core claimrule add`, you can use the ESXi Shell or vSphere CLI. For more information, see *Getting Started with ESXCLI* and *ESXCLI Reference*.

Examples in this topic demonstrate how to enable HPP and set up the path selection schemes (PSS).

Note: Enabling the HPP is not supported on PXE booted ESXi hosts.

1. Create an HPP claim rule by running the `esxcli storage core claimrule add` command.

Use one of the following methods to add the claim rule.

Method	Description
Based on the NVMe controller model	<pre>esxcli storage core claimrule add --type vendor --nvme-controller-model</pre> For example, <code>esxcli storage core claimrule add --rule 429 --type vendor --nvme-controller-model "ABCD*" --plugin HPP</code>
Based on the PCI vendor ID and subvendor ID	<pre>esxcli storage core claimrule add --type vendor --pci-vendor-id --pci-sub-vendor-id</pre> For example, <code>esxcli storage core claimrule add --rule 429 --type vendor --pci-vendor-id 8086 --pci-sub-vendor-id 8086 --plugin HPP.</code>

2. Configure the PSS.

Use one of the following methods.

Method	Description
Set the PSS based on the device ID	<pre>esxcli storage hpp device set</pre> For example, <code>esxcli storage hpp device set --device=device --pss=FIXED --path=preferred path</code>
Set the PSS based on the vendor/model	Use the <code>--config-string</code> option with the <code>esxcli storage core claimrule add</code> command.

Method	Description
	For example, <code>esxcli storage core claimrule add -r 914 -t vendor -V vendor -M model -P HPP --config-string "pss=LB-Latency,latency-eval-time=40000"</code>

3. Reboot your host for your changes to take effect.

Set Latency Sensitive Threshold

When you use the HPP for your storage devices, set the latency sensitive threshold for the device, so that I/O can avoid the I/O scheduler.

By default, ESXi passes every I/O through the I/O scheduler. However, using the scheduler might create internal queuing, which is not efficient with the high-speed storage devices.

You can configure the latency sensitive threshold and enable the direct submission mechanism that helps I/O to bypass the scheduler. With this mechanism enabled, the I/O passes directly from PSA through the HPP to the device driver.

For the direct submission to work properly, the observed average I/O latency must be lower than the latency threshold you specify. If the I/O latency exceeds the latency threshold, the system stops the direct submission and temporarily reverts to using the I/O scheduler. The direct submission is resumed when the average I/O latency drops below the latency threshold again.

You can set the latency threshold for a family of devices claimed by HPP. Set the latency threshold using the vendor and model pair, the controller model, or PCIe vendor ID and sub vendor ID pair.

1. Set the latency sensitive threshold for the device by running the following command:

```
esxcli storage core device latencythreshold set -t value in milliseconds
```

Use one of the following options.

Option	Example
Vendor/model	Set the latency sensitive threshold parameter for all devices with the indicated vendor and model: <code>esxcli storage core device latencythreshold set -v 'vendor1' -m 'model1' -t 10</code>
NVMe controller model	Set the latency sensitive threshold for all NVMe devices with the indicated controller model: <code>esxcli storage core device latencythreshold set -c 'controller_model1' -t 10</code>
PCIe vendor/subvendor ID	Set the latency sensitive threshold for devices with 0x8086 as PCIe vendor ID and 0x8086 as PCIe sub vendor ID. <code>esxcli storage core device latencythreshold set -p '8086' -s '8086' -t 10</code>

2. Verify that the latency threshold is set:

```
esxcli storage core device latencythreshold list
```

Device	Latency Sensitive Threshold
naa.55cd2e404c1728aa	0 milliseconds
naa.500056b34036cdfd	0 milliseconds
naa.55cd2e404c172bd6	50 milliseconds

3. Monitor the status of the latency sensitive threshold. Check VMkernel logs for the following entries:

- Latency Sensitive Gatekeeper turned on for device *device*. Threshold of *XX* msec is larger than max completion time of *YYY* msec
- Latency Sensitive Gatekeeper turned off for device *device*. Threshold of *XX* msec is exceeded by command completed in *YYY* msec

ESXi esxcli HPP Commands

You can use the ESXi Shell or vSphere CLI commands to configure and monitor the high-performance plug-in.

See *Getting Started with ESXCLI* for an introduction, and *ESXCLI Reference* for details of the `esxcli` command use.

Command	Description	Options
<code>esxcli storage hpp path list</code>	List the paths currently claimed by the high-performance plug-in.	<code>-d --device=device</code> Display information for a specific device. <code>-p --path=path</code> Limit the output to a specific path.
<code>esxcli storage hpp device list</code>	List the devices currently controlled by the high-performance plug-in.	<code>-d --device=device</code> Show a specific device.
<code>esxcli storage hpp device set</code>	Configure settings for an HPP device.	<code>-B --bytes=long</code> Maximum bytes on the path, after which the path is switched. <code>--cfg-file</code> Update the configuration file and runtime with the new setting. If the device is claimed by another PSS, ignore any errors when applying to runtime configuration. <code>-d --device=device</code> The HPP device upon which to operate. Use any of the UIDs that the device reports. Required. <code>-I --iops=long</code> Maximum IOPS on the path, after which the path is switched. <code>-T --latency-eval-time=long</code> Control at what interval, in ms, the latency of paths must be evaluated. <code>-L --mark-device-local=bool</code> Set HPP to treat the device as local or not. <code>-M --mark-device-ssd=bool</code> Specify whether or not the HPP treats the device as an SSD. <code>-p --path=str</code> The path to set as the preferred path for the device. <code>-P --pss=pss_name</code> The path selection scheme to assign to the device. If you do not specify the value, the system selects the default. For the description of path selection schemes, see VMware High Performance Plug-In and Path Selection Schemes . Options include: • <code>FIXED</code> Use the <code>-p --path=str</code> suboption to set the preferred path.

Command	Description	Options
		• LB-Bytes Use the <code>-B --bytes=long</code> suboption to specify the input. • LB-IOPs Use the <code>-I --iops=long</code> suboption to specify the input. • LB-Latency Suboptions include: <code>-T --latency-eval-time=long</code> <code>-S --sampling-ios-per-path=long</code> • LB-RR Default Suboptions include: <code>-B --bytes=long</code> <code>-I --iops=long</code> <code>-S --sampling-ios-per-path=long</code> Control how many sample I/Os must be issued on each path to calculate latency of the path. <code>-U --use-ano=bool</code> Set the option to <code>true</code> to include non-optimized paths in the set of active paths used to issue I/Os on this device. Otherwise, set the option to <code>false</code>.
<code>esxcli storage hpp device usermarkedssd list</code>	List the devices that were marked or unmarked as SSD by user.	<code>-d --device=device</code> Limit the output to a specific device.

Viewing and Managing Storage Paths on ESXi Hosts

When you start your ESXi host or rescan your storage adapter, the host discovers all physical paths to storage devices available to the host. Based on a set of claim rules, the host determines which multipathing module, the NMP, HPP, or an MPP, owns the paths to a particular device.

The module that owns the device becomes responsible for managing the multipathing support for the device. By default, the host performs a periodic path evaluation every five minutes and assigns unclaimed paths to the appropriate module.

For the paths managed by the NMP module, a second set of claim rules is used. These rules assign an SATP and PSP modules to each storage device and determine which Storage Array Type Policy and Path Selection Policy to apply.

Use the vSphere Client to view the Storage Array Type Policy and Path Selection Policy assigned to a specific storage device. You can also check the status of all available paths for this storage device. If needed, you can change the default Path Selection Policy using the client.

To change the default multipathing module or SATP, modify claim rules using the vSphere CLI.

You can find some information about modifying claim rules in [Using Claim Rules to Control ESXi Multipathing Modules](#).

To learn more about multipathing modules, see [Multipathing Concepts and Acronyms](#).

View Storage Device Paths

View which multipathing policies the host uses for a specific storage device and the status of all available paths for this storage device.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Devices**.
4. Select the storage device whose paths you want to view.
5. Click the **Properties** tab and review the module that owns the device, for example, NMP or HPP.
Under Multipathing Policies, you can also see the Path Selection Policy and, if applicable, the Storage Array Type Policy assigned to the device.
6. Click the **Paths** tab to review all paths available for the storage device and the status of each path. The following path status information can appear:

Status	Description
Active (I/O)	Working path or multiple paths that currently transfer data.
Standby	Paths that are inactive. If the active path fails, they can become operational and start transferring I/O.
Disabled	Paths that are deactivated by the administrator.
Dead	Paths that are no longer available for processing I/O. A physical medium failure or array misconfiguration can cause this status.

If you are using the **Fixed** path policy, you can see which path is the preferred path. The preferred path is marked with an asterisk (*) in the Preferred column.

View Datastore Paths

Review the paths that connect to storage devices backing your VMFS datastores.

1. In the vSphere Client, navigate to the datastore.
2. Click the **Configure** tab.
3. Click **Connectivity and Multipathing**.
4. Select a host to view multipathing details for its devices.
5. Under Multipathing Policies, review the module that owns the device, such as NMP. You can also see the Path Selection Policy and Storage Array Type Policy assigned to the device.

For example, you might see the following:

Path Selection Policy	Preferred Path
Storage Array Type Policy	VMW_SATP_LOCAL
Owner Plugin	NMP

6. Under Paths, review the device paths and the status of each path. The following path status information can appear:

Status	Description
Active (I/O)	Working path or multiple paths that currently transfer data.
Standby	Paths that are inactive. If the active path fails, they can become operational and start transferring I/O.
Disabled	Paths that are deactivated by the administrator.

Status	Description
Dead	Paths that are no longer available for processing I/O. A physical medium failure or array misconfiguration can cause this status.

If you are using the **Fixed** path policy, you can see which path is the preferred path. The preferred path is marked with an asterisk (*) in the Preferred column.

6. Select a path policy and configure its settings. Your options change depending on the type of a storage device you use.
 - For information about path policies for SCSI devices, see [VMware Path Selection Plug-Ins and Policies](#).
 - For information about path mechanisms for NVMe devices, see [VMware High Performance Plug-In and Path Selection Schemes](#).
7. To save your settings and exit the dialog box, click **OK**.

Change Default Parameters for Latency Round Robin

On the ESXi host, you can activate the latency mechanism for the Round Robin path selection policy. The mechanism considers I/O bandwidth and path latency to select an optimal path for I/O. When using the latency mechanism, the Round Robin policy can dynamically select the optimal path and achieve better load balancing results.

Set the path selection policy to Round Robin. See [Change the Path Selection Policy](#).

When latency is activated, usage of paths for the I/Os depends on the workload. When a large number of paths is available, but the workload is not high enough, the system might not use some paths for I/Os.

You use the `esxcli` command to change the default parameters of the latency mechanism or deactivate the mechanism.

1. Configure the latency mechanism by using the following command.

```
esxcli storage nmp psp roundrobin deviceconfig set --type=latency --device=device ID
```

The command takes the following parameters:

Parameter	Description
-S --num-sampling-cycles= <i>sampling value</i>	When <code>--type</code> is set to <code>latency</code> , this parameter controls how many I/Os to use to calculate the average latency of each path. The default value of this parameter is 16.
-T --latency-eval-time= <i>time in ms</i>	When <code>--type</code> is set to <code>latency</code> , this parameter controls the frequency at which the latency of paths is updated. Default is 3 minutes.

2. Verify whether the latency Round Robin and its parameters are configured correctly.

```
esxcli storage nmp psp roundrobin deviceconfig get --device=device ID
```

or

```
esxcli storage nmp device list --device=device ID
```

The following sample output shows the path's configuration:

```
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config: {policy=latency,latencyEvalTime=180000,samplingCycles=16,curSamplingCycle=16,useANO=0; CurrentPath=vmhbal:C0:T0:L0: NumIOsPending=0,latency=0}
```

To deactivate the latency mechanism, in the Advanced System Settings for your host, change the `Misc.EnablePSPLatencyPolicy` parameter to 0.

Deactivate Storage Paths

You can temporarily deactivate paths for maintenance or other reasons.

You deactivate a path using the Paths panel. You have several ways to access the Paths panel, from a datastore, a storage device, an adapter, or a Virtual Volumes Protocol Endpoint view.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click one of the following items:
 - **Storage Adapters**
 - **Storage Devices**
 - **Protocol Endpoints**
4. In the right pane, select the item whose paths you want to deactivate, an adapter, storage device, or Protocol Endpoint, and click the **Paths** tab.
5. Select the path to deactivate and click **Disable**.
The path's status changes to Disabled.

Using Claim Rules to Control ESXi Multipathing Modules

Claim rules determine which multipathing module owns the paths to a particular storage device. They also define the type of multipathing support that the host provides to the device.

To learn more about multipathing modules, see [Multipathing Concepts and Acronyms](#).

The rules fall into these categories:

- **Core Claim Rules**. These claim rules determine which multipathing module, the NMP, HPP, or a third-party MPP, claims the specific device.
- **SATP Claim Rules**. Depending on the device type, these rules assign a particular SATP submodule that provides vendor-specific multipathing management to the device.

You can use the `esxcli` commands to add or change the core and SATP claim rules. Typically, you add the claim rules to load a third-party MPP or to hide a LUN from your host. Changing claim rules might be necessary when default settings for a specific device are not sufficient.

For more information about commands available to manage PSA claim rules, see the *Getting Started with ESXCLI*.

For a list of storage arrays and corresponding SATPs and PSPs, see the Storage/SAN section of the *vSphere Compatibility Guide*.

Multipathing Considerations

Specific considerations apply when you manage storage multipathing plug-ins and claim rules.

The following considerations help you with multipathing:

- If no SATP is assigned to the device by the claim rules, the default SATP for iSCSI or FC devices is VMW_SATP_DEFAULT_AA. The default PSP is VMW_PSP_FIXED.
- When the system searches the SATP rules to locate a SATP for a given device, it searches the driver rules first. If there is no match, the vendor/model rules are searched, and finally the transport rules are searched. If no match occurs, NMP selects a default SATP for the device.
- If VMW_SATP_ALUA is assigned to a specific storage device, but the device is not ALUA-aware, no claim rule match occurs for this device. The device is claimed by the default SATP based on the device's transport type.

- The default PSP for all devices claimed by VMW_SATP_ALUA is VMW_PSP_MRU. The VMW_PSP_MRU selects an active/optimized path as reported by the VMW_SATP_ALUA, or an active/unoptimized path if there is no active/optimized path. This path is used until a better path is available (MRU). For example, if the VMW_PSP_MRU is currently using an active/unoptimized path and an active/optimized path becomes available, the VMW_PSP_MRU will switch the current path to the active/optimized one.
- While VMW_PSP_MRU is typically selected for ALUA arrays by default, certain ALUA storage arrays need to use VMW_PSP_FIXED. To check whether your storage array requires VMW_PSP_FIXED, see the *VMware Compatibility Guide* or contact your storage vendor. When using VMW_PSP_FIXED with ALUA arrays, unless you explicitly specify a preferred path, the ESXi host selects the most optimal working path and designates it as the default preferred path. If the host selected path becomes unavailable, the host selects an alternative available path. However, if you explicitly designate the preferred path, it will remain preferred no matter what its status is.
- By default, the PSA claim rule 101 masks Dell array pseudo devices. Do not delete this rule, unless you want to unmask these devices.

List Multipathing Claim Rules for the Host

Use the `esxcli` command to list available multipathing claim rules.

Claim rules indicate whether the NMP, HPP, or a third-party MPP manages a given physical path. Each claim rule identifies a set of paths based on the following parameters:

- Vendor/model strings
- Transportation, such as SATA, IDE, Fibre Channel
- Adapter, target, or LUN location
- Device driver, for example, Mega-RAID

List the multipathing claim rules by running the `esxcli storage core claimrule list --claimrule-class=MP` command.

If you do not use the `claimrule-class` option, the MP rule class is implied.

Sample Output of the `esxcli storage core claimrule list` Command

Rule	Class	Rule	Class	Type	Plugin	Matches
MP		10	runtime	vendor	HPP	vendor=NVMe model=*
MP		10	file	vendor	HPP	vendor=NVMe model=*
MP		50	runtime	transport	NMP	transport=usb
MP		51	runtime	transport	NMP	transport=sata
MP		52	runtime	transport	NMP	transport=ide
MP		53	runtime	transport	NMP	transport=block
MP		54	runtime	transport	NMP	transport=unknown
MP		101	runtime	vendor	MASK_PATH	vendor=DELL model=Universal Xport
MP		101	file	vendor	MASK_PATH	vendor=DELL model=Universal Xport
MP		200	runtime	vendor	MPP_1	vendor=NewVend model=*
MP		200	file	vendor	MPP_1	vendor=NewVend model=*
MP		201	runtime	location	MPP_2	adapter=vmhba41 channel=* target=* lun=*
MP		201	file	location	MPP_2	adapter=vmhba41 channel=* target=* lun=*
MP		202	runtime	driver	MPP_3	driver=megaraid
MP		202	file	driver	MPP_3	driver=megaraid
MP		65535	runtime	vendor	NMP	vendor=* model=*

This example indicates the following:

- The NMP claims all paths connected to storage devices that use the USB, SATA, IDE, and Block SCSI transportation.

- The rules for HPP, MPP_1, MPP_2, and MPP_3 have been added, so that the modules can claim specified devices. For example, the HPP claims all devices with vendor NVMe. All devices handled by the inbox nvme driver are claimed regardless of the actual vendor. The MPP_1 module claims all paths connected to any model of the NewVend storage array.
- You can use the MASK_PATH module to hide unused devices from your host. By default, the PSA claim rule 101 masks Dell array pseudo devices with a vendor string `DELL` and a model string `Universal Xport`.
- The Rule Class column in the output describes the category of a claim rule. It can be MP (multipathing plug-in), Filter, or VAAI.
- The Class column shows which rules are defined and which are loaded. The `file` parameter in the Class column indicates that the rule is defined. The `runtime` parameter indicates that the rule has been loaded into your system. For a user-defined claim rule to be active, two lines with the same rule number must exist, one line for the rule with the `file` parameter and another line with `runtime`. Several default system-defined claim rules have only one line with the Class of `runtime`. You cannot modify these rules.
- The default rule 65535 assigns all unclaimed paths to the NMP. Do not delete this rule.

Add Multipathing Claim Rules

Use the `esxcli` commands to add a multipathing PSA claim rule to the set of claim rules on the system. For the new claim rule to be active, you first define the rule and then load it into your system.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Examples when you add a PSA claim rule include the following:

- You load a new third-party MPP and must define the paths that this module claims.
- You must enable the native HPP.

Warning: You cannot create rules where two different plug-ins claim paths to the same device. Your attempts to create these claim rules fail with a warning in `vmkernel.log`.

1. To define a new claim rule, use the following command:

```
esxcli storage core claimrule add
```

The command takes the following options:

Option	Description
<code>-A --adapter=<adapter></code>	Adapter of the paths to use. Valid only if <code>--type</code> is <code>location</code> .
<code>-u --autoassign</code>	Adds a claim rule based on its characteristics. The rule number is not required.
<code>-C --channel=<channel></code>	Channel of the paths to use. Valid only if <code>--type</code> is <code>location</code> .
<code>-c --claimrule-class=<cl></code>	Claim rule class to use in this operation. You can specify MP (default), Filter, or VAAI. To configure hardware acceleration for a new array, add two claim rules, one for the VAAI filter and another for the VAAI plug-in. See Add Hardware Acceleration Claim Rules for detailed instructions.
<code>-d --device=<device_uid></code>	UID of the device. Valid only when <code>--type</code> is <code>device</code> .
<code>-D --driver=<driver></code>	Driver for the HBA of the paths to use. Valid only if <code>--type</code> is <code>driver</code> .
<code>-f --force</code>	Force claim rules to ignore validity checks and install the rule anyway.
<code>--force-reserved</code>	Override protection of reserved rule ID ranges. Reserved claim rules are the rules with an ID below 100. You can use them to reassign local devices to specific plug-ins, for example, the NVMe device to HPP.

Option	Description
<code>--if-unset=<str></code>	Run this command if this advanced user variable is not set to 1.
<code>-i --iqn=<iscsi_name></code>	iSCSI Qualified Name for the target. Valid only when <code>--type</code> is <code>target</code> .
<code>-L --lun=<lun_id></code>	LUN of the paths. Valid only if <code>--type</code> is <code>location</code> . LUN ID must not be higher than the value of the advanced configuration option <code>/Disk/MaxLUN</code> .
<code>-M --model=<model></code>	Model of the paths to use. Valid only if <code>--type</code> is <code>vendor</code> . Valid values are values of the Model string from the SCSI inquiry string. Run <code>vicfg-scsidevs <conn_options> -l</code> on each device to see model string values.
<code>-P --plugin=<plugin></code>	PSA plug-in to use. The values are <code>NMP</code> , <code>MASK_PATH</code> , or <code>HPP</code> . Third parties can also provide their own PSA plug-ins. Required.
<code>-r --rule=<rule_ID></code>	Rule ID to use. The rule ID indicates the order in which the claim rule is to be evaluated. User-defined claim rules are evaluated in numeric order starting with 101. You can run <code>esxcli storage core claimrule list</code> to determine which rule IDs are available.
<code>-T --target=<target></code>	Target of the paths to use. Valid only if <code>--type</code> is <code>location</code> .
<code>-R --transport=<transport></code>	Transport of the paths to use. Valid only if <code>--type</code> is <code>transport</code> . The following values are supported. <code>block</code> — block storage <code>fc</code> — Fibre Channel <code>iscsivendor</code> — iSCSI <code>iscsi</code> — not currently used <code>ide</code> — IDE storage <code>sas</code> — SAS storage <code>sata</code> — SATA storage <code>usb</code> — USB storage <code>parallel</code> — parallel <code>fcoe</code> — FCoE <code>unknown</code>
<code>-t --type=<type></code>	Type of matching to use for the operation. Valid values are the following. Required. <code>vendor</code> <code>location</code> <code>driver</code> <code>transport</code> <code>device</code> <code>target</code>
<code>-V --vendor=<vendor></code>	Vendor of the paths to use. Valid only if <code>--type</code> is <code>vendor</code> . Valid values are values of the vendor string from the SCSI inquiry string. Run <code>vicfg-scsidevs <conn_options> -l</code> on each device to see vendor string values.
<code>--wwnn=<wwnn></code>	World-Wide Node Number for the target.
<code>--wwpn=<wwpn></code>	World-Wide Port Number for the target.
<code>-a --xcopy-use-array-values</code>	Use the array reported values to construct the XCOPY command to be sent to the storage array. This applies to VAAI claim rules only.

Option	Description
-s --xcopy-use-multi-segs	Use multiple segments when issuing an XCOPY request. Valid only if <code>--xcopy-use-array-values</code> is specified.
-m --xcopy-max-transfer-size	Maximum data transfer size in MB when you use a transfer size different than array reported. Valid only if <code>--xcopy-use-array-values</code> is specified.
-k --xcopy-max-transfer-size-kib	Maximum transfer size in KiB for the XCOPY commands when you use a transfer size different than array reported. Valid only if <code>--xcopy-use-array-values</code> is specified.

2. To load the new claim rule into your system, use the following command:

```
esxcli storage core claimrule load
```

This command loads all newly created multipathing claim rules from the `esx.conf` configuration file into the VMkernel. The command has no options.

3. To apply claim rules that are loaded, use the following command:

```
esxcli storage core claimrule run
```

The command takes the following options:

Option	Description
-A --adapter=<adapter>	If <code>--type</code> is <code>location</code> , name of the HBA for the paths to run the claim rules on. To run claim rules on paths from all adapters, omit this option.
-C --channel=<channel>	If <code>--type</code> is <code>location</code> , indicate the channel of the paths to use in this operation. To run claim rules on paths with any channel number, omit this option.
-c --claimrule-class=<cl>	Claim rule class to use in this operation.
-d --device=<device_uid>	UID of the device.
-L --lun=<lun_id>	If <code>--type</code> is <code>location</code> , indicate the LUN of the paths to run claim rules on. To run claim rules on paths with any LUN, omit this option.
-p --path=<path_uid>	If <code>--type</code> is <code>path</code> , this option indicates the unique path identifier (UID) or the runtime name of a path to run claim rules on.
-T --target=<target>	If <code>--type</code> is <code>location</code> , indicate the target of the paths to run claim rules on. To run claim rules on paths with any target number, omit this option.
-t --type=<location path all>	Type of claim to perform. By default, uses <code>all</code> , which means claim rules run without restriction to specific paths. Valid values are <code>location</code> , <code>path</code> , and <code>all</code> .
-w --wait	You can use this option only if you also use <code>--type all</code>. If the option is included, the claim waits for paths to settle before running the claim operation. In that case, the system does not start the claiming process until it is likely that all paths on the system have appeared before starting the claim process. After the claiming process has started, the command does not return until device registration has completed. If you add or remove paths during the claiming or the discovery process, this option might not work correctly.

Defining Multipathing Claim Rules

In the following example, you add and load rule # 500. The rule claims all paths with the NewMod model string and the NewVend vendor string for the NMP plug-in.

```
# esxcli storage core claimrule add -r 500 -t vendor -V NewVend -M NewMod -P NMP
```

```
# esxcli storage core claimrule load
```

After you run the `esxcli storage core claimrule list` command, you can see the new claim rule appearing on the list.

The following output indicates that the claim rule 500 has been loaded into the system and is active.

Rule	Class	Rule	Class	Type	Plugin	Matches
...	...	...	...	...	...	...
MP		500	runtime	vendor	NMP	vendor=NewVend model=NewMod
MP		500	file	vendor	NMP	vendor=NewVend model=NewMod

Delete Multipathing Claim Rules

Use the `esxcli` commands to remove a multipathing PSA claim rule from the set of claim rules on the system.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

1. Delete a claim rule from the set of claim rules.

```
esxcli storage core claimrule remove
```

Note: By default, the PSA claim rule 101 masks Dell array pseudo devices. Do not delete this rule, unless you want to unmask these devices.

The command takes the following options:

Option	Description
<code>-c --claimrule-class=<str></code>	Indicate the claim rule class (MP, Filter, VAAI).
<code>-P --plugin=<str></code>	Indicate the plug-in.
<code>-r --rule=<long></code>	Indicate the rule ID.

This step removes the claim rule from the File class.

2. Remove the claim rule from the system.

```
esxcli storage core claimrule load
```

This step removes the claim rule from the Runtime class.

Mask Paths

You can prevent the host from accessing storage devices or LUNs or from using individual paths to a LUN. Use the `esxcli` commands to mask the paths. When you mask paths, you create claim rules that assign the MASK_PATH plug-in to the specified paths.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

1. Check what the next available rule ID is.

```
esxcli storage core claimrule list
```

The claim rules that you use to mask paths have rule IDs in the range from 101 through 200. If this command shows that rules 101 and 102 exist, you can specify 103 for the rule to add.

2. Assign the MASK_PATH plug-in to a path by creating a new claim rule for the plug-in.

```
esxcli storage core claimrule add -P MASK_PATH
```

3. Load the MASK_PATH claim rule into your system.

```
esxcli storage core claimrule load
```

4. Verify that the MASK_PATH claim rule was added correctly.

```
esxcli storage core claimrule list
```

5. If a claim rule for the masked path exists, remove the rule.

```
esxcli storage core claiming unclaim
```

6. Run the path claiming rules.

```
esxcli storage core claimrule run
```

After you assign the MASK_PATH plug-in to a path, the path state becomes irrelevant and is no longer maintained by the host. As a result, commands that display the masked path's information might show the path state as dead.

Masking a LUN

In this example, you mask the LUN 20 on targets T1 and T2 accessed through storage adapters vmhba2 and vmhba3.

1.

```
#esxcli storage core claimrule list
```
2.

```
#esxcli storage core claimrule add -P MASK_PATH -r 109 -t location -A vmhba2 -C 0 -T 1 -L 20
#esxcli storage core claimrule add -P MASK_PATH -r 110 -t location -A vmhba3 -C 0 -T 1 -L 20
#esxcli storage core claimrule add -P MASK_PATH -r 111 -t location -A vmhba2 -C 0 -T 2 -L 20
#esxcli storage core claimrule add -P MASK_PATH -r 112 -t location -A vmhba3 -C 0 -T 2 -L 20
```
3.

```
#esxcli storage core claimrule load
```
4.

```
#esxcli storage core claimrule list
```
5.

```
#esxcli storage core claiming unclaim -t location -A vmhba2
#esxcli storage core claiming unclaim -t location -A vmhba3
```
6.

```
#esxcli storage core claimrule run
```

Unmask Paths

When you need the host to access the masked storage device, unmask the paths to the device.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Note: When you run an unclaim operation using a device property, for example, device ID or vendor, the paths claimed by the MASK_PATH plug-in are not unclaimed. The MASK_PATH plug-in does not track any device property of the paths that it claims.

1. Delete the MASK_PATH claim rule.

```
esxcli storage core claimrule remove -r rule#
```

2. Verify that the claim rule was deleted correctly.

```
esxcli storage core claimrule list
```

3. Reload the path claiming rules from the configuration file into the VMkernel.

```
esxcli storage core claimrule load
```

4. Run the `esxcli storage core claiming unclaim` command for each path to the masked storage device.

For example:

```
esxcli storage core claiming unclaim -t location -A vmhba0 -C 0 -T 0 -L 149
```

5. Run the path claiming rules.

```
esxcli storage core claimrule run
```

Your host can now access the previously masked storage device.

Define NMP SATP Rules

The NMP SATP claim rules define which SATP manages a storage device. Usually, you can use the default SATPs provided for storage devices. If default settings are not sufficient, use the `esxcli` commands to change the SATP for a specific device.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

You might need to create an SATP rule when you install a third-party SATP for a specific storage array.

1. To add a claim rule for a specific SATP, run the `esxcli storage nmp satp rule add` command. The command takes the following options.

Option	Description
<code>-b --boot</code>	This rule is a system default rule added at boot time. Do not modify <code>esx.conf</code> or add to a host profile.
<code>-c --claim-option=string</code>	Set the claim option string when adding a SATP claim rule.
<code>-e --description=string</code>	Set the claim rule description when adding a SATP claim rule.
<code>-d --device=string</code>	Set the device when adding SATP claim rules. Device rules are mutually exclusive with vendor/model and driver rules.
<code>-D --driver=string</code>	Set the driver string when adding a SATP claim rule. Driver rules are mutually exclusive with vendor/model rules.
<code>-f --force</code>	Force claim rules to ignore validity checks and install the rule anyway.
<code>-h --help</code>	Show the help message.
<code>-M --model=string</code>	Set the model string when adding SATP a claim rule. Vendor/Model rules are mutually exclusive with driver rules.
<code>-o --option=string</code>	Set the option string when adding a SATP claim rule.
<code>-P --psp=string</code>	Set the default PSP for the SATP claim rule.
<code>-O --psp-option=string</code>	Set the PSP options for the SATP claim rule.
<code>-s --satp=string</code>	The SATP for which a new rule is added.
<code>-R --transport=string</code>	Set the claim transport type string when adding a SATP claim rule.
<code>-t --type=string</code>	Set the claim type when adding a SATP claim rule.
<code>-V --vendor=string</code>	Set the vendor string when adding SATP claim rules. Vendor/Model rules are mutually exclusive with driver rules.

Note: When searching the SATP rules to locate a SATP for a given device, the NMP searches the driver rules first. If there is no match, the vendor/model rules are searched, and finally the transport rules. If there is still no match, NMP selects a default SATP for the device.

2. Reboot your host.

Defining an NMP SATP Rule

The following sample command assigns the `VMW_SATP_INV` plug-in to manage storage arrays with vendor string `NewVend` and model string `NewMod`.

```
# esxcli storage nmp satp rule add -V NewVend -M NewMod -s VMW_SATP_INV
```

When you run the `esxcli storage nmp satp list -s VMW_SATP_INV` command, you can see the new rule on the list of `VMW_SATP_INV` rules.

Scheduling Queues for Virtual Machine I/Os in vSphere Environment

By default, vSphere provides a mechanism that creates scheduling queues for every virtual machine file. Each file, for example `.vmdk`, gets its own bandwidth controls. You can use the vSphere Client or the `esxcli` commands to turn off the mechanism.

This mechanism ensures that I/O for a particular virtual machine file goes into its own separate queue and avoids interfering with I/Os from other files.

Edit Per File I/O Scheduling in the vSphere Client

The advanced `VMkernel.Boot.isPerFileSchedModelActive` parameter controls the per file I/O scheduling mechanism on VMFS and NFS 3 datastores. On the ESXi host, the mechanism is turned on by default. You can turn off the mechanism using the **Advanced System Settings** dialog box.

If you turn off the per file I/O scheduling model, your host reverts to a legacy scheduling mechanism. The legacy scheduling maintains only one I/O queue for each virtual machine and storage device pair. All I/Os between the virtual machine and its virtual disks are moved into this queue. As a result, I/Os from different virtual disks might interfere with each other in sharing the bandwidth and affect each other's performance.

Note: Do not deactivate per file scheduling if you have the HPP plug-in and the latency sensitive threshold parameter configured for high-speed local devices. Turning off per file scheduling might cause unpredictable behavior.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. Edit the value of the `VMkernel.Boot.isPerFileSchedModelActive` parameter.

Option	Description
True	Turn on the per file scheduling mechanism. By default, the parameter is activated.
False	Deactivate the per file scheduling mechanism.

5. Reboot the host for the changes to take effect.

Use esxcli Commands to Edit Per File I/O Scheduling

You can use the `esxcli` commands to change the I/O scheduling capability for VMFS, NFS 3, and NFS 4.1 datastores on your ESXi host. The capability is turned on by default.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

To set or edit per file I/O scheduling, run the following commands:

Option	Description
<code>esxcli system settings kernel set -s isPerFileSchedModelActive -v FALSE</code>	Turn off per file I/O scheduling for VMFS and NFS 3.
<code>esxcli system settings kernel set -s isPerFileSchedModelActive -v TRUE</code>	Turn on per file I/O scheduling for VMFS and NFS 3.
<code>esxcli system module parameters list -m nfs41client</code>	Display the current status of NFS 4.1 file based scheduler.
<code>esxcli system module parameters set -m nfs41client -p fileBasedScheduler=0</code>	Turn off file based scheduler for NFS 4.1.
<code>esxcli system module parameters set -m nfs41client -p fileBasedScheduler=1</code>	Turn on file based scheduler for NFS 4.1.

Raw Device Mapping in vSphere

Raw device mapping (RDM) in vSphere provides a mechanism for a virtual machine to have direct access to a LUN on the physical storage subsystem.


The following topics contain information about RDMs in vSphere and provide instructions on how to create and manage RDMs.

Introduction to Raw Device Mapping in vSphere

An RDM is a mapping file in a separate vSphere VMFS volume that acts as a proxy for a raw physical storage device. With RDM, a virtual machine can access and use the storage device directly. The RDM contains metadata for managing and redirecting disk access to the physical device.

The file gives you some of the advantages of direct access to a physical device, but keeps some advantages of a virtual disk in VMFS. As a result, it merges the VMFS manageability with the raw device access.

Figure 107: Raw Device Mapping



Typically, you use VMFS datastores for most virtual disk storage. On certain occasions, you might use raw LUNs or logical disks located in a SAN.

For example, you might use raw LUNs with RDMs in the following situations:

- When SAN snapshot or other layered applications run in the virtual machine. The RDM enables backup offloading systems by using features inherent to the SAN.
- In any MSCS clustering scenario that spans physical hosts, such as virtual-to-virtual clusters and physical-to-virtual clusters. In this case, cluster data and quorum disks are configured as RDMs rather than as virtual disks on a shared VMFS.

Think of an RDM as a symbolic link from a VMFS volume to a raw LUN. The mapping makes LUNs appear as files in a VMFS volume. The RDM, not the raw LUN, is referenced in the virtual machine configuration. The RDM contains a reference to the raw LUN.

Two compatibility modes are available for RDMs:

- In the virtual compatibility mode, the RDM acts like a virtual disk file. The RDM can use snapshots.
- In the physical compatibility mode, the RDM offers direct access to the SCSI device for those applications that require lower-level control.

Benefits of Raw Device Mapping

An RDM provides a number of benefits, but it should not be used in every situation. In general, virtual disk files are preferable to RDMs for manageability. However, when you need raw devices, you must use the RDM.

RDM offers several benefits.

User-Friendly Persistent Names

Provides a user-friendly name for a mapped device. When you use an RDM, you do not need to refer to the device by its device name. You refer to it by the name of the mapping file, for example:

```
/vmfs/volumes/myVolume/myVMDirectory/myRawDisk.vmdk
```

Dynamic Name Resolution

Stores unique identification information for each mapped device. VMFS associates each RDM with its current SCSI device, regardless of changes in the physical configuration of the server because of adapter hardware changes, path changes, device relocation, and so on.

Distributed File Locking

Makes it possible to use VMFS distributed locking for raw SCSI devices. Distributed locking on an RDM makes it safe to use a shared raw LUN without losing data when two virtual machines on different servers try to access the same LUN.

File Permissions

Makes file permissions possible. The permissions of the mapping file are enforced at file-open time to protect the mapped volume.

File System Operations

Makes it possible to use file system utilities to work with a mapped volume, using the mapping file as a proxy. Most operations that are valid for an ordinary file can be applied to the mapping file and are redirected to operate on the mapped device.


Snapshots

Makes it possible to use virtual machine snapshots on a mapped volume. Snapshots are not available when the RDM is used in physical compatibility mode.

vMotion

Lets you migrate a virtual machine with vMotion. The mapping file acts as a proxy to allow vCenter Server to migrate the virtual machine by using the same mechanism that exists for migrating virtual disk files.

Figure 108: vMotion of a Virtual Machine Using Raw Device Mapping



SAN Management Agents

Makes it possible to run some SAN management agents inside a virtual machine. Similarly, any software that needs to access a device by using hardware-specific SCSI commands can be run in a virtual machine. This kind of software is called SCSI target-based software. When you use SAN management agents, select a physical compatibility mode for the RDM.

N-Port ID Virtualization (NPIV)

Makes it possible to use the NPIV technology that allows a single Fibre Channel HBA port to register with the Fibre Channel fabric using several worldwide port names (WWPNs). This ability makes the HBA port appear as multiple virtual ports, each having its own ID and virtual port name. Virtual machines can then claim each of these virtual ports and use them for all RDM traffic.

Note: You can use NPIV only for virtual machines with RDM disks.

VMware works with vendors of storage management software to ensure that their software functions correctly in environments that include ESXi. Some applications of this kind are:

- SAN management software
- Storage resource management (SRM) software
- Snapshot software
- Replication software

Such software uses a physical compatibility mode for RDMs so that the software can access SCSI devices directly.

Various management products are best run centrally (not on the ESXi machine), while others run well on the virtual machines. VMware does not certify these applications or provide a compatibility matrix. To find out whether a SAN management application is supported in an ESXi environment, contact the SAN management software provider.

RDM Considerations and Limitations

Certain considerations and limitations exist when you use RDMs.

- The RDM is not available for direct-attached block devices or certain RAID devices. The RDM uses a SCSI serial number to identify the mapped device. Because block devices and some direct-attach RAID devices do not export serial numbers, they cannot be used with RDMs.
- If you are using the RDM in physical compatibility mode, you cannot use a snapshot with the disk. Physical compatibility mode allows the virtual machine to manage its own, storage-based, snapshot or mirroring operations. Virtual machine snapshots are available for RDMs with virtual compatibility mode.
- You cannot map to a disk partition. RDMs require the mapped device to be a whole LUN.
- If you use vMotion to migrate virtual machines with RDMs, make sure to maintain consistent LUN IDs for RDMs across all participating ESXi hosts.

Characteristics of Raw Device Mapping in vSphere

An RDM is a special mapping file in a vSphere VMFS volume that manages metadata for its mapped device. The mapping file is presented to the management software as an ordinary disk file, available for the usual file-system operations. To the virtual machine, the storage virtualization layer presents the mapped device as a virtual SCSI device.

Key contents of the metadata in the mapping file include the location of the mapped device (name resolution), the locking state of the mapped device, permissions, and so on.

RDM Virtual and Physical Compatibility Modes

You can use RDMs in virtual compatibility or physical compatibility modes. Virtual mode specifies full virtualization of the mapped device. Physical mode specifies minimal SCSI virtualization of the mapped device, allowing the greatest flexibility for SAN management software.

In virtual mode, the VMkernel sends only READ and WRITE to the mapped device. The mapped device appears to the guest operating system exactly the same as a virtual disk file in a VMFS volume. The real hardware characteristics are hidden. If you are using a raw disk in virtual mode, you can realize the benefits of VMFS such as advanced file locking for data protection and snapshots for streamlining development processes. Virtual mode is also more portable across storage hardware than physical mode, presenting the same behavior as a virtual disk file.

In physical mode, the VMkernel passes all SCSI commands to the device, with one exception: the REPORT LUNs command is virtualized so that the VMkernel can isolate the LUN to the owning virtual machine. Otherwise, all physical characteristics of the underlying hardware are exposed. Physical mode is useful to run SAN management agents or other SCSI target-based software in the virtual machine. Physical mode also allows virtual-to-physical clustering for cost-effective high availability.

VMFS5 and VMFS6 support greater than 2 TB disk size for RDMs in virtual and physical modes.

Dynamic Name Resolution

The RDM file supports dynamic name resolution when a path to a raw device changes.


VMFS uniquely identifies all mapped storage devices, and the identification is stored in its internal data structures. Any change in the path to a raw device, such as a Fibre Channel switch failure or the addition of a new HBA, can change the

device name. Dynamic name resolution resolves these changes and automatically associates the original device with its new name.

Raw Device Mapping with Virtual Machine Clusters

Use an RDM with virtual machine clusters that require access to the same raw LUN for failover scenarios. The setup is similar to that of a virtual machine cluster that accesses the same virtual disk file, but an RDM replaces the virtual disk file.

Figure 109: Access from Clustered Virtual Machines



Comparing Available SCSI Device Access Modes

The ways of accessing a SCSI-based storage device include a virtual disk file on a VMFS datastore, virtual mode RDM, and physical mode RDM.

The following table provides a comparison of features available with the different modes.

Table 223: Features Available with Virtual Disks and Raw Device Mappings

ESXi Features	Virtual Disk File	Virtual Mode RDM	Physical Mode RDM
SCSI Commands Passed Through	No	No	Yes REPORT LUNs is not passed through
vCenter Server Support	Yes	Yes	Yes
Snapshots	Yes	Yes	No
Distributed Locking	Yes	Yes	Yes
Clustering	Cluster-in-a-box only	Cluster-in-a-box cluster-across-boxes	Physical-to-virtual clustering cluster-across-boxes
SCSI Target-Based Software	No	No	Yes

Use virtual disk files for the cluster-in-a-box type of clustering. If you plan to reconfigure your cluster-in-a-box clusters as cluster-across-boxes clusters, use virtual mode RDMs for the cluster-in-a-box clusters.

Create vSphere Virtual Machines with RDMs

When you give your vSphere virtual machine direct access to a raw SAN LUN, you create an RDM disk that resides on a VMFS datastore and points to the LUN. You can create the RDM as an initial disk for a new virtual machine or add it to an existing virtual machine. When creating the RDM, you specify the LUN to be mapped and the datastore on which to put the RDM.

Although the RDM disk file has the same `.vmdk` extension as a regular virtual disk file, the RDM contains only mapping information. The actual virtual disk data is stored directly on the LUN.

This procedure assumes that you are creating a new virtual machine. For information, see the *vSphere Virtual Machine Administration* documentation.

1. Create a virtual machine.
 - a) Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host, and select **New Virtual Machine**.
 - b) Select **Create a new virtual machine** and click **Next**.
 - c) Follow the steps required to create a virtual machine.
2. On the Customize Hardware page, click the **Virtual Hardware** tab.
3. Optional: To delete the default virtual hard disk that the system created for your virtual machine, move your cursor over the disk and click the **Remove** icon.
4. Add an RDM disk.
 - a) Click **Add New Devices** and select **RDM Disk** from the list.
 - b) From the list of LUNs, select a target raw LUN and click **OK**.
 The system creates an RDM disk that maps your virtual machine to the target LUN. The RDM disk is shown on the list of virtual devices as a new hard disk.
5. Configure the RDM disk.
 - a) Click the **New Hard Disk** triangle to expand the properties for the RDM disk.
 - b) Select a location for the RDM.
 You can place the RDM on the same datastore where your virtual machine configuration files reside, or select a different datastore.

Note: To use vMotion for virtual machines with enabled NPIV, make sure that the RDM files and the virtual machine files are located on the same datastore. You cannot perform Storage vMotion when NPIV is enabled.

 - c) Select a compatibility mode.

Option	Description
Physical	Allows the guest operating system to access the hardware directly. Physical compatibility is useful if you are using SAN-aware applications on the virtual machine. However, a virtual machine with a physical compatibility RDM cannot be cloned, made into a template, or migrated if the migration involves copying the disk.

Option	Description
Virtual	Allows the RDM to behave as if it were a virtual disk, so you can use such features as taking snapshots, cloning, and so on. When you clone the disk or make a template out of it, the contents of the LUN are copied into a <code>.vmdk</code> virtual disk file. When you migrate a virtual compatibility mode RDM, you can migrate the mapping file or copy the contents of the LUN into a virtual disk.

d) If you selected virtual compatibility mode, select a disk mode.

Disk modes are not available for RDM disks using physical compatibility mode.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent - Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.
Independent - Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset.

6. Complete your virtual machine configuration.

Deactivate SCSI INQUIRY Cache for Improved Performance in vSphere

Certain virtual machines with RDMs must obtain the SCSI INQUIRY information from the LUN instead of using SCSI INQUIRY data cached by ESXi.

Certain guest operating systems or applications that run in the virtual machines with the RDMs display unpredictable behavior.

This behavior might be caused by cached SCSI INQUIRY data that interferes with specific guest operating systems and applications.

When the ESXi host first connects to a target storage device, it issues the SCSI INQUIRY command to obtain basic identification data from the device. By default, ESXi caches the received SCSI INQUIRY data (Standard, page 80, and page 83), and the data remains unchanged afterwards. Responses for subsequent SCSI INQUIRY commands are returned from the cache.

However, specific guest operating systems running in virtual machines with RDMs must query the LUN instead of using SCSI INQUIRY data cached by ESXi. In these cases, you can configure the VM to ignore the SCSI INQUIRY cache.

Use one of the following methods.

Option	Description
Modify the <code>.vmx</code> file of the virtual machine with the RDM	Use this method for the VMs with hardware version 8 or later. - Add the following parameter to the file: <pre>scsix:y.ignoreDeviceInquiryCache = "true"</pre> where <code>x</code> is the SCSI controller number and <code>y</code> is the SCSI target number of the RDM. - Reboot the VM.
Use the <code>esxcli</code> command	Because you configure the setting at a host level, no VM hardware version limitations apply.

Option	Description
	<pre>esxcli storage core device inquirycache set --device device id --ignore true</pre> No VM reboot is required.

No matter which method you use to set the SCSI INQUIRY cache parameter to true, the VM starts contacting the LUN directly for the SCSI INQUIRY data.


ignoreDeviceInquiryCache parameter in vmx	ignore inquirycache parameter in esxcli	Inquiry request serviced from
True	True	LUN
False (default if the parameter is not present)	True	LUN
True	False	LUN
False (default if the parameter is not present)	False	Cache

Storage Policy Based Management in vSphere

In vSphere, Storage Policy Based Management (SPBM) plays a major role by helping to align storage with application demands of your virtual machines. It provides a storage policy framework that serves as a single unified control panel across a broad range of data services and storage solutions.

As an abstraction layer, SPBM abstracts storage services delivered by Virtual Volumes, vSAN, I/O filters, or other storage entities.

Rather than integrating with each individual type of storage and data services, SPBM provides a universal framework for different types of storage entities.



SPBM offers the following mechanisms:

- Advertisement of storage capabilities and data services that storage arrays and other entities, such as I/O filters, offer.

- Bidirectional communications between ESXi and vCenter Server on one side, and storage arrays and entities on the other.
- Virtual machine provisioning based on VM storage policies.

About Virtual Machine Storage Policies

Virtual machine storage policies are essential to virtual machine provisioning through SPBM. The policies control which type of storage is provided for the virtual machine and how the virtual machine is placed within storage. They also determine data services that the virtual machine can use.

vSphere offers default storage policies. See [Default Storage Policies in vSphere](#).

In addition, you can define policies and assign them to the virtual machines. See [Workflow for vSphere VM Storage Policies](#).

You use the VM Storage Policies interface to create a storage policy. When you define the policy, you specify various storage requirements for applications that run on the virtual machines. You can also use storage policies to request specific data services, such as caching or replication, for virtual disks.

You apply the storage policy when you create, clone, or migrate the virtual machine. After you apply the storage policy, the SPBM mechanism assists you with placing the virtual machine in a matching datastore. In certain storage environments, SPBM determines how the virtual machine storage objects are provisioned and allocated within the storage resource to guarantee the required level of service. The SPBM also enables requested data services for the virtual machine and helps you to monitor policy compliance.

Default Storage Policies in vSphere

When you provision a virtual machine on a datastore in vSphere, you must assign to the virtual machine a compatible VM storage policy. If you do not configure and explicitly assign the storage policy to the virtual machine, the system uses a default storage policy.

VMware-Provided Default Storage Policy

The generic default storage policy that ESXi provides applies to all datastores and does not include rules specific to any storage type.

In addition, ESXi offers the default storage policies for object-based datastores, vSAN or Virtual Volumes. These policies guarantee the optimum placement for the virtual machine objects within the object-based storage.

For information about the default storage policy for Virtual Volumes, see [Virtual Volumes and VM Storage Policies](#).

VMFS and NFS datastores do not have specific default policies and can use the generic default policy or a custom policy you define for them.

User-Defined Default Storage Policies

You can create a VM storage policy that is compatible with vSAN or Virtual Volumes. You can then designate this policy as the default for vSAN and Virtual Volumes datastores. The user-defined default policy replaces the default storage policy that VMware provides.

Each vSAN and Virtual Volumes datastore can have only one default policy at a time. However, you can create a single storage policy with multiple placement rule sets, so that it matches multiple vSAN and Virtual Volumes datastores. You can designate this policy as the default policy for all datastores.

When the VM storage policy becomes the default policy for a datastore, you cannot delete the policy unless you disassociate it from the datastore.

Change the Default Storage Policy for a Datastore

For Virtual Volumes and vSAN datastores, VMware provides storage policies that are used as the default during the virtual machine provisioning. You can change the default storage policy for a selected Virtual Volumes or vSAN datastore.

Create a storage policy that is compatible with Virtual Volumes or vSAN. You can create a policy that matches both types of storage.

Note: Do not designate a storage policy with replication rules as a default storage policy. Otherwise, the policy prevents you from selecting replication groups.

1. In the vSphere Client, navigate to the datastore.
2. Click the **Configure** tab, and click **General**.
3. In the Default Storage Policy pane, click **Edit**.
4. From the list of available storage policies, select a policy to designate as the default and click **OK**.

The selected storage policy becomes the default policy for the datastore. The system assigns this policy to any virtual machine objects that you provision on the datastore when no other policy is selected.

Workflow for vSphere VM Storage Policies

The entire process of creating and managing storage policies in vSphere typically includes several steps. Whether you must perform a specific step might depend on the type of storage or data services that your environment offers.

Step	Description
Populate the VM Storage Policies interface with appropriate data.	The VM Storage Policies interface is populated with information about datastores and data services that are available in your storage environment. This information is obtained from storage providers and datastore tags. - For entities represented by storage providers, verify that an appropriate provider is registered. Entities that use the storage provider include vSAN, Virtual Volumes, and I/O filters. Depending on the type of storage entity, some providers are self-registered. Other providers must be manually registered. See Use Storage Providers to Populate the VM Storage Policies Interface and Register Storage Providers for Virtual Volumes. - Tag datastores that are not represented by storage providers. You can also use tags to indicate a property that is not communicated through the storage provider, such as geographical location or administrative group. See Assign Tags to Datastores.
Create predefined storage policy components.	A storage policy component describes a single data service, such as replication, that must be provided for the virtual machine. You can define the component in advance and associate it with multiple VM storage policies. The components are reusable and interchangeable. See Create Storage Policy Components.
Create VM storage policies.	When you define storage policies for virtual machines, you specify storage requirements for applications that run on the virtual machines. See Creating and Managing vSphere VM Storage Policies.

Step	Description
Apply the VM storage policy to the virtual machine.	You can apply the storage policy when deploying the virtual machine or configuring its virtual disks. See Assign Storage Policies to Virtual Machines .
Check compliance for the VM storage policy.	Verify that the virtual machine uses the datastore that is compliant with the assigned storage policy. See Check Compliance for a VM Storage Policy .

To create and manage your storage policies, you use the VM Storage Policy interface of the vSphere Client.

Populating the VM Storage Policies Interface in vSphere

Before you start creating VM storage policies, you must populate the VM Storage Policy interface with information about storage entities and data services that are available in your vSphere storage environment.

This information is obtained from storage providers, also called VASA providers. Another source is datastore tags.

Storage Capabilities and Services

Certain datastores, for example, Virtual Volumes and vSAN, are represented by the storage providers. Through the storage providers, the datastores can advertise their capabilities in the VM Storage Policy interface. These datastore capabilities, data services, and other characteristics with ranges of values populate the VM Storage Policy interface.

You use these characteristics when you define datastore-based placement and service rules for your storage policy.

Data Services

I/O filters on your hosts are also represented by the storage providers. The storage provider delivers information about the data services of the filters to the VM Storage Policy interface. You use this information when defining the rules for host-based data services, also called common rules. Unlike the datastore-specific rules, these rules do not define storage placement and storage requirements for the virtual machine. Instead, they activate the requested I/O filter data services for the virtual machine.

Tags

Generally, VMFS and NFS datastores are not represented by a storage provider. They do not display their capabilities and data services in the VM Storage Policies interface. You can use tags to encode information about these datastores. For example, you can tag your VMFS datastores as VMFS-Gold and VMFS-Silver to represent different levels of service.

For Virtual Volumes and vSAN datastores, you can use tags to encode information that is not advertised by the storage provider, such as geographical location (Palo Alto), or administrative group (Accounting).

Similar to the storage capabilities and characteristics, all tags associated with the datastores appear in the VM Storage Policies interface. You can use the tags when you define the tag-based placement rules.

Use Storage Providers to Populate the VM Storage Policies Interface

For entities represented by storage (VASA) providers, verify that an appropriate provider is registered. After the storage providers are registered, the VM Storage Policies interface becomes populated with information about datastores and data services that the providers represent.

Register the storage providers that require manual registration. For more information, see the appropriate documentation:

- [Administering VMware vSAN](#)
- [Working with VMware vSphere Virtual Volumes](#)
- [Filtering Virtual Machine I/O in vSphere](#)

Entities that use the storage provider include vSAN, Virtual Volumes, and I/O filters. Depending on the type of the entity, some providers are self-registered. Other providers, for example, the Virtual Volumes storage provider, must be manually registered. After the storage providers are registered, they deliver the following data to the VM Storage Policies interface:

- Storage capabilities and characteristics for such datastores as Virtual Volumes and vSAN.
- Data services the I/O filters provide.

1. Browse to the vCenter Server instance.
2. Click the **Configure** tab, and click **Storage Providers**.
3. In the Storage Providers list, view the storage providers registered with vCenter Server.
The list shows general information including the name of the storage provider, its URL and status, storage entities that the provider represents, and so on.
4. To display more details, select a specific storage provider or its component from the list.

Assign Tags to Datastores

Use tags to encode information about a datastore. The tags are helpful when your datastore is not represented by a storage provider and does not advertise its services in the VM Storage Policies interface. You can also use the tags to indicate a property that is not communicated through a storage provider, such as a geographical location or administrative group.

Required privileges:

- **vSphere Tagging > Create vSphere Tag** on the root vCenter Server instance
- **vSphere Tagging > Create vSphere Tag Category** on the root vCenter Server instance
- **vSphere Tagging > Assign or Unassign vSphere Tag** on the root vCenter Server instance

You can apply a new tag that contains general storage information to a datastore. For more details about the tags, their categories, and how to manage the tags, see the *vCenter Server and Host Management* documentation.

1. In the vSphere Client, create a category for storage tags.
 - a) From the Home menu, click **Tags & Custom Attributes**.
 - b) Click the **Tags** tab and click **Categories**.
 - c) Click the **Add Category** icon.
 - d) Specify the category properties. See the following example.

Category Property	Example
Category Name	Storage Location
Description	Category for tags related to location of storage
Tags Per Object	Many tags
Associable Object Types	Datastore and Datastore Cluster

- e) Click **OK**.
2. Create a storage tag.
 - a) On the **Tags** tab, click **Tags**.
 - b) Click the **Add Tag** icon.
 - c) Specify the properties for the tag. See the following example.

Tag Property	Example
Name	Texas
Description	Datastore located in Texas
Category	Storage Location

- d) Click **OK**.
3. Apply the tag to the datastore.
 - a) Navigate to the datastore.
 - b) Right-click the datastore, and select **Tags & Custom Attributes > Assign Tag**.
 - c) From the list of tags, select an appropriate tag, for example, Texas in the Storage Location category, and click **Assign**.

The new tag is assigned to the datastore and appears on the datastore **Summary** tab in the **Tags** pane.

When creating a VM storage policy, you can reference the tag to include the tagged datastore in the list of compatible storage resources. See [Create a VM Storage Policy for Tag-Based Placement](#).

Or you can exclude the tagged datastore from the VM storage policy. For example, your VM storage policy can include Virtual Volumes datastores located in Texas and California, but exclude datastores located in Nevada.

To learn more about how to use tags in VM storage policies, watch the following video.

About Rules and Rule Sets in vSphere Storage Policies

After the VM Storage Policies interface is populated with the appropriate data, you can start creating your vSphere storage policies. Creating a policy involves defining specific storage placement rules and rules to configure data services.

Rules

The rule is a basic element of the VM storage policy. Each individual rule is a statement that describes a single requirement for virtual machine storage and data services.

Rule Sets

Within a storage policy, individual rules are organized into collections of rules, or rule sets. Typically, the rule sets can be in one of the following categories: rules for host-based services and datastore-specific rules.

Datastore-Specific Rule Sets

Each rule set must include placement rules that describe requirements for virtual machine storage resources. All placement rules within a single rule set represent a single storage entity. These rules can be based on storage capabilities or tags. In addition, the datastore-specific rule set can include optional rules or storage policy components that describe data services to provide for the virtual machine. Generally, these rules request such services as caching, replication, other services provided by storage systems.

To define the storage policy, one datastore-specific set is required. Additional rule sets are optional. A single policy can use multiple sets of rules to define alternative storage placement parameters, often from several storage providers.

Placement Rules: Capability-Based

Placement rules specify a particular storage requirement for the VM and enable SPBM to distinguish compatible datastores among all datastores in the inventory. These rules also describe how the virtual machine storage objects are allocated within the datastore to receive the required level of service. For example, the rules can list Virtual Volumes as a destination and define the maximum recovery point objective (RPO) for the Virtual Volumes objects.

When you provision the virtual machine, these rules guide the decision that SPBM makes about the virtual machine placement. SPBM finds the Virtual Volumes datastores that can match the rules and satisfy the storage requirements of the virtual machine. See [Create a VM Storage Policy for Virtual Volumes](#).

Placement Rules: Tag-Based

Tag-based rules reference datastore tags. These rules can define the VM placement, for example, request as a target all datastores with the VMFS-Gold tag. You can also use the tag-based rules to fine-tune your VM placement request further. For example, exclude datastores with the Palo Alto tag from the list of your Virtual Volumes datastores. See [Create a VM Storage Policy for Tag-Based Placement](#).

Rules for Host-Based Services

This rule set activates data services provided by the host. The set for host-based services can include rules or storage policy components that describe particular data services, such as encryption or replication.

Unlike datastore-specific rules, this set does not include placement rules. Rules for host-based services are generic for all types of storage and do not depend on the datastore. See [Create a VM Storage Policy for Host-Based Data Services](#).


Table 224: Structure of a VM Storage Policy

Rules for Host-Based Services	Datastore-Specific Rule Sets
Rules or predefined storage policy components to activate data services installed on ESXi hosts. For example, replication by I/O filters.	Capability-based or tag-based placement rules that describe requirements for virtual machine storage resources. For example, Virtual Volumes placement.
	Rules or predefined storage policy components that activate data services provided by storage. For example, caching by Virtual Volumes.

Relationships Between Rules and Rule Sets

The boolean operator **OR** defines the relationship between the datastore-specific rule sets within the policy. The **AND** operator defines the relationship between all rules within a single rule set. The policy can contain only a rule set for host-based services, or only a datastore-specific rule set, or both.

If the rule set for host-based services is not present, meeting all the rules of a single datastore-specific rule set is sufficient to satisfy the entire policy. If the rule set for host-based services is present, the policy matches the datastore that satisfies the host services rules and all rules in one of the datastore-specific sets.



Creating and Managing vSphere VM Storage Policies

To create and manage storage policies for your virtual machines in vSphere, you use the VM Storage Policies interface.

Create a VM Storage Policy for VMFS Datastore

Create a custom storage policy compatible with VMFS datastores. In the storage policy, you can specify disk or volume provisioning options.

- Be familiar with provisioning policies. See [About Virtual Disk Provisioning Policies](#).
- Required privileges: **VM storage policies > Update** and **VM storage policies > View**.

1. Open the **Create VM Storage Policy** wizard.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **VM Storage Policies**.
 - c) Click **Create**.

2. Enter the policy name and description, and click **Next**.

Option	Action
vCenter Server	Select the vCenter Server instance.
Name	Enter the name of the storage policy, for example VMFS Storage Policy.
Description	Enter the description of the storage policy.

3. On the **Policy structure** page under **Datastore specific rules**, enable rules for a target storage entity, VMFS storage. You can enable rules for several datastores. Multiple rule sets allow a single policy to define alternative storage placement parameters, often from several storage providers.

4. On the **VMFS** rules page, define storage placement rules for the target VMFS datastore.

- a) Click the **Placement** tab.
- b) From the **Volume allocation** drop-down menu, select one of the following options.

Conserve space when possible	Use to create a virtual disk of volume in thin format. The option neither guarantees space reservation nor runtime performance. This option gives fast provisioning time but the space is allocated on demand and the blocks are also initialized on demand.
Fully initialized	Use to create a disk or volume in thick provision eager zeroed format. The option guarantees that the required space is reserved on the datastore and all the blocks of the disk are zeroed. This option gives the best runtime performance, but provisioning time is proportional to the size of the disk.
Reserve space	Use to create a disk or volume in thick provision lazy zeroed format. The option guarantees that the required space is reserved on the datastore. The blocks of the disk are not zeroed, so there are no guarantee on the runtime performance.

- c) To fine-tune your placement request further, click the **Tags** tab and add a tag-based rule.

Tag-based rules can filter datastores by including or excluding specific placement criteria. For example, your VM storage policy can include VMFS datastores located in Texas and California, but exclude datastores located in Nevada.

Create a VM Storage Policy for Host-Based Data Services

To define the VM storage policy in the vSphere Client, use the **Create VM Storage Policy** wizard. In this task, you create rules for data services offered by ESXi hosts. The VM storage policy that includes these rules activates specified data services for the virtual machine.

- For information about encrypting your virtual machines, see the *vSphere Security* documentation.
- For information about I/O filters, see [Filtering Virtual Machine I/O in vSphere](#).
- For information about storage policy components, see [About vSphere Storage Policy Components](#).
- Required privileges: **VM storage policies > Update** and **VM storage policies > View**.

Available data services include encryption, I/O control, caching, and so on. Certain data services, such as encryption, are provided by VMware. Others can be offered by third-party I/O filters that you install on your host.

The data services are usually generic for all types of storage and do not depend on a datastore. Adding datastore-specific rules to the storage policy is optional.

If you add datastore-specific rules, and both the I/O filters on the host and storage offer the same type of service, for example, encryption, your policy can request this service from both providers. As a result, the virtual machine data is encrypted twice, by the I/O filter and your storage. However, replication provided by Virtual Volumes and replication provided by the I/O filter cannot coexist in the same storage policy.

- Open the **Create VM Storage Policy** wizard.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **VM Storage Policies**.
 - c) Click **Create**.
- Enter the policy name and description, and click **Next**.

Option	Action
vCenter Server	Select the vCenter Server instance.

Option	Action
Name	Enter the name of the storage policy.
Description	Enter the description of the storage policy.

3. On the **Policy structure** page under **Host based services**, click **Enable host based rules**.
4. On the **Host based services** page, define rules and configure data services provided by your host.
 - a) Click the tab for the data service category, for example, **Encryption**.
 - b) Define custom rules for the data service category or use predefined components.

Option	Description
Disabled	This is a default setting.
Use storage policy component	Select a storage policy component from the drop-down menu. This option is available only if you have predefined components in your database.
Custom	Define custom rules for the data service category by specifying an appropriate provider and values for the rules.

Note: You can activate several data services. If you use encryption with other data services, use the custom encryption option and set the **Allow I/O filters before encryption** parameter to **True**. The parameter ensures that other services, such as replication, can analyze clear text data before it is encrypted.

5. On the **Storage compatibility** page, review the list of datastores that match this policy.
To be compatible with the policy for host-based services, datastores must be connected to the host that provides these services. If you add datastore-specific rule sets to the policy, the compatible datastores must also satisfy storage requirements of the policy.
6. On the **Review and finish** page, review the storage policy settings and click **Finish**.
To change any settings, click **Back** to go to the relevant page.

The new VM storage policy for host-based data services appears on the list.

Create a VM Storage Policy for Virtual Volumes

To define the VM storage policy in the vSphere Client, use the **Create VM Storage Policy** wizard. In this task, you create a custom storage policy compatible with Virtual Volumes. When you define the VM storage policy for Virtual Volumes, you create rules to configure storage and data services provided by the Virtual Volumes datastore. The rules are applied when the VM is placed on the Virtual Volumes datastore. The custom storage policy can replace the default No Requirements storage policy for Virtual Volumes that VMware provides.

- Verify that the Virtual Volumes storage provider is available and active. See [Register Storage Providers for Virtual Volumes](#).
- Make sure that the VM Storage Policies interface is populated with information about storage entities and data services that are available in your storage environment. See [Populating the VM Storage Policies Interface in vSphere](#).
- Define appropriate storage policy components. See [Create Storage Policy Components](#).

- Required privileges: **VM storage policies > Update** and **VM storage policies > View**.

The procedure assumes that you are creating the VM storage policy for Virtual Volumes. For information about the vSAN storage policy, see the *Administering VMware vSAN* documentation.

- Open the **Create VM Storage Policy** wizard.
 - Click **Menu > Policies and Profiles**.
 - Under **Policies and Profiles**, click **VM Storage Policies**.
 - Click **Create**.

- Enter the policy name and description, and click **Next**.

Option	Action
vCenter Server	Select the vCenter Server instance.
Name	Enter the name of the storage policy, for example Virtual Volumes Storage Policy.
Description	Enter the description of the storage policy.

- On the **Policy structure** page under Datastore specific rules, enable rules for a target storage entity, such as Virtual Volumes storage.

You can enable rules for several datastores. Multiple rule sets allow a single policy to define alternative storage placement parameters, often from several storage providers.

- On the *Virtual Volumes* rules page, define storage placement rules for the target Virtual Volumes datastore.

- Click the **Placement** tab and click **Add Rule**.

- From the Add Rule drop-down menu, select available capability and specify its value.

For example, you can specify the number of read operations per second for the Virtual Volumes objects.

You can include as many rules as you need for the selected storage entity. Verify that the values you provide are within the range of values that the Virtual Volumes datastore advertises.

- To fine-tune your placement request further, click the **Tags** tab and add a tag-based rule.

Tag-based rules can filter datastores by including or excluding specific placement criteria. For example, your VM storage policy can include Virtual Volumes datastores located in Texas and California, but exclude datastores located in Nevada.

- Optional: Define rules to configure datastore-specific services.

The data services, such as encryption, caching, or replication, are offered by the storage. The VM storage policy that references data services, requests these services for the VM when the VM is placed to the Virtual Volumes datastore.

- Click the tab for the data service category, for example, **Replication**.
- Define custom rules for the data service category or use predefined components.

Option	Description
Disabled	This is a default setting.
Use storage policy component	Select a storage policy component from the drop-down menu. This option is available only if you have predefined components in your database.
Custom	Define custom rules for the data service category by specifying an appropriate provider and values for the rules.

- On the **Storage compatibility** page, review the list of datastores that match this policy.

If the policy includes several rule sets, the datastore must satisfy at least one rule set and all rules within this set.

- On the **Review and finish** page, review the storage policy settings and click **Finish**.

To change any settings, click **Back** to go to the relevant page.

The new VM storage policy compatible with Virtual Volumes appears on the list.

You can now associate this policy with a virtual machine, or designate the policy as default.

Create a VM Storage Policy for Tag-Based Placement

Tag-based rules reference the tags that you assign to the datastores and can filter the datastores to be used for placement of the VMs. To define tag-based placement in the vSphere Client, use the **Create VM Storage Policy** wizard.

- Make sure that the VM Storage Policies interface is populated with information about storage entities and data services that are available in your storage environment. See [Populating the VM Storage Policies Interface in vSphere](#).
- Required privileges: **VM storage policies > Update** and **VM storage policies > View**.

Note: At a datastore cluster level, you can use a tag-based storage policy only when all datastores in the cluster have consistent tags.

1. Open the **Create VM Storage Policy** wizard.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **VM Storage Policies**.
 - c) Click **Create**.

2. Enter the policy name and description, and click **Next**.

Option	Action
vCenter Server	Select the vCenter Server instance.
Name	Enter the name of the storage policy.
Description	Enter the description of the storage policy.

3. On the **Policy structure** page under Datastore specific rules, enable tag-based placement rules.
4. On the **Tag based placement** page, create the tag rules.
 - a) Click **Add Tag Rule** and define tag-based placement criteria. Use the following as an example.

Option	Example
Tag category	Level of Service
Usage option	Use storage tagged with
Tags	Gold

All datastores with the Gold tag become compatible as the storage placement target.

- b) Optional: Add more tag-based rules.
5. On the **Storage compatibility** page, review the list of datastores that match this policy.
 6. On the **Review and finish** page, review the storage policy settings and click **Finish**.
To change any settings, click **Back** to go to the relevant page.

The new VM storage policy compatible with tagged datastores appears on the list.

Edit or Clone a VM Storage Policy

If storage requirements for virtual machines and virtual disks change, you can modify the existing storage policy. You can also create a copy of the existing VM storage policy by cloning it. While cloning, you can optionally select to customize the original storage policy.

Required privilege: **StorageProfile.View**

1. In the vSphere Client, navigate to the storage policy.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **VM Storage Policies**.
2. Select the storage policy, and click one of the following icons:
 - **Edit**
 - **Clone**
3. Optional: Modify the policy and click **OK**.
4. If editing the storage policy that is used by a virtual machine, reapply the policy to the virtual machine.

Option	Description
Manually later	If you select this option, the compliance status for all virtual disks and virtual machine home objects associated with the storage policy changes to Out of Date. To update configuration and compliance, manually reapply the storage policy to all associated entities. See Reapply Virtual Machine Storage Policy .
Now	Update virtual machine and compliance status immediately after editing the storage policy.

About vSphere Storage Policy Components

A VM storage policy in vSphere can include one or several reusable and interchangeable building blocks, called storage policy components. Each component describes a particular data service to be provided for the virtual machine. You can define the policy components in advance and associate them with multiple VM storage policies.

You cannot assign the predefined component directly to a virtual machine or virtual disk. Instead, you must add the component to the VM storage policy, and assign the policy to the virtual machine.

The component describes one type of service from one service provider. The services can vary depending on the providers that you use, but generally belong in one of the following categories.

- Compression
- Caching
- Encryption
- Replication

When you create the storage policy component, you define the rules for one specific type and grade of service.

The following example shows that virtual machines VM1 and VM2 have identical placement requirements, but must have different grades of replication services. You can create the storage policy components with different replication parameters and add these components to the related storage policies.

Table 225: Storage Policy Components

Virtual Machine	Placement Rules	Storage Policy Component
VM1 requires replication every 2 hours	Virtual Volumes Datastore	2-hour Replication
VM2 requires replication every 4 hours	Virtual Volumes Datastore	4-hour Replication

The provider of the service can be a storage system, an I/O filter, or another entity. If the component references an I/O filter, the component is added to the host-based rules of the storage policy. Components that reference entities other than the I/O filters, for example, a storage system, are added to the datastore-specific rule sets.

When you work with the components, follow these guidelines:

- Each component can include only one set of rules. All characteristics in this rule set belong to a single provider of the data services.
- If the component is referenced in the VM storage policy, you cannot delete the component. Before deleting the component, you must remove it from the storage policy or delete the storage policy.
- When you add components to the policy, you can use only one component from the same category, for example caching, per a set of rules.

Create Storage Policy Components

A storage policy component describes a single data service, such as replication, that must be provided for the virtual machine. You can define the component in advance and associate it with multiple VM storage policies. The components are reusable and interchangeable.

1. In the vSphere Client, open the **New Storage Policy Component** dialog box.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **Storage Policy Components**.
2. Click **Create Storage Policy Component**.
3. Select the vCenter Server instance.
4. Enter a name, for example, 4-hour Replication, and a description for the policy component.
Make sure that the name does not conflict with the names of other components or storage policies.
5. Select the category of the service, for example, **Replication**.
6. Select the service provider.
7. Define rules for the selected category.
For example, if you are configuring 4-hour replication, set the Recovery Point Objective (RPO) value to 4.
For encryption based on I/O filters, set the **Allow I/O filters before encryption** parameter. Encryption provided by storage does not require this parameter.

Option	Description
False (default)	Does not allow the use of other I/O filters before the encryption filter.
True	Allows the use of other I/O filters before the encryption filter. Other filters, such as replication, can analyze clear text data before it is encrypted.

8. Click **OK**.

The new component appears in the list of storage policy components.

You can add the component to the VM storage policy. If the data service that the component references is provided by the I/O filters, you add the component to the host-based rules of the storage policy. Components that reference entities other than the I/O filters, for example, a storage system, are added to the datastore-specific rule sets.

Edit or Clone Storage Policy Components

You can modify the existing storage policy components. You can also create a copy of the existing component by cloning it.

1. In the vSphere Client, navigate to the storage policy component to edit or clone.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **Storage Policy Components**.
2. Select the component and click one of the following icons.

Option	Description
Edit Settings	When editing, you cannot change the category of the data service and the provider. For example, if the original component references replication provided by I/O filters, these settings must remain unchanged.
Clone	When cloning, you can customize any settings of the original component.

3. Modify appropriate values, and click **OK**.
4. If a VM storage policy that is assigned to a virtual machine references the policy component you edit, reapply the storage policy to the virtual machine.

Menu Item	Description
Manually later	If you select this option, the compliance status for all virtual disks and virtual machine home objects associated with the storage policy changes to Out of Date. To update configuration and compliance, manually reapply the storage policy to all associated entities. See Reapply Virtual Machine Storage Policy .
Now	Update virtual machine and compliance status immediately after editing the storage policy.

vSphere Storage Policies and Virtual Machines

After you define a VM storage policy in vSphere, you can apply it to a virtual machine. You apply the storage policy when provisioning the virtual machine or configuring its virtual disks. Depending on its type and configuration, the policy might serve different purposes. The policy can select the most appropriate datastore for the virtual machine and enforce the required level of service. Or it can enable specific data services for the virtual machine and its disks.

If you do not specify the storage policy, the system uses a default storage policy that is associated with the datastore. If your storage requirements for the applications on the virtual machine change, you can modify the storage policy that was originally applied to the virtual machine.

Assign Storage Policies to Virtual Machines

You can assign a VM storage policy in an initial deployment of a virtual machine or when you perform other virtual machine operations, such as cloning or migrating.

This topic describes how to assign the VM storage policy when you create a virtual machine. For information about other deployment methods that include cloning, deployment from a template, and so on, see the *vSphere Virtual Machine Administration* documentation.

You can apply the same storage policy to the virtual machine configuration file and all its virtual disks. If storage requirements for your virtual disks and the configuration file are different, you can associate different storage policies with the VM configuration file and the selected virtual disks.

1. Start the virtual machine provisioning process and follow the appropriate steps.
2. Assign the same storage policy to all virtual machine files and disks.
 - a) On the **Select storage** page, select a storage policy from the **VM Storage Policy** drop-down menu.
Based on its configuration, the storage policy separates all datastores into compatible and incompatible. If the policy references data services offered by a specific storage entity, for example, Virtual Volumes, the compatible list includes datastores that represent only that type of storage.
 - b) Select an appropriate datastore from the list of compatible datastores.
The datastore becomes the destination storage resource for the virtual machine configuration file and all virtual disks.
 - c) If you use the replication service with Virtual Volumes, specify the replication group.
Replication groups indicate which VMs and virtual disks must be replicated together to a target site.

Option	Description
Preconfigured replication group	Replication groups that are configured in advance on the storage side. vCenter Server and ESXi discover the replication groups, but do not manage their life cycle.
Automatic replication group	Virtual Volumes creates a replication group and assigns all VM objects to this group.

3. Change the VM storage policy for the virtual disk.
Use this option if requirements for storage placement are different for virtual disks. You can also use this option to enable I/O filter services, such as caching and replication, for your virtual disks.
 - a) On the **Customize hardware** page, expand the **New hard disk** pane.
 - b) From the **VM storage policy** drop-down menu, select the storage policy to assign to the virtual disk.
 - c) Optional: Change the storage location of the virtual disk.
Use this option to store the virtual disk on a datastore other than the datastore where the VM configuration file resides.
4. Complete the virtual machine provisioning process.

After you create the virtual machine, the **Summary** tab displays the assigned storage policies and their compliance status. If storage placement requirements for the configuration file or the virtual disks change, you can later modify the virtual policy assignment.

Change Storage Policy Assignment for Virtual Machine Files and Disks

If your storage requirements for the applications on the virtual machine change, you can edit the storage policy that was originally applied to the virtual machine.

You can edit the storage policy for a powered-off or powered-on virtual machine.

When changing the VM storage policy assignment, you can apply the same storage policy to the virtual machine configuration file and all its virtual disks. You can also associate different storage policies with the VM configuration file

and the virtual disks. You might apply different policies when, for example, storage requirements for your virtual disks and the configuration file are different.

1. In the vSphere Client, browse to the virtual machine.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **VM Storage Policies**.
 - c) Click the storage policy you want to change, and click **VM Compliance**.

You can see the list of virtual machines that use this storage policy.

- d) Click the virtual machine whose policy you want to modify.
2. Click the **Configure** tab and click **Policies**.
3. Click **Edit VM Storage Policies**.
4. Specify the VM storage policy for your virtual machine.

Option	Actions
Apply the same storage policy to all virtual machine objects	Select the policy from the VM storage policy drop-down menu.
Apply different storage policies to the VM home object and virtual disks	1. Turn on the Configure per disk option. 2. Select the object, for example, VM home. 3. In the VM Storage Policy column, select the policy from the drop-down menu.

5. If you use Virtual Volumes policy with replication, configure the replication group.
Replication groups indicate which VMs and virtual disks must be replicated together to a target site.
All storage objects of a VM should belong to the same replication group. You cannot assign different replication groups to different storage objects of a VM.
6. Click **OK** to save the VM storage policy changes.

The storage policy is assigned to the virtual machine and its disks.

Check Compliance for a VM Storage Policy

You can check whether a virtual machine uses a datastore that is compatible with the storage requirements specified in the VM storage policy.

Verify that the virtual machine has a storage policy that is associated with it.

1. In the vSphere Client, navigate to the virtual machine.
2. Click the **Configure** tab and click **Policies**.
3. Click **Check VM Storage Policy Compliance**.
The system verifies the compliance.
4. View the compliance status.

Compliance Status	Description
Compliant	The datastore that the virtual machine or virtual disk uses has the storage capabilities compatible with the policy requirements.
Noncompliant	The datastore that the virtual machine or virtual disk uses does not have the storage capabilities compatible with the policy requirements. You can migrate the virtual machine files and virtual disks to compliant datastores.

Compliance Status	Description
Out of Date	The status indicates that the policy has been edited, but the new requirements have not been communicated to the datastore where the virtual machine objects reside. To communicate the changes, reapply the policy to the objects that are out of date.
Not Applicable	This storage policy references datastore capabilities that are not supported by the datastore where virtual machine resides.

When you cannot bring the noncompliant datastore into compliance, migrate the files or virtual disks to a compatible datastore. See [Find Compatible Storage Resource for Noncompliant Virtual Machine](#).

If the status is Out of Date, reapply the policy to the objects. See [Reapply Virtual Machine Storage Policy](#).

Find Compatible Storage Resource for Noncompliant Virtual Machine

Determine which datastore is compatible with the storage policy that is associated with your virtual machine.

Occasionally, a storage policy that is assigned to a virtual machine can be in the noncompliant status. This status indicates that the virtual machine or its disks use datastores that are incompatible with the policy. You can migrate the virtual machine files and virtual disks to compatible datastores.

Use this task to determine which datastores satisfy the requirements of the policy.

1. Verify that the storage policy for the virtual machine is in the noncompliant state.
 - a) In the vSphere Client, navigate to the virtual machine.
 - b) Click the **Summary** tab.

The VM Storage Policy Compliance panel on the VM Storage Policies pane shows the Noncompliant status.

2. Navigate to the noncompliant storage policy.
 - a) Click **Menu > Policies and Profiles**.
 - b) Under **Policies and Profiles**, click **VM Storage Policies**.
3. Display the list of compatible datastores for the noncompliant storage policy.
 - a) Click the storage policy.
 - b) Click **Storage Compatibility**.

The list of datastores that match the requirements of the policy appears.

You can migrate the virtual machine or its disks to one of the datastores in the list.

Reapply Virtual Machine Storage Policy

After you edit a storage policy that is already associated with a virtual machine object, you must reapply the policy. By reapplying the policy, you communicate new storage requirements to the datastore where the virtual machine object resides.

The compliance status for a virtual machine is Out of Date. The status indicates that the policy has been edited, but the new requirements have not been communicated to the datastore.

1. In the vSphere Client, navigate to the virtual machine.
2. Click the **Configure** tab and click **Policies**.
3. Verify that the compliance status is Out of Date.
4. Click **Reapply VM Storage Policy**.
5. Check the compliance status.

Compliance Status	Description
Compliant	The datastore that the virtual machine or virtual disk uses has the storage capabilities that the policy requires.
Noncompliant	The datastore supports specified storage requirements, but cannot currently satisfy the storage policy. For example, the status might become Noncompliant when physical resources of the datastore are unavailable. You can bring the datastore into compliance by making changes in the physical configuration of your host cluster. For example, by adding hosts or disks to the cluster. If additional resources satisfy the storage policy, the status changes to Compliant. When you cannot bring the noncompliant datastore into compliance, migrate the files or virtual disks to a compatible datastore. See Find Compatible Storage Resource for Noncompliant Virtual Machine.
Not Applicable	The storage policy references datastore capabilities that are not supported by the datastore.

Using Storage Providers in vSphere

A storage provider is a software component that is offered by VMware or developed by a third party through vSphere APIs for Storage Awareness (VASA). The storage provider can also be called VASA provider. The storage providers integrate with various storage entities that include external physical storage and storage abstractions, such as vSAN and Virtual Volumes. Storage providers can also support software solutions, for example, I/O filters.

Storage Provider Concepts

Generally, vCenter Server and ESXi use the storage providers to obtain information about storage configuration, status, and storage data services offered in your environment. This information appears in the vSphere Client. The information helps you to make appropriate decisions about virtual machine placement, to set storage requirements, and to monitor your storage environment.

Persistence Storage Providers

Storage providers that manage arrays and storage abstractions, are called persistence storage providers. Providers that support Virtual Volumes or vSAN belong to this category. In addition to storage, persistence providers can provide other data services, such as replication.

Data Service Providers

Another category of providers is I/O filter storage providers, or data service providers. These providers offer data services that include host-based caching, compression, and encryption.

Both persistence storage and data service providers can belong to one of these categories.

Built-in Storage Providers


Built-in storage providers are offered by VMware. Typically, they do not require registration. For example, the storage providers that support vSAN or I/O filters are built-in and become registered automatically.

Third-Party Storage Providers

When a third party offers a storage provider, you typically must register the provider. An example of such a provider is the Virtual Volumes provider. You use the vSphere Client to register and manage each storage provider component.

For more information, [Manage Storage Providers for vSphere Virtual Volumes](#).

The following graphic illustrates how different types of storage providers facilitate communications between vCenter Server and ESXi and other components of your storage environment. For example, the components might include storage arrays, Virtual Volumes storage, and I/O filters.



Storage Providers and Data Representation

Information that the storage provider supplies can be divided into the following categories:

- Storage data services and capabilities. This type of information is essential for such functionalities as vSAN, Virtual Volumes, and I/O filters. The storage provider that represents these functionalities integrates with the Storage Policy Based Management (SPBM) mechanism. The storage provider collects information about data services that are offered by underlying storage entities or available I/O filters.

You reference these data services when you define storage requirements for virtual machines and virtual disks in a storage policy. Depending on your environment, the SPBM mechanism ensures appropriate storage placement for a virtual machine or enables specific data services for virtual disks. For details, see [Creating and Managing vSphere VM Storage Policies](#).

- **Storage status.** This category includes reporting about status of various storage entities. It also includes alarms and events for notifying about configuration changes.
This type of information can help you troubleshoot storage connectivity and performance problems. It can also help you to correlate array-generated events and alarms to corresponding performance and load changes on the array.
- **Storage DRS information** for the distributed resource scheduling on block devices or file systems. This information helps to ensure that decisions made by Storage DRS are compatible with resource management decisions internal to the storage systems.

Storage Provider Requirements and Considerations

When you use third-party storage provider, vendors are responsible for supplying storage providers. The VMware VASA program defines an architecture that integrates third-party storage providers into the vSphere environment, so that vCenter Server and ESXi hosts can communicate with the storage providers.

To use storage providers, follow these requirements:

- Make sure that every storage provider you use is certified by VMware and properly deployed. For information about deploying the storage providers, contact your storage vendor.
- Make sure that the storage provider is compatible with the vCenter Server and ESXi versions. See *VMware Compatibility Guide*.
- Do not install the VASA provider on the same system as vCenter Server.
- If your environment contains older versions of storage providers, existing functionality continues to work. However, to use new features, upgrade your storage provider to a new version.

Managing Storage Providers

You might need to register certain types of storage providers to establish their connection to ESXi. You can then review and manage the storage providers.

Register Storage Providers

Use the vSphere Client to register a separate storage provider for each host in a cluster.

Verify that the storage provider component is installed on the storage side and obtain its credentials from your storage administrator.

After registration, vCenter Server can detect and use the functionality of the latest VASA version.

Note: Built-in storage providers do not require registration. For example, the storage providers that support vSAN or I/O filters are build-in. They become registered and appear on the list of storage providers automatically. For more information, see the *Administering VMware vSAN* documentation and [Filtering Virtual Machine I/O in vSphere](#).

1. Navigate to vCenter Server.
2. Click the **Configure** tab, and click **Storage Providers**.
3. Click the **Add** icon.
4. Enter connection information for the storage provider, including the name, URL, and credentials.
5. Click **OK**.

vCenter Server registers the storage provider and establishes a secure connection with it.

To troubleshoot registration of your storage provider, see the VMware Knowledge Base article <https://kb.vmware.com/s/article/49798>.

View Storage Provider Information

After you register a storage provider component with vCenter Server, the storage provider appears on the storage providers list. Certain storage providers are self-registered and automatically appear on the list after you set up the entity they represent, for example, vSAN or I/O filters.

Use the vSphere Client to view general storage provider information and details for each storage component.

1. Navigate to vCenter Server.
2. Click the **Configure** tab, and click **Storage Providers**.
3. In the Storage Providers list, view the storage providers registered with vCenter Server.
The list shows general information including the name of the storage provider, its URL and status, and version of VASA APIs.
4. To display additional details, click the icon next to the provider's name.

Note: A single storage provider can support storage systems from multiple different vendors.

You can display information about provider's certificates and storage entities it represents.

Re-authenticate Host VASA Client


Starting from vSphere 8.0 Update 3, you can monitor the storage provider connection and authentication status on each ESXi host. If there is an authentication issue in a specific host, you can re-authenticate the specific host against the storage provider.

Register storage providers. See [Register Storage Providers for Virtual Volumes](#).

Use the vSphere Client to monitor the storage providers.

1. Navigate to the ESXi host in vCenter Server.
2. Click the **Configure** tab, and click **Storage Providers**.
The **Storage Providers** tab lists the storage provider registered with the ESXi host. You can view general information including the name of the storage provider, status, and version of the VASA API. The **Client Certificate** lists the VASA client certificate details.

- If the status indicates an authentication error for a provider, click the row action (three dots) icon next to the provider's name and click **Re-authenticate Host VASA Clients**.



The VASA client re-authorizes the ESXi host with the storage provider.

Manage Storage Providers

Use the vSphere Client to perform several management operations on the registered storage providers.

- Navigate to vCenter Server.
- Click the **Configure** tab, and click **Storage Providers**.
- Click the icon to perform one of the following global actions for all providers.

Icon	Description
Synchronize Storage Providers	Synchronize all storage providers with the current state of the environment.
Re-authenticate Host VASA Clients	Re-authenticate all hosts with Virtual Volumes datastores against their respective storage providers. This icon is available only for Virtual Volumes providers. Note: This operation is time consuming.

4. To manage a specific storage provider, navigate to the provider and click the **Row action** (three dots) icon next to the provider's name.

The screenshot shows the vSphere configuration interface for Storage Providers. The left sidebar has a 'Settings' section with 'Storage Providers' selected. The main content area is titled 'Storage Providers' and has tabs for 'STORAGE PROVIDERS' and 'PROVIDERS ARRAYS'. Below these are buttons for 'ADD', 'SYNCHRONIZE STORAGE PROVIDERS', and 'RE-AUTHENTICATE HOST VASA CLIENT'. A table lists storage providers with columns for 'Storage Provider', 'Status', and 'URL'. A context menu is open for the first provider, 'vasa nvme', which is currently 'Offline'. The menu options are: Rescan, Remove, Refresh certificate, Upgrade Storage Provider, and Re-authenticate vCenter Server.

	Storage Provider	Status	URL
⋮	vasa nvme	Offline	https://192.168.1.101:5443/
⋮		Online	https://192.168.1.101:5443/
⋮		Online	https://192.168.1.101:5443/
⋮		Online	https://192.168.1.101:5443/
⋮		Online	https://192.168.1.101:5443/

5. Select one of the following from the **Row actions** menu.

Availability of the menu options depends on the type of the storage provider you use. For example, you cannot manually unregister or upgrade certain storage providers supplied by VMware, such as vSAN or I/O filter storage providers.

Option	Description
Rescan	Update storage data for the provider. vCenter Server periodically updates storage data in its database. The updates are partial and reflect only those changes that storage providers communicate to vCenter Server at that moment. When needed, you can perform a full database synchronization for the selected storage provider.
Remove	Unregister storage providers that you do not use. After this operation, vCenter Server closes the connection and removes the storage provider from its configuration.
Refresh certificate	vCenter Server warns you when a certificate assigned to a storage provider is about to expire. You can refresh the certificate to continue using the provider. If you fail to refresh the certificate before it expires, vCenter Server discontinues using the provider.
Upgrade Storage Provider	When vCenter Server notifies you about a new available VASA version, accept this version to enable it in your vSphere environment.

Option	Description
	vCenter Server activates new capabilities exposed by the storage provider and communicates this information to all hosts in the inventory. When you perform this action, you must enter credentials of the storage provider. You cannot select specific VASA version for an upgrade if several versions are available. vCenter Server upgrades only to the latest version. Note: You have the flexibility to upgrade at your own pace. For more information, see vSphere Virtual Volumes Security.
Re-authenticate vCenter Server	Re-authenticate vCenter Server against the storage provider. This action re-authorizes vCenter Server with the storage provider if it has lost authentication and would need credentials. You must enter credentials of the storage provider.

Working with VMware vSphere Virtual Volumes

VMware vSphere Virtual Volumes, also known as vVols, virtualizes storage devices by abstracting physical hardware resources into logical pools of capacity. The Virtual Volumes functionality changes the storage management paradigm from managing space inside datastores to managing abstract storage objects handled by storage arrays. With Virtual Volumes, an individual virtual machine, not the datastore, becomes a unit of storage management, while storage hardware gains complete control over virtual disk content, layout, and management.

Virtual Volumes maps virtual disks and their derivatives, clones, snapshots, and replicas, directly to objects, called virtual volumes, on a storage system. This mapping allows vSphere to offload intensive storage operations such as snapshot, cloning, and replication to the storage system.

vSphere Virtual Volumes Concepts and Architecture

Learn about basic concepts and different components of the vSphere Virtual Volumes functionality.

Watch the following video to learn about Virtual Volumes concepts.

Virtual Volume Objects

Virtual volumes are encapsulations of virtual machine files, virtual disks, and their derivatives.

Virtual volumes are stored natively inside a storage system that is connected to your ESXi hosts through Ethernet or SAN. They are exported as objects by a storage system and are managed entirely by hardware on the storage side. Typically, a unique GUID identifies a virtual volume. Virtual volumes are not preprovisioned, but created automatically when you perform virtual machine management operations. These operations include a VM creation, cloning, and snapshotting. ESXi and vCenter Server associate one or more virtual volumes to a virtual machine.

Types of Virtual Volumes

The system creates the following types of virtual volumes for the core elements that make up the virtual machine:

Data-vVol

A data virtual volume that corresponds directly to each virtual disk .vmdk file. As virtual disk files on traditional datastores, virtual volumes are presented to virtual machines as SCSI or NVMe disks. Data-vVol can be either thick or thin-provisioned.

Config-vVol

A configuration virtual volume, or a home directory, represents a small directory that contains metadata files for a virtual machine. The files include a .vmx file, descriptor files for virtual disks, log files, and so forth. The configuration virtual volume is formatted

with a file system. When ESXi uses the SCSI or NVMe protocol to connect to storage, configuration virtual volumes are formatted with VMFS. With NFS protocol, configuration virtual volumes are presented as an NFS directory. Typically, it is thin-provisioned. Starting with vSphere 7.0 Update 2, partners can increase the config-vVol to above 4 GB. Work with your Virtual Volumes partner on implementing this if it is supported by your partner and applicable to your environment.

vSphere 8.0 Update 2 supports space reclamation for config-vVols that reside on Virtual Volumes datastores accessed through the SCSI or NVMe protocols. For more information, see [Reclaim Space on the vSphere Virtual Volumes Datastores](#)

Swap-vVol

Created when a VM is first powered on. It is a virtual volume to hold copies of VM memory pages that cannot be retained in memory. Its size is determined by the VM's memory size. It is thick-provisioned by default.

Snapshot-vVol

A virtual memory volume to hold the contents of virtual machine memory for a snapshot. Thick-provisioned.

For more information, see [vSphere Virtual Volumes Snapshots](#).

Other

A virtual volume for specific features. For example, a digest virtual volume is created for Content-Based Read Cache (CBRC).

Typically, a VM creates a minimum of three virtual volumes, data-vVol, config-vVol, and swap-vVol. The maximum depends on how many virtual disks and snapshots reside on the VM.

By using different virtual volumes for different VM components, you can apply and manipulate storage policies at the finest granularity level. For example, a virtual volume that contains a virtual disk can have a richer set of services than the virtual volume for the VM boot disk.

Disk Provisioning

The Virtual Volumes functionality supports a concept of thin and thick-provisioned virtual disks. However, from the I/O prospective, implementation and management of thin or thick provisioning by the arrays is transparent to the ESXi host. ESXi offloads to the storage arrays any functions related to thin provisioning.

You select the thin or thick type for your virtual disk at the VM creation time. If your disk is thin and resides on a Virtual Volumes datastore, you cannot change its type later by inflating the disk.

Shared Disks

You can place a shared disk on a Virtual Volumes storage that supports SCSI Persistent Reservations for Virtual Volumes. You can use this disk as a quorum disk and eliminate RDMS in the MSCS clusters. For more information, see the *vSphere Resource Management* documentation.

Virtual Volumes Storage Providers

A Virtual Volumes storage provider, also called a VASA provider, is a software component that acts as a storage awareness service for vSphere. The provider mediates out-of-band communication between vCenter Server and ESXi hosts on one side and a storage system on the other.

The storage provider is implemented through VMware APIs for Storage Awareness (VASA) and is used to manage all aspects of Virtual Volumes storage.

The storage provider delivers information from the underlying storage container. The storage container capabilities appear in vCenter Server and the vSphere Client. Then, in turn, the storage provider communicates virtual machine storage

requirements, which you can define in the form of a storage policy, to the storage layer. This integration process ensures that a virtual volume created in the storage layer meets the requirements outlined in the policy.

Typically, vendors are responsible for supplying storage providers that can integrate with vSphere and provide support to Virtual Volumes. Every storage provider must be certified by VMware and properly deployed. For information about deploying and upgrading the Virtual Volumes storage provider to a version compatible with current ESXi release, contact your storage vendor.

After you deploy the storage provider, you must register it in vCenter Server. See [Register Storage Providers for Virtual Volumes](#). To upgrade your storage providers or for other operations that you can perform, see [Manage Storage Providers for vSphere Virtual Volumes](#).

Virtual Volumes Storage Containers

Unlike traditional block or file based storage, the Virtual Volumes functionality does not require preconfigured storage on a storage side. Instead, Virtual Volumes uses a storage container. It is a pool of raw storage capacity or an aggregation of storage capabilities that a storage system can provide to virtual volumes.

A storage container is a part of the logical storage fabric and is a logical unit of the underlying hardware. The storage container logically groups virtual volumes based on management and administrative needs. For example, the storage container can contain all virtual volumes created for a tenant in a multitenant deployment, or a department in an enterprise deployment. Each storage container serves as a virtual volume store and virtual volumes are allocated out of the storage container capacity.

Typically, a storage administrator on the storage side defines storage containers. The number of storage containers, their capacity, and their size depend on a vendor-specific implementation. At least one container for each storage system is required.

Note: A single storage container cannot span different physical arrays.

After you register a storage provider associated with the storage system, vCenter Server discovers all configured storage containers along with their storage capability profiles, protocol endpoints, and other attributes. A single storage container can export multiple capability profiles. As a result, virtual machines with diverse needs and different storage policy settings can be a part of the same storage container.

Initially, all discovered storage containers are not connected to any specific host, and you cannot see them in the vSphere Client. To mount a storage container, you must map it to a Virtual Volumes datastore.

Static Protocol Endpoints

With SCSI or NFS transports, ESXi hosts use a logical I/O proxy, called the protocol endpoint, to communicate with virtual volumes. ESXi uses protocol endpoints to establish a data path on demand from virtual machines to their respective virtual volumes.

Note: Information in this section applies only to static protocol endpoints that use the SCSI or NFS transports. For specifics about the NVMe protocol endpoints, see [NVMe and Virtual Volumes in vSphere](#).

Each virtual volume is bound to a specific protocol endpoint. When a virtual machine on the host performs an I/O operation, the protocol endpoint directs the I/O to the appropriate virtual volume. Typically, a storage system requires just a few protocol endpoints. A single protocol endpoint can connect to hundreds or thousands of virtual volumes.

On the storage side, a storage administrator configures protocol endpoints, one or several per storage container. The protocol endpoints are a part of the physical storage fabric. The storage system exports the protocol endpoints with associated storage containers through the storage provider. After you map the storage container to a Virtual Volumes datastore, the ESXi host discovers the protocol endpoints and they become visible in the vSphere Client. The protocol endpoints can also be discovered during a storage rescan. Multiple hosts can discover and mount the protocol endpoints.

In the vSphere Client, the list of available protocol endpoints looks similar to the host storage devices list. Different storage transports can be used to expose the protocol endpoints to ESXi. When the SCSI-based transport is used, the protocol endpoint represents a proxy LUN defined by a T10-based LUN WWN. For the NFS protocol, the protocol endpoint is a mount point, such as an IP address and a share name. You can configure multipathing on the SCSI-based protocol endpoint, but not on the NFS-based protocol endpoint. No matter which protocol you use, the storage array can provide multiple protocol endpoints for availability purposes.

Protocol endpoints are managed per array. ESXi and vCenter Server assume that all protocol endpoints reported for an array are associated with all containers on that array. For example, if an array has two containers and three protocol endpoints, ESXi assumes that virtual volumes on both containers can be bound to all three protocol points.

For information about viewing the static protocol endpoints in the vSphere Client, see [Review and Manage Static Protocol Endpoints](#).

Binding and Unbinding Virtual Volumes

At the time of creation, a virtual volume is a passive entity and is not immediately ready for I/O. To access the virtual volume, ESXi or vCenter Server send a bind request.

The storage system replies with a protocol endpoint ID that becomes an access point to the virtual volume. The protocol endpoint accepts all I/O requests to the virtual volume. This binding exists until ESXi sends an unbind request for the virtual volume.

For later bind requests on the same virtual volume, the storage system can return different protocol endpoint IDs.

When you use the NVMe protocol, the bind virtual volume response provides the NVMe subsystem NQN and the namespace ID (nsid) of the namespace virtual volume object. The ESXi host uses this information and resolves this to the ANA group within the subsystem. Corresponding to this ANA group, a virtual protocol endpoint (vPE) is created if it does not exist. It is used to direct all I/O requests to Virtual Volumes.

When receiving concurrent bind requests to a virtual volume from multiple ESXi hosts, the storage system can return the same or different endpoint bindings to each requesting ESXi host. In other words, the storage system can bind different concurrent hosts to the same virtual volume through different endpoints.

The unbind operation removes the I/O access point for the virtual volume. The storage system might unbind the virtual volume from its protocol endpoint immediately, or after a delay, or take some other action. A bound virtual volume cannot be deleted until it is unbound.

Virtual Volumes Datastores

A Virtual Volumes datastore represents a storage container in vCenter Server and the vSphere Client.

After vCenter Server discovers storage containers exported by storage systems, you must mount them as Virtual Volumes datastores. The Virtual Volumes datastores are not formatted in a traditional way like, for example, VMFS datastores. You must still create them because all vSphere functionalities, including FT, HA, DRS, and so on, require the datastore construct to function properly.

You use the datastore creation wizard in the vSphere Client to map a storage container to a Virtual Volumes datastore. The Virtual Volumes datastore that you create corresponds directly to the specific storage container.

From a vSphere administrator perspective, the Virtual Volumes datastore is similar to any other datastore and is used to hold virtual machines. Like other datastores, the Virtual Volumes datastore can be browsed and lists virtual volumes by virtual machine name. Like traditional datastores, the Virtual Volumes datastore supports unmounting and mounting. However, such operations as upgrade and resize are not applicable to the Virtual Volumes datastore. The Virtual Volumes datastore capacity is configurable by the storage administrator outside of vSphere.

You can use the Virtual Volumes datastores with traditional VMFS and NFS datastores and with vSAN.

Note: The size of a virtual volume must be a multiple of 1 MB, with a minimum size of 1 MB. As a result, all virtual disks that you provision on a Virtual Volumes datastore must be an even multiple of 1 MB. If the virtual disk you migrate to the Virtual Volumes datastore is not an even multiple of 1 MB, extend the disk to the nearest even multiple of 1 MB.

To create a Virtual Volumes datastore, see [Create a Virtual Volumes Datastore in vSphere Environment](#).

Virtual Volumes and VM Storage Policies

A virtual machine that runs on a Virtual Volumes datastore requires a VM storage policy.

A VM storage policy is a set of rules that contains placement and quality-of-service requirements for a virtual machine. The policy enforces appropriate placement of the virtual machine within Virtual Volumes storage and guarantees that storage can satisfy virtual machine requirements.

You use the VM Storage Policies interface to create a Virtual Volumes storage policy. When you assign the new policy to the virtual machine, the policy enforces that the Virtual Volumes storage meets the requirements.

Virtual Volumes Default Storage Policy

For Virtual Volumes, VMware provides a default storage policy that contains no rules or storage requirements, called Virtual Volumes No Requirements Policy. This policy is applied to the VM objects when you do not specify another policy for the virtual machine on the Virtual Volumes datastore. With the No Requirements policy, storage arrays can determine the optimum placement for the VM objects.

The default No Requirements policy that VMware provides has the following characteristics:

- You cannot delete, edit, or clone this policy.
- The policy is compatible only with the Virtual Volumes datastores.
- You can create a VM storage policy for Virtual Volumes and designate it as the default.

Virtual Volumes and Storage Protocols

A Virtual Volumes storage system provides protocol endpoints that are discoverable on the physical storage fabric. ESXi hosts use the protocol endpoints to connect to virtual volumes on the storage. Operation of the protocol endpoints depends on storage protocols that expose the endpoints to ESXi hosts.

Virtual Volumes supports NFS version 3 and 4.1, iSCSI, Fibre Channel, FCoE, NVMe over Fibre Channel, and NVMe over TCP.

Irrespective of which storage protocol is used, protocol endpoints provide uniform access to both SAN and NAS storage. A virtual volume, like a file on other traditional datastore, is presented to a virtual machine as a SCSI or an NVMe disk.

Note:

A storage container is dedicated to SAN (SCSI or NVMe) or NAS and cannot be shared across those protocol types. An array can present one storage container with SCSI protocol endpoints and a different container with NFS protocol endpoints. The container cannot use a combination of SCSI, NVMe, and NFS storage access protocols.

Virtual Volumes and SCSI-Based Transports

On disk arrays, Virtual Volumes supports Fibre Channel, FCoE, and iSCSI protocols.

When the SCSI-based protocol is used, the protocol endpoint represents a proxy LUN defined by a T10-based LUN WWN.

As any block-based LUNs, the protocol endpoints are discovered using standard LUN discovery commands. The ESXi host periodically rescans for new devices and asynchronously discovers block-based protocol endpoints. The protocol

endpoint can be accessible by multiple paths. Traffic on these paths follows well-known path selection policies, as is typical for LUNs.

On SCSI-based disk arrays at VM creation time, ESXi makes a virtual volume and formats it as VMFS. This small virtual volume stores all VM metadata files and is called the config-vVol. The config-vVol functions as a VM storage locator for vSphere.

Virtual Volumes on disk arrays supports the same set of SCSI commands as VMFS and use ATS as a locking mechanism.

CHAP Support for iSCSI Endpoints

Virtual Volumes supports Challenge Handshake Access Protocol (CHAP) with iSCSI targets. This support allows ESXi hosts to share CHAP initiator credentials with Virtual Volumes storage providers, also called VASA providers, and for Virtual Volumes storage providers to raise system events notifying vCenter Server of changes to CHAP target credentials on the storage array.

Each ESXi host can have multiple HBAs and each HBA can have properties configured on it. One of these properties is the authentication method that the HBA must use. Authentication is optional, but if implemented, it must be supported by both the initiator and target. CHAP is an authentication method that can be used both directions between initiator and target.

For more information about different CHAP authentication methods, see [Selecting CHAP Authentication Method](#). To configure CHAP on your ESXi host, see [Configuring CHAP Parameters for iSCSI or iSER Storage Adapters on ESXi Host](#).

Virtual Volumes and NFS Transports

With NAS storage, a protocol endpoint is an NFS share that the ESXi host mounts using IP address or DNS name and a share name. Virtual Volumes supports NFS version 3 and 4.1 to access NAS storage. Both IPv4 and IPv6 formats are supported.

No matter which version you use, a storage array can provide multiple protocol endpoints for availability purposes.

In addition, NFS version 4.1 introduces trunking mechanisms that enable load balancing and multipathing.

Virtual Volumes on NAS devices supports the same NFS Remote Procedure Calls (RPCs) that ESXi hosts use when connecting to NFS mount points.

On NAS devices, a config-vVol is a directory subtree that corresponds to a config-vVolID. The config-vVol must support directories and other operations that are necessary for NFS.

Virtual Volumes and NVMe


Virtual Volumes supports NVMe protocols, including NVMe over Fibre Channel and NVMe over TCP. A virtual volume object is mapped to a namespace within an NVMe subsystem. ANA groups within the NVMe subsystem are viewed as virtual protocol endpoints on the ESXi host.

Virtual protocol endpoints are used for path state management as the ANA group state changes. The ESXi host discovers the ANA groups dynamically, as and when required. This means that the virtual protocol endpoint is created only when the ESXi host needs I/O access to a namespace virtual volume within the NVMe subsystem. Config-vVols on NVMe are similar to SCSI that are formatted with VMFS. They are also used to store the VM metadata files.

To configure NVMe with Virtual Volumes on your ESXi host, see [NVMe and Virtual Volumes in vSphere](#).

Virtual Volumes Architecture

An architectural diagram provides an overview of how all components of the Virtual Volumes functionality interact with each other.



Watch the video for information about Virtual Volumes architecture.

vSphere Virtual Volumes Security

Learn about security improvements that vSphere Virtual Volumes storage provider, also called VASA provider, version 5 offers in vSphere 8.0 Update 1 and later releases. You can also learn about the security model in vSphere 8.0 and earlier releases.

VASA 5 Support and Security with vSphere 8.0 Update 1 and Later

All storage providers of VASA 5 and later versions use a stricter authentication mechanism, which requires ESXi to be authenticated in context of vCenter Server. VASA 5 improves security and offers a significantly modified management model that includes the following major changes:

- For each vCenter Server that registers with the array using VASA 5 or later, the VASA provider creates a dedicated web server instance, or a virtual host, which can be either a virtual web server instance or a completely separate instance. The VASA client in vCenter Server relies on certificate based authentication and authorization to access its dedicated virtual host created on the array. All VASA client certificates, including vCenter Server and ESXi certificates, get registered with the virtual host. This creates strong isolation between different vCenter Server systems when the systems are authenticated. In addition, VASA providers can offer separate resource access and improved isolation to different vCenter Server systems.
- With VASA 5, the VASA client uses a dedicated certificate for VASA communications. Each vCenter Server provisions a certificate for the VASA provider, which is managed through a dedicated virtual host specific to vCenter Server. All ESXi hosts that support VASA 5 use the dedicated virtual host created by their managing vCenter Server.
- vCenter Server provisions the VASA client certificate for each new ESXi host 8.0 Update 1 or later and synchronizes the public key of the certificate with the VASA provider. In contrast with the previous security model that authenticated the CA issuer for the client certificate, the VASA provider now identifies and authorizes an individual client using the public key.
- To comply with VMware security requirements, vSphere does not trust self-signed certificates for TLS communications. The only exception is during a short period when the VASA provider gets registered and for backward compatibility. An array administrator can use a Custom CA certificate for the VASA provider to override the self-signed certificate at array for backward compatibility and bootstrapping.
- To avoid losing access to the VASA provider, follow these guidelines. For information, see [Manage Storage Providers for vSphere Virtual Volumes](#).
 - Do not unregister and re-register your VASA provider to upgrade. Instead, use a proper upgrade mechanism for your VASA provider. When vCenter Server notifies you about a new available VASA version, make sure to accept this version within a reasonable time. To upgrade from the vSphere Client, use the **Upgrade Storage Provider** option.
 - Regularly refresh the VASA provider certificate. Make sure to refresh the certificate within a reasonable time after vCenter Server warns you that the certificate assigned to the VASA provider is about to expire. Use the vSphere Client **Refresh certificate** option.
 - When you renew the VMCA root certificate or vCenter Server client certificate, SMS might lose connection with the VASA provider. If the provider is offline, use the **Re-authenticate vCenter Server** option.
 - If a host loses authentication, you can remediate the authentication failure by using the **Re-authenticate Host VASA Clients** option.

Security Certificates with vSphere 8.0 and Earlier Releases

vSphere includes the VMware Certificate Authority (VMCA). By default, the VMCA creates all internal certificates used in vSphere environment. It generates certificates for newly added ESXi hosts and storage VASA providers that manage or represent Virtual Volumes storage systems.

Communication with the VASA provider is protected by SSL certificates. These certificates can come from the VASA provider or from the VMCA.

- Certificates can be directly provided by the VASA provider for long-term use. They can be either self-generated and self-signed, or derived from an external Certificate Authority.
- Certificates can be generated by the VMCA for use by the VASA provider.

When a host or VASA provider is registered, VMCA follows these steps automatically, without involvement from the vSphere administrator.


1. When a VASA provider is first added to the vCenter Server storage management service (SMS), it produces a self-signed certificate.
2. After verifying the certificate, the SMS requests a Certificate Signing Request (CSR) from the VASA provider.
3. After receiving and validating the CSR, the SMS presents it to the VMCA on behalf of the VASA provider, requesting a CA signed certificate.
The VMCA can be configured to function as a standalone CA, or as a subordinate to an enterprise CA. If you set up the VMCA as a subordinate CA, the VMCA signs the CSR with the full chain.
4. The signed certificate with the root certificates is passed to the VASA provider. The VASA provider can authenticate all future secure connections originating from the SMS on vCenter Server and on ESXi hosts.

vSphere Virtual Volumes Snapshots

In vSphere, snapshots preserve the state and data of a virtual machine at the time you take the snapshot. Snapshots are useful when you must revert repeatedly to the same virtual machine state, but you do not want to create multiple virtual machines. Virtual Volumes snapshots serve many purposes. You can use them to create a quiesced copy for backup or archival purposes, or to create a test and rollback environment for applications. You can also use them to provision application images instantly.

In Virtual Volumes environment, snapshots are managed by ESXi and vCenter Server, but are performed by the storage array.

Each snapshot creates an extra virtual volume object, a snapshot virtual volume, that holds the contents of virtual machine memory. Original VM data is copied to this object, and it remains read-only, which prevents the guest operating system from writing to snapshot. You cannot resize the snapshot virtual volume. Typically, when you replicate the VM, its snapshot virtual volume is also replicated.



The base virtual volume remains active, or read-write. When another snapshot is created, it preserves the new state and data of the virtual machine at the time you take the snapshot.

When you delete snapshots, only the base virtual volume remains, while snapshot virtual volume objects are discarded. The base virtual volume represents the most current state of the virtual machine. Unlike snapshots on the traditional datastores, snapshot virtual volumes do not need to commit their contents to the base virtual volume.



For information about creating and managing snapshots, see the *vSphere Virtual Machine Administration* documentation.

vSphere Virtual Volumes and Shared Virtual Disks

Virtual Volumes support shared virtual disks that can be used with such clustering software as Windows Server Failover Clustering (WSFC). Starting with vSphere 8.0 Update 2, Virtual Volumes support the hot extension functionality for shared virtual disks on the Virtual Volumes datastore. With this functionality, you can increase the size of a shared disk without bringing the cluster down.

To learn more about the support Virtual Volumes provide to shared disks with WSFC, see the following information:

- [VMware vSphere® Virtual Volumes Support for WSFC](#)
- [Virtual Volumes \(vVols\) now supports WSFC](#)
- [vVols and Microsoft WSFC](#)

For information about setting up and extending the shared virtual disk on the Virtual Volumes datastore, see [Add Hard Disks to the First Node of a WSFC Across Physical Hosts with vVols](#).

Before You Enable vSphere Virtual Volumes

To work with vSphere Virtual Volumes, ensure that your storage and vSphere environment are set up correctly.

Prepare Storage System for Virtual Volumes

To prepare your storage system environment for Virtual Volumes, follow these guidelines. For additional information, contact your storage vendor.

- The storage system or storage array that you use must support Virtual Volumes and integrate with the vSphere components through vSphere APIs for Storage Awareness (VASA). The storage array must support thin provisioning and snapshotting.
- The Virtual Volumes storage provider must be deployed.
- The following components must be configured on the storage side:
 - Protocol endpoints. Except NVMe.
 - Storage containers
 - Storage profiles
 - Replication configurations if you plan to use Virtual Volumes with replication. See [Requirements for Replication with Virtual Volumes](#).
- If you use NVMe, create NVM subsystems and associate Virtual Volumes Host NQNs with relevant NVM subsystems. This configuration step depends on partner specific NVMe design. For more information, contact your array vendor.

Prepare vSphere Environment

- Ensure to follow appropriate setup guidelines for the type of storage you use, NVMe over Fibre Channel, NVMe over TCP, Fibre Channel, FCoE, iSCSI, or NFS. If necessary, install and configure storage adapters on your ESXi hosts.
 - If you use iSCSI, activate the software iSCSI adapters on your ESXi hosts. Configure Dynamic Discovery and enter the IP address of your Virtual Volumes storage system. See [Configure the Software iSCSI Adapter with ESXi](#).
 - If you use NVMe, follow the requirements for NVMe over Fibre Channel and NVMe over TCP. See [Requirements for VMware NVMe Storage](#).
- Synchronize all components in the storage array with vCenter Server and all ESXi hosts. Use Network Time Protocol (NTP) to do this synchronization.

For more information, contact your vendor and see *VMware Compatibility Guide*

Synchronize vSphere Storage Environment with a Network Time Server

If you use Virtual Volumes, configure Network Time Protocol (NTP) to make sure all ESXi hosts on the vSphere network are synchronized.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, select **Time Configuration**.
4. Click **Edit** and set up the NTP server.
 - a) Select **Use Network Time Protocol (Enable NTP client)**.
 - b) Set the NTP Service Startup Policy.
 - c) Enter the IP addresses of the NTP server to synchronize with.
 - d) Click **Start** or **Restart** in the NTP Service Status section.
5. Click **OK**.

The host synchronizes with the NTP server.

Configure vSphere Virtual Volumes

To configure Virtual Volumes in the vSphere environment, follow several steps.

- Be familiar with concepts and requirements related to Virtual Volumes in the vSphere environment. See [vSphere Virtual Volumes Concepts and Architecture](#) and [NVMe and Virtual Volumes in vSphere](#).
 - Follow guidelines in [Before You Enable vSphere Virtual Volumes](#).
1. [Register Storage Providers for Virtual Volumes](#).
 2. [Create a Virtual Volumes Datastore in vSphere Environment](#).

Register Storage Providers for Virtual Volumes

Your Virtual Volumes environment must include storage providers, also called VASA providers. Typically, third-party vendors develop storage providers through the VMware APIs for Storage Awareness (VASA). Storage providers facilitate communication between vSphere and the storage side. Use the vSphere Client to register the Virtual Volumes storage providers.

Verify that an appropriate version of the Virtual Volumes storage provider is installed on the storage side. Obtain credentials of the storage provider.

After registration, the Virtual Volumes provider communicates with vCenter Server. The provider reports characteristics of underlying storage and data services, such as replication, that the storage system provides. The characteristics appear in the VM Storage Policies interface and can be used to create a VM storage policy compatible with the Virtual Volumes

datastore. After you apply this storage policy to a virtual machine, the policy is pushed to Virtual Volumes storage. The policy enforces optimal placement of the virtual machine within Virtual Volumes storage and guarantees that storage can satisfy virtual machine requirements. If your storage provides extra services, such as caching or replication, the policy enables these services for the virtual machine.

For more information about the storage providers, see [Virtual Volumes Storage Providers](#).


1. Navigate to vCenter Server.
2. Click the **Configure** tab, and click **Storage Providers**.
3. Click the **Add** icon.
4. Enter connection information for the storage provider, including the name, URL, and credentials.
5. To complete the registration, click **OK**.

vCenter Server

discovers and registers the

Virtual Volumes

storage provider.



6. Verify information available for your provider.
- Switch between **Storage Providers** and **Providers Arrays** to display information about your provider and the array it represents.

Create a Virtual Volumes Datastore in vSphere Environment

You use the **New Datastore** wizard to create a Virtual Volumes datastore.

For more information about the Virtual Volumes datastore, see [Virtual Volumes Datastores](#).

1. In the vSphere Client object navigator, browse to a host, a cluster, or a data center.
2. From the right-click menu, select **Storage > New Datastore**.
3. Select **vVol** as the datastore type.
4. Enter the datastore name and select a backing storage container from the list of storage containers.
Make sure to use the name that does not duplicate another datastore name in your data center environment.
If you mount the same Virtual Volumes datastore to several hosts, the name of the datastore must be consistent across all hosts.

5. Select the hosts that require access to the datastore.
6. Review the configuration options and click **Finish**.

New Datastore

1 Type

2 Name and container selection

3 Hosts accessibility

4 Ready to complete

Ready to complete

Review your selections before finishing the wizard

▼ **Name and container selection**

Datastore name	Datastore
Datastore type	vVol
Storage container name:	coke2
Storage container UUID:	vvol:c0b1db-
Storage array(s):	xVP SCSI Array 2
Storage provider(s):	xvp-vasa4, xvp-vasa3

▼ **Hosts accessibility**

Hosts	10.161.16 10.161.16 10.161.18 10.161.17
-------	---

After you create the Virtual Volumes datastore, you can perform the following operations:

- Rename the datastore, browse datastore files, unmount the datastore, and so on. See [Basic Administrative Operations for vSphere Datastores](#).
- Provision virtual machines on the Virtual Volumes datastore. See [Provision VMs on vSphere Virtual Volumes Datastores](#) and the *vSphere Virtual Machine Administration* documentation.
- Review and manage protocol endpoints. See [Review and Manage Protocol Endpoints in vSphere](#).
- Manage or upgrade storage providers. See [Manage Storage Providers for vSphere Virtual Volumes](#).

Provision VMs on vSphere Virtual Volumes Datastores

Follow this procedure to provision virtual machines on a vSphere Virtual Volumes datastore.

Note: All virtual disks that you provision on a Virtual Volumes datastore must be an even multiple of 1 MB.

A virtual machine that runs on a Virtual Volumes datastore requires an appropriate VM storage policy.

After you provision the virtual machine, you can perform typical VM management tasks. For information, see the *vSphere Virtual Machine Administration* documentation.

1. Define a VM storage policy for Virtual Volumes.

VMware provides a default No Requirements storage policy for Virtual Volumes. If you need, you can create a custom storage policy compatible with Virtual Volumes.

See [Create a VM Storage Policy for Virtual Volumes](#).

2. Assign the Virtual Volumes storage policy to the virtual machine.

To guarantee that the Virtual Volumes datastore fulfills specific storage requirements when allocating a virtual machine, associate the Virtual Volumes storage policy with the virtual machine.

See [Assign Storage Policies to Virtual Machines](#).

3. Change default storage policy for a Virtual Volumes datastore.

For virtual machines provisioned on Virtual Volumes datastores, VMware provides a default No Requirements policy. You cannot edit this policy, but you can designate a newly created policy as default.

See [Change the Default Storage Policy for a Datastore](#).

Review and Manage Protocol Endpoints in vSphere

Learn how you can review and manage protocol endpoints that the ESXi host uses to access Virtual Volumes on storage arrays. How the protocol endpoints are represented in the vSphere Client depends on whether the host uses the SCSI and NFS protocols or the NVMe protocol to access Virtual Volumes.

Review and Manage Static Protocol Endpoints

With the SCSI or NFS protocols, the protocol endpoints represent a proxy LUN or an NFS mount point. These protocol points are called static. After you map a storage container to a Virtual Volumes datastore, these protocol endpoints

become visible in the vSphere Client. The list of available static protocol endpoints looks similar to the host storage devices list. You can review properties of the protocol endpoints and modify specific settings.

For information about static protocol endpoints that use SCSI or NFS transports, see [Static Protocol Endpoints](#).

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Protocol Endpoints**.
4. Click the **Protocol Endpoints** tab.
5. To view details for a specific item, select this item from the list.

The screenshot shows the vSphere Client interface. The left sidebar displays the navigation tree with the following structure:

- SC- [redacted]
- > a
- vcqaDC
 - cls
 - 10.161.16 [redacted]** (selected)
 - 10.161.16 [redacted]
 - 10.161.17 [redacted]
 - 10.161.18 [redacted]
 - New Cluster

The main pane shows the configuration for the selected host **10.161.16 [redacted]**. The **Configure** tab is active, and the **Protocol Endpoints** section under **Storage** is selected. The right pane shows the **PROTOCOL EN** table with the following entries:

	Name
☒	SCST_B i.316 [redacted]
☐	SCST_B i.373 [redacted]

Below the table, there is a **Manage Columns** button and a **Properties** section with a **General** tab. The **Runtime na** property is visible.

6. Use tabs to access additional information and modify properties for the selected protocol endpoint.

Tab	Description
Properties	View the item properties and characteristics. For SCSI (block) items, view and edit multipathing policies.
Paths (SCSI protocol endpoints only)	Display paths available for the protocol endpoint. Deactivate or activate a selected path. Change the Path Selection Policy.
Datastores	Display a corresponding Virtual Volumes datastore. Perform datastore management operations.

Change the Path Selection Policy for a SCSI Protocol Endpoint

If your ESXi host uses SCSI-based transport to communicate with protocol endpoints representing a storage array, you can modify default multipathing policies assigned to protocol endpoints. Use the **Edit Multipathing Policies** dialog box to change a path selection policy.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Protocol Endpoints**.
4. Select the protocol endpoint whose paths you want to change and click the **Properties** tab.
5. Under Multipathing Policies, select **Edit Multipathing** from the **Actions** menu.

The screenshot shows the vSphere Client interface. The left sidebar displays a tree view of the environment, with the path **SC- > a > vcqaDC > cls > 10.161.1.100** selected. The main pane is divided into three sections:

- Summary:** Shows the selected protocol endpoint, **SCST_BIO iSCSI Disk (eui.634.00000000000000000000000000000000)**.
- Monitor:** Shows the selected protocol endpoint, **SCST_BIO iSCSI Disk (eui.634.00000000000000000000000000000000)**.
- Configure:** Shows the selected protocol endpoint, **SCST_BIO iSCSI Disk (eui.634.00000000000000000000000000000000)**.

The **Configure** tab is active, and the **Protocol Endpoints** section is expanded. The **Properties** sub-tab is selected, showing the **General** properties of the endpoint. The **Multipathing Policies** section is expanded, and the **Path Selection Policy** is highlighted. The **Actions** menu is open, showing the **Edit Multipathing** option.

6. Select a path policy and configure its settings.

For information about path policies for SCSI devices, see [VMware Path Selection Plug-Ins and Policies](#).

7. To save your settings and exit the dialog box, click **OK**.

Review NVMe Protocol Endpoints

Unlike the static protocol endpoint, the NVMe protocol endpoint is virtual. This protocol endpoint is a logical representation of the ANA group where the virtual volumes reside. Until a VM is powered on, the NVMe protocol point doesn't exist. Once

the VM is powered on, the NVMe protocol point is created, so that the host can access the virtual volumes in the ANA group. Use the vSphere Client to review NVMe virtual protocol endpoints.

For conceptual information that applies to Virtual Volumes with NVMe storage, see [NVMe and Virtual Volumes in vSphere](#).

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Protocol Endpoints**.
4. Click the **NVMe Protocol Endpoints** tab and review the list of available NVMe protocol endpoints and their information.

You can see such details as target NVMe subsystems that are tracked for the NVMe Virtual Volumes datastore.

The screenshot shows the vSphere Client interface. The top navigation bar includes the vSphere Client logo and a search bar. The left sidebar shows the inventory tree with the path: SC- > a > vcqaDC > cls > 10.161. The main content area is divided into three sections: a left-hand navigation pane, a central configuration pane, and a right-hand details pane.

Left-hand navigation pane: Shows the configuration categories for the selected host (10.161). The categories are: Storage, Networking, Virtual Machines, and System. The 'Storage' category is expanded, showing 'Protocol Endpoints' as the selected option.

Central configuration pane: Displays the 'Protocol Endpoints' configuration page. It includes a 'Target NQN' field with the value 'nqn.2016-06.com.vvol:storage.nvme-spdk-0050!0-subsys-1'. Below this is a 'Manage Columns' button.

Right-hand details pane: Displays the 'Datastores' section. It includes a 'Quick Filter' field and a table listing the available NVMe virtual volumes.

Name	Status
nvme vvol	✓ Normal

- To display additional details, select a protocol endpoint from the list.

You can display a Virtual Volumes datastore that corresponds to the specific protocol endpoint.

Manage Storage Providers for vSphere Virtual Volumes

After you register storage providers for Virtual Volumes, use the vSphere Client to manage the providers.

Register storage providers. See [Register Storage Providers for Virtual Volumes](#).

For more information about the storage providers, see [Virtual Volumes Storage Providers](#).

- Navigate to vCenter Server.
- Click the **Configure** tab, and click **Storage Providers**.
- Click the icon to perform one of the following global actions for all Virtual Volumes providers.

Icon	Description
Synchronize Storage Providers	Synchronize all storage providers with the current state of the environment.
Re-authenticate Host VASA Clients	Re-authenticate all hosts against their respective storage providers. Applies to ESXi hosts version 8.0 Update 1 or later. The hosts must support VASA 5. Note: This operation is time consuming. To re-authenticate only specific hosts, see Re-authenticate Host VASA Client .

- To manage a specific storage provider, navigate to the provider and click the **Row action** (three dots) icon next to the provider's name.

The screenshot shows the vSphere Client interface for managing storage providers. The 'Configure' tab is active, and the 'Storage Providers' section is selected in the left sidebar. The main content area shows a table of storage providers. A context menu is open over the first row, displaying the following actions:

- Rescan
- Remove
- Refresh certificate
- Upgrade Storage Provider
- Re-authenticate vCenter Server

Storage Provider	Status	URL
vasa nvme	Offline	https://192.168.1.101:443/vasa/
	Online	https://192.168.1.102:443/vasa/
	Online	https://192.168.1.103:443/vasa/
	Online	https://192.168.1.104:443/vasa/
	Online	https://192.168.1.105:443/vasa/

5. Select one of the following from the **Row actions** menu.

Option	Description
Rescan	Update storage data for the provider. vCenter Server periodically updates storage data in its database. The updates are partial and reflect only those changes that storage providers communicate to vCenter Server at that moment. When needed, you can perform a full database synchronization for the selected storage provider.
Remove	Unregister storage providers that you do not use. After this operation, vCenter Server closes the connection and removes the storage provider from its configuration.
Refresh certificate	vCenter Server warns you when a certificate assigned to a storage provider is about to expire. You can refresh the certificate to continue using the provider. If you fail to refresh the certificate before it expires, vCenter Server discontinues using the provider.
Upgrade Storage Provider	When vCenter Server notifies you about a new available VASA version, accept this version to enable it in your vSphere environment. vCenter Server activates new capabilities exposed by the storage provider and communicates this information to all hosts in the inventory. When you perform this action, you must enter credentials of the storage provider. You cannot select specific VASA version for an upgrade if several versions are available. vCenter Server upgrades only to the latest version. Note: You have the flexibility to upgrade at your own pace. For more information, see vSphere Virtual Volumes Security .
Re-authenticate vCenter Server	Re-authenticate vCenter Server against the storage provider. This action re-authorizes vCenter Server with the storage provider if it has lost authentication and would need credentials. You must enter credentials of the storage provider.

Re-authenticate Host VASA Client

Starting from vSphere 8.0 Update 3, you can monitor the storage provider connection and authentication status on each ESXi host. If there is an authentication issue in a specific host, you can re-authenticate the specific host against the storage provider.


Register storage providers. See [Register Storage Providers for Virtual Volumes](#).

Use the vSphere Client to monitor the storage providers.

1. Navigate to the ESXi host in vCenter Server.
2. Click the **Configure** tab, and click **Storage Providers**.

The **Storage Providers** tab lists the storage provider registered with the ESXi host. You can view general information including the name of the storage provider, status, and version of the VASA API. The **Client Certificate** lists the VASA client certificate details.

- If the status indicates an authentication error for a provider, click the row action (three dots) icon next to the provider's name and click **Re-authenticate Host VASA Clients**.



The VASA client re-authorizes the ESXi host with the storage provider.


NVMe and Virtual Volumes in vSphere

In vSphere, Virtual Volumes supports NVMe over Fibre Channel and NVMe over TCP protocols.

When you use this protocol, after you register the storage provider (VASA), the system automatically completes the underlying setup. You only create the datastore. The virtual protocol endpoints and connections are handled by the storage provider.

Note: Starting with vSphere 8.0 U2, third-party MPPs support NVMe with Virtual Volumes. For more information about MPPs, see [Multipathing Concepts and Acronyms](#).

Be familiar with the following concepts specific to the configuration of NVMe with Virtual Volumes.



Virtual Volume Namespace

Generally, a namespace can be considered an NVMe equivalent of a LUN. Each type of virtual volume, such as data, configuration, or swap virtual volume, creates a namespace with a 1:1 virtual volume to namespace ratio. The namespace limits depend on the array vendor.

Asymmetrical Namespace Access (ANA) Group

With SCSI or NFS Virtual Volumes implementation, the storage container logically groups the virtual volume objects. With NVMe, grouping of the virtual volumes depends on the array. Generally, at the array, the virtual volumes namespaces are grouped into an ANA group. The array determines the number of ANA groups, each having a unique ANAGRPID within the NVM subsystem. Virtual volume namespaces are allocated and active only on BIND request to the VASA provider. The BIND request also adds the namespace to the ANA group. A namespace remains allocated and active until the last host unbinds the virtual volume.

NVMe Virtual Protocol Endpoint

With SCSI or NFS protocols, protocol endpoints represent a proxy LUN or an NFS mount point. These protocol points are called static.

Unlike a static protocol point, the NVMe protocol endpoint is virtual. This protocol endpoint is a logical representation of the ANA group where the virtual volumes reside. Until a VM is powered on, the NVMe protocol point doesn't exist. Once the VM is powered on, the NVMe protocol point is created, so that the host can access the virtual volumes in the ANA group.

For information about viewing the NVMe virtual protocol endpoints in the vSphere Client, see [Review NVMe Protocol Endpoints](#).

Configure NVMe with Virtual Volumes

Follow these steps to configure NVMe with Virtual Volumes on your ESXi host.

- Make sure that your storage provider (VASA) supports NVMe with Virtual Volumes.
- Make sure that your storage array supports the NVMe with Virtual Volumes capability.
- Prepare your vSphere environment and a storage system. For more information, see [Before You Enable vSphere Virtual Volumes](#) and [Requirements for VMware NVMe Storage](#).
- If you use NVMe over TCP, configure adapters for NVMe over TCP storage. See [Configuring NVMe over TCP on ESXi](#).

1. Register the storage provider.

For information, see [Register Storage Providers for Virtual Volumes](#).

vSphere Client

Search in all environments

wdc-rdops-vm03-dhcp-14-22.eng.vmware.com

vcqa

Cluster

w1-hs4-l1503.eng.vmware.com

w1-hs4-l1504.eng.vmware.com

wdc-rdops-vm03-dhcp-14-22.eng.vmware.com

Summary

Monitor

Settings

General

Licensing

Message of the Day

Advanced Settings

Authentication Proxy

vCenter HA

Security

Trust Authority

Key Providers

Alarm Definitions

Scheduled Tasks

Storage Providers

vSphere Zones

vSAN

Update

Internet Connectivity

VMware by Broadcom

2655

2. Create a Virtual Volumes datastore.

When creating the datastore, select the storage containers that support NVMe capability. For information, see [Create a Virtual Volumes Datastore in vSphere Environment](#).



 nvme_vvol

Summary

Monitor



✓  wdc-rdops-vm03-dhcp-14-22.eng.vmware.com



vcqa



 datastore1



 datastore1 (1)

 nvme_vvol

 sharedvmfs4

Details



Ty

H

vi

V

A.

pl

LC

Tags



3. If you use NVMe over TCP, add a controller.

See [Add Controllers for NVMe over Fabrics](#).

Note: You do not need to configure the controller for NVMe over Fibre Channel. After you create the datastore, the system discovers controllers.



Search in all environments



▼



w1-hs4-l1504.eng.vmware.com



 w1-hs4-l1504

Monitor

Storage Adapters

I/O Filters

TCP/IP configuration

Swap File Location

Host Profile


vSphere Virtual Volumes and Replication

vSphere Virtual Volumes support replication and disaster recovery. With the array-based replication, you can off-load replication of virtual machines to your storage array and use full replication capabilities of the array. You can replicate a single VM object, such as a virtual disk. You can also group several VM objects or virtual machines to replicate them as a single unit.

Array-based replication is policy driven. After you configure your Virtual Volumes storage for replication, information about replication capabilities and replication groups is delivered from the array by the storage provider. This information shows in the VM Storage Policy interface of vCenter Server.

You use the VM storage policy to describe replication requirements for your virtual machines. The parameters that you specify in the storage policy depend on how your array implements replication. For example, your VM storage policy might include such parameters as the replication schedule, replication frequency, or recovery point objective (RPO). The policy might also indicate the replication target, a secondary site where your virtual machines are replicated, or specify whether replicas must be deleted.

By assigning the replication policy during VM provisioning, you request replication services for your virtual machine. After that, the array takes over the management of all replication schedules and processes.



Requirements for Replication with Virtual Volumes

When you enable Virtual Volumes with replication, in addition to general Virtual Volumes requirements, your environment must satisfy several specific prerequisites.

For general Virtual Volumes requirements, see [Before You Enable vSphere Virtual Volumes](#).

Storage Requirements

Implementation of Virtual Volumes replication depends on your array and might be different for storage vendors. Generally, the following requirements apply to all vendors.

- The storage arrays that you use to implement replication must be compatible with Virtual Volumes.

- The arrays must integrate with the version of the storage (VASA) provider compatible with Virtual Volumes replication.
- The storage arrays must be replication capable and configured to use vendor-provided replication mechanisms. Typical configurations usually involve one or two replication targets. Any required configurations, such as pairing of the replicated site and the target site, must be also performed on the storage side.
- When applicable, replication groups and fault domains for Virtual Volumes must be preconfigured on the storage side.

For more information, contact your vendor and see *VMware Compatibility Guide*.


vSphere Requirements

- Use the vCenter Server and ESXi versions that support Virtual Volumes storage replication. vCenter Server and ESXi hosts that are older than 6.5 release do not support replicated Virtual Volumes storage. Any attempts to create a replicated VM on an incompatible host fail with an error. For information, see *VMware Compatibility Guide*.
- If you plan to migrate a virtual machine, make sure that target resources, such as the ESXi hosts and Virtual Volumes datastores, support storage replication.

Virtual Volumes and Replication Groups

When your storage offers replication services, in addition to storage containers and protocol endpoints, your storage administrator can configure replication groups on the storage side.

vCenter Server and ESXi can discover replication groups, but do not manage their life cycle. Replication groups, also called consistency groups, indicate which VMs and virtual disks must be replicated together to a target site. A single VM cannot span multiple replication groups.




If no preconfigured groups are available, Virtual Volumes can use an automatic method. With the automatic method, Virtual Volumes creates a replication group on demand and associates this group with a Virtual Volumes object being provisioned. If you use the automatic replication group, all components of a virtual machine are assigned to the group. You cannot mix preconfigured and automatic replication groups for components of the same virtual machine.

Virtual Volumes and Fault Domains


In the Virtual Volumes environment, fault domains define how specific replication groups must be combined when being replicated from a source to a target site.

Fault domains are configured and reported by the storage array, and are not exposed in the vSphere Client. The Storage Policy Based Management (SPBM) mechanism discovers fault domains and uses them for validation purposes during a virtual machine creation.

For example, provision two virtual machines, one associated with replication group Anaheim: B, the second associated with replication group Anaheim: C. SPBM validates the provisioning because each VM is in its own replication group with corresponding target fault domains.



Now provision two VMs, one associated with replication group Anaheim: B, the second associated with replication group Anaheim: D. This configuration is invalid. Both replication groups replicate to the New-York fault domain, however, only one replicates to the Boulder fault domain.



Virtual Volumes Replication Workflow

If information about replication capabilities of the Virtual Volumes storage array shows in vCenter Server, you can activate replication for your virtual machines.

The workflow to activate replication for your virtual machines includes steps typical for the virtual machine provisioning on Virtual Volumes storage.

1. Define the VM storage policy compatible with replication storage. The datastore-based rules of the policy must include the replication component. See [Create a VM Storage Policy for Virtual Volumes](#).
After you configure the storage policy that includes replication, vCenter Server discovers available replication groups.
2. Assign the replication policy to your virtual machine. If configured, select a compatible replication group, or use the automatic assignment. See [Assign Storage Policies to Virtual Machines](#).

Replication Guidelines and Considerations

When you use replication with Virtual Volumes, specific considerations apply.

- You can apply the replication storage policy only to a configuration virtual volume and a data virtual volume. Other VM objects inherit the replication policy in the following way:
 - The memory virtual volume inherits the policy of the configuration virtual volume.
 - The digest virtual volume inherits the policy of the data virtual volume.
 - The swap virtual volume, which exists while a virtual machine is powered on, is excluded from replication.

- If you do not apply the replication policy to a VM disk, the disk is not replicated.
- The replication storage policy should not be used as a default storage policy for a datastore. Otherwise, the policy prevents you from selecting replication groups.
- Replication preserves snapshot history. If a snapshot was created and replicated, you can recover to the application consistent snapshot.
- You can replicate a linked clone. If a linked clone is replicated without its parent, it becomes a full clone.
- If a descriptor file belongs to a virtual disk of one VM, but resides in the VM home of another VM, both VMs must be in the same replication group. If the VMs are located in different replication groups, both of these replication groups must be failed over at the same time. Otherwise, the descriptor might become unavailable after the failover. As a result, the VM might fail to power on.
- In your Virtual Volumes with replication environment, you might periodically run a test failover workflow to ensure that the recovered workloads are functional after a failover.

The resulting test VMs that are created during the test failover are fully functional and suitable for general administrative operations. However, certain considerations apply:

- All VMs created during the test failover must be deleted before the test failover stops. The deletion ensures that any snapshots or snapshot-related virtual volumes that are part of the VM, such as the snapshot virtual volume, do not interfere with stopping of the test failover.
- You can create full clones of the test VMs.
- You can create fast clones only if the policy applied to the new VM contains the same replication group ID as the VM being cloned. Attempts to place the child VM outside of the replication group of the parent VM fail.

Best Practices for Working with vSphere Virtual Volumes

Observe the following recommendations when you use Virtual Volumes with ESXi and vCenter Server.

Guidelines and Limitations when Using Virtual Volumes

For the best experience with Virtual Volumes functionality, you must follow specific guidelines.

Virtual Volumes supports the following capabilities, features, and VMware products:

- With Virtual Volumes, you can use advanced storage services that include replication, encryption, deduplication, and compression on individual virtual disks. Contact your storage vendor for information about services they support with Virtual Volumes.
- Virtual Volumes functionality supports backup software that uses vSphere APIs - Data Protection. Virtual volumes are modeled on virtual disks. Backup products that use vSphere APIs - Data Protection are as fully supported on virtual volumes as they are on VMDK files on a LUN. Snapshots that the backup software creates using vSphere APIs - Data Protection look as non-vVols snapshots to vSphere and the backup software.

Note: Virtual Volumes do not support SAN transport mode. vSphere APIs - Data Protection automatically selects an alternative data transfer method.

For more information about integration with the vSphere Storage APIs - Data Protection, consult your backup software vendor.

- Virtual Volumes supports vSphere features such as vSphere vMotion, Storage vMotion, snapshots, linked clones, and DRS.
- You can use clustering products, such as Oracle Real Application Clusters, with Virtual Volumes. To use these products, you activate the multiwrite setting for a virtual disk stored on the Virtual Volumes datastore.

For more details, see the knowledge base article at <http://kb.vmware.com/kb/2112039>. For a list of features and products that Virtual Volumes functionality supports, see *VMware Product Interoperability Matrixes*.

Virtual Volumes Limitations

Improve your experience with Virtual Volumes by knowing the following limitations:

- Because the Virtual Volumes environment requires vCenter Server, you cannot use Virtual Volumes with a standalone host.
- Virtual Volumes functionality does not support RDMS.
- A Virtual Volumes storage container cannot span multiple physical arrays. Some vendors present multiple physical arrays as a single array. In such cases, you still technically use one logical array.
- Host profiles that contain Virtual Volumes datastores are vCenter Server specific. After you extract this type of host profile, you can attach it only to hosts and clusters managed by the same vCenter Server as the reference host. For more information, see *Configuration Maximums* document.

Virtual Volumes Limitations with NVMe

The following features are not supported for NVMe over Virtual Volumes specific to vSphere 8.0 release:

- Replication and Site Recovery Manager operations.
- NVMe reservations.
- NVMe deallocate and unmap operations.
- Disk sharing across VMs.
- NVMe in-band migration.

Best Practices for Storage Container Provisioning

Follow these best practices when provisioning storage containers on the Virtual Volumes array side.

Creating Containers Based on Your Limits

Because storage containers apply logical limits when grouping virtual volumes, the container must match the boundaries that you want to apply.

Examples might include a container created for a tenant in a multitenant deployment, or a container for a department in an enterprise deployment.

- Organizations or departments, for example, Human Resources and Finance
- Groups or projects, for example, Team A and Red Team
- Customers

Putting All Storage Capabilities in a Single Container

Storage containers are individual datastores. A single storage container can export multiple storage capability profiles. As a result, virtual machines with diverse needs and different storage policy settings can be a part of the same storage container.

Changing storage profiles must be an array-side operation, not a storage migration to another container.

Avoiding Over-Provisioning Your Storage Containers

When you provision a storage container, the space limits that you apply as part of the container configuration are only logical limits. Do not provision the container larger than necessary for the anticipated use. If you later increase the size of the container, you do not need to reformat or repartition it.

Using Storage-Specific Management UI to Provision Protocol Endpoints

Every storage container needs protocol endpoints (PEs) that are accessible to ESXi hosts.

When you use block storage, the PE represents a proxy LUN defined by a T10-based LUN WWN. For NFS storage, the PE is a mount point, such as an IP address or DNS name, and a share name. For NVMe storage, the PE is an ANA group and proxy namespaces are virtual.

Typically, configuration of PEs is array-specific. When you configure PEs, you might need to associate them with specific storage processors, or with certain hosts. To avoid errors when creating PEs, do not configure them manually. Instead, when possible, use storage-specific management tools.

No Assignment of IDs Above Disk.MaxLUN to Protocol Endpoint LUNs

By default, an ESXi host can access LUN IDs that are within the range of 0 to 1023. If the ID of the protocol endpoint LUN that you configure is 1024 or greater, the host might ignore the PE.

If your environment uses LUN IDs that are greater than 1023, change the number of scanned LUNs through the `Disk.MaxLUN` parameter. See [Change the Number of Scanned Storage Devices](#).

Best Practices for Virtual Volumes Performance

To ensure optimal Virtual Volumes performance results, follow these recommendations.

Using Different VM Storage Policies for Individual Virtual Volume Components

By default, all components of a virtual machine in the Virtual Volumes environment get a single VM storage policy. However, different components might have different performance characteristics, for example, a database virtual disk and a corresponding log virtual disk. Depending on performance requirements, you can assign different VM storage policies to individual virtual disks and to the VM home file, or config-vVol.

When you use the vSphere Client, you cannot change the VM storage policy assignment for swap-vVol, memory-vVol, or snapshot-vVol.

See [Create a VM Storage Policy for Virtual Volumes](#).

Getting a Host Profile with Virtual Volumes

The best way to get a host profile with Virtual Volumes is to configure a reference host and extract its profile. If you manually edit an existing host profile in the vSphere Client and attach the edited profile to a new host, you might trigger compliance errors. Other unpredictable problems might occur. For more details, see the [VMware Knowledge Base article 2146394](#).

Monitoring I/O Load on Individual Protocol Endpoint

- All virtual volume I/O goes through protocol endpoints (PEs). Arrays select protocol endpoints from several PEs that are accessible to an ESXi host. Arrays can do load balancing and change the binding path that connects the virtual volume and the PE. See [Binding and Unbinding Virtual Volumes](#).
- On block storage, ESXi gives a large queue depth to I/O because of a potentially high number of virtual volumes. The `Scsi.ScsiVVolPESNRO` parameter controls the number of I/O that can be queued for PEs. You can configure the parameter on the Advanced System Settings page of the vSphere Client.

Monitoring Array Limitations

A single VM might occupy multiple virtual volumes. See [Virtual Volume Objects](#).

Suppose that your VM has two virtual disks, and you take two snapshots with memory. Your VM might occupy up to 10 Virtual Volumes objects: a config-vVol, a swap-vVol, two data-vVols, four snapshot-vVols, and two memory snapshot-vVols.

Ensuring that Storage Provider Is Available

To access Virtual Volumes storage, your ESXi host requires a storage provider (VASA provider). It is either built into the storage array firmware or deployed as an appliance in the form of a VM. If your VASA provider is a VM based appliance, the storage provider is always available. To ensure that the storage provider is always available, follow these guidelines:

- Do not migrate a storage provider VM to Virtual Volumes storage.
- Back up your storage provider VM.
- When appropriate, use vSphere HA or Site Recovery Manager to protect the storage provider VM.

Using Stretched Storage Clustering with Virtual Volumes

Starting with vSphere 8.0 Update 3, Virtual Volumes support Stretched Storage Clustering (SSC) through SCSI.

SSC is a specific storage configuration that is commonly referred to as Metro Storage Clustering. A metro cluster is a host cluster that spans multiple sites located in different locations (usually two). The different sites that make up these types of clusters are almost always located in the same area if not city due to the latency requirements. The hosts may be in different data centers but still access the same shared storage. SSC configurations allow fully active and workload-balanced data centers to be used to their full potential and for fast recovery in the event of a host or even full site failure. These configurations are usually implemented in environments where disaster and downtime avoidance is a key requirement.

Virtual Volumes support the uniform host access SSC configuration where vSphere hosts from both sites are all connected to a storage node in the storage cluster across all sites. For information about specific requirements and best practices for an SSC implementation, see <https://core.vmware.com/vsphere-virtual-volumes-vvols>.

Using the `esxcli storage vvol` Command to Troubleshoot vSphere Virtual Volumes

Learn about the command options that the `esxcli storage vvol` takes. The topics also provide solutions to problems that you might encounter when using vSphere Virtual Volumes.

Virtual Volumes and `esxcli` Commands

You can use the `esxcli storage vvol` commands to troubleshoot your Virtual Volumes environment.

The following command options are available.

Table 226: `esxcli storage vvol` commands

Namespace	Command Option	Description
<code>esxcli storage core device</code>	<code>list</code>	Identify protocol endpoints. The output entry <code>Is VVOL PE: true</code> indicates that the storage device is a protocol endpoint.
<code>esxcli storage vvol nvme info</code>	<code>get</code>	Get information for the Virtual Volumes NVMe host information.
<code>esxcli storage vvol storagearray</code>	<code>list</code>	List containers or protocol endpoints on the storage array.
<code>esxcli storage vvol storagecontainer</code>	<code>list</code>	List all available storage containers.
	<code>abandonedvvol scan</code>	Scan the specified storage container for abandoned Virtual Volumes.

Namespace	Command Option	Description
esxcli storage vvol daemon	unbindall	Unbind all virtual volumes from all VASA providers known to the ESXi host.
esxcli storage vvol protocolendpoint	list	List all protocol endpoints that your host can access.
	startforceunbind	Unbind all Virtual Volumes on the PE with zero open handle count.
esxcli storage vvol stats		Track performance statistics. For information, see Collecting Statistical Information for Virtual Volumes .
esxcli storage vvol vasacontext	get	Show the VASA context (VC UUID) associated with the host.
esxcli storage vvol vasaprovider	list	List all storage (VASA) providers associated with the host.
esxcli storage vvol virtualvolume	unmap	Reclaim the space by unmapping free blocks from VMFS config volumes. For information, see Reclaim Space on the vSphere Virtual Volumes Datastores .
esxcli storage vvol vmstats	get	Get the Virtual Volumes information and statistics for a specific virtual machine.

Collecting Statistical Information for Virtual Volumes

You can use the `vvol stats` command in your ESXi host to track performance statistics.

The following command options are available.

Command	Description	Options
esxcli storage vvol stats get	Get statistics for all VASA providers (default), or for specified namespace or entity in the given namespace.	-e --entity= <i>str</i> Enter entity ID. -n --namespace= <i>str</i> Enter node namespace expression. -r --raw Use raw format output.
esxcli storage vvol stats list	List all the statistics nodes (default), or nodes under a specified namespace.	-n --namespace= <i>str</i> Enter node namespace expression.
esxcli storage vvol stats enable	Activate statistics tracking for the complete namespace.	
esxcli storage vvol stats disable	Deactivate statistics tracking for the complete namespace.	
esxcli storage vvol stats add	Activate statistics tracking for a specific entity under a specific namespace.	-e --entity= <i>str</i> Enter entity ID. -n --namespace= <i>str</i> Enter node namespace expression.
esxcli storage vvol stats remove	Removes specific entity for statistics tracking under the specified namespace.	-e --entity= <i>str</i> Enter entity ID. -n --namespace= <i>str</i> Enter node namespace expression.
esxcli storage vvol stats reset	Reset the statistics counter for the specified statistics namespace or entity.	-e --entity= <i>str</i> Enter entity ID. -n --namespace= <i>str</i> Enter node namespace expression.

Virtual Volumes Datastore Is Inaccessible

After you create a Virtual Volumes datastore, it remains inaccessible.

The vSphere Client shows the datastore as inaccessible. You cannot use the datastore for virtual machine provisioning.

This problem might occur when you fail to configure protocol endpoints for the SCSI-based storage container or NVMe ANA groups mapped to the virtual datastore. Like traditional LUNs, SCSI protocol endpoints need to be configured so that an ESXi host can detect them.

Before creating virtual datastores for SCSI-based containers, make sure to configure protocol endpoints on the storage side.

Failures When Migrating VMs or Deploying VM OVF to Virtual Volumes Datastores

Your attempts to migrate a virtual machine or to deploy a VM OVF to Virtual Volumes datastores fail.

An OVF template or a VM being migrated from a nonvirtual datastore might include additional large files, such as ISO disk images, DVD images, and image files. If these additional files cause the configuration virtual volume to exceed its 4-GB limit, migration or deployment to a virtual datastore fails.

The configuration virtual volume, or config-vVol, contains various VM-related files. On traditional nonvirtual datastores, these files are stored in the VM home directory. Similar to the VM home directory, the config-vVol typically includes the VM configuration file, virtual disk and snapshot descriptor files, log files, lock files, and so on.

On virtual datastores, all other large-sized files, such as virtual disks, memory snapshots, swap, and digest, are stored as separate virtual volumes.

Config-vVols are created as 4-GB virtual volumes. Generic content of the config-vVol usually consumes only a fraction of this 4-GB allocation, so config-vVols are typically thin-provisioned to conserve backing space. Any additional large files, such as ISO disk images, DVD images, and image files, might cause the config-vVol to exceed its 4-GB limit. If such files are included in an OVF template, deployment of the VM OVF to Virtual Volumes storage fails. If these files are part of an existing VM, migration of that VM from a traditional datastore to Virtual Volumes storage also fails.

Note: Starting with vSphere 7.0 Update 2, partners can increase the config-vVol to above 4 GB. Work with your Virtual Volumes partner on implementing this if it is supported by your partner and applicable to your environment.

- For VM migration. Before migrating a VM from a traditional datastore to a virtual datastore, remove excess content from the VM home directory to keep the config-vVol under the 4-GB limit.
- For OVF deployment. Because you cannot deploy an OVF template that contains excess files directly to a virtual datastore, first deploy the VM to a nonvirtual datastore. Remove any excess content from the VM home directory, and migrate the resulting VM to Virtual Volumes storage.

Filtering Virtual Machine I/O in vSphere

I/O filters are software components that can be installed on ESXi hosts and can offer additional data services to virtual machines. The filters process I/O requests, which move between the guest operating system of a virtual machine and virtual disks.

The I/O filters can be offered by VMware or created by third parties through vSphere APIs for I/O Filtering (VAIO).

About I/O Filters

I/O filters can gain direct access to the virtual machine I/O path. You can enable the I/O filter for an individual virtual disk. The I/O filters are independent of the storage topology.

VMware offers certain categories of I/O filters. In addition, third-party vendors can create the I/O filters. Typically, they are distributed as packages that provide an installer to deploy the filter components on vCenter Server and ESXi host clusters.

After the I/O filters are deployed, vCenter Server configures and registers an I/O filter storage provider, also called a VASA provider, for each host in the cluster. The storage providers communicate with vCenter Server and make data services offered by the I/O filter visible in the VM Storage Policies interface. You can reference these data services when defining common rules for a VM policy. After you associate virtual disks with this policy, the I/O filters are enabled on the virtual disks.

Datastore Support

I/O filters can support all datastore types including the following:

- VMFS
- NFS 3
- NFS 4.1
- vVol
- vSAN

System Requirements for I/O Filters

To be able to use I/O filters in your environment, you must follow specific requirements.

- Use the latest version of ESXi and vCenter Server compatible with I/O filters. Older versions might not support I/O filters, or provide only partial support.
- Check for any additional requirements that individual partner solutions might have. In specific cases, your environment might need flash devices, extra physical memory, or network connectivity and bandwidth. For information, contact your vendor or your VMware representative.
- Web server to host partner packages for filter installation. The server must remain available after initial installation. When a new host joins the cluster, the server pushes appropriate I/O filter components to the host.

I/O Filter Guidelines and Best Practices


When you use I/O filters in your environment, follow specific guidelines and best practices.

- Because I/O filters are datastore-agnostic, all types of datastores, including VMFS, NFS, Virtual Volumes, and vSAN, are compatible with I/O filters.
- I/O filters support RDMs in virtual compatibility mode. No support is provided to RDMs in physical compatibility mode.
- You cannot change or assign the I/O filter policy while migrating or cloning a virtual machine. You can change the policy after you complete the migration or cloning.
- When you clone or migrate a virtual machine with I/O filter policy from one host to another, make sure that the destination host has a compatible filter installed. This requirement applies to migrations initiated by an administrator or by such functionalities as HA or DRS.
- When you convert a template to a virtual machine, and the template is configured with I/O filter policy, the destination host must have the compatible I/O filter installed.
- If you use vCenter Site Recovery Manager to replicate virtual disks, the resulting disks on the recovery site do not have the I/O filter policies. You must create the I/O filter policies in the recovery site and reattach them to the replicated disks.
- If your virtual machine has a snapshot tree associated with it, you cannot add, change, or remove the I/O filter policy for the virtual machine.

Using Cache I/O Filters with Flash Storage Devices on ESXi Hosts

A cache I/O filter can use a local flash device on your ESXi host to cache virtual machine data.

If your caching I/O filter uses local flash devices, you need to configure a virtual flash resource, also known as VFFS volume. You configure the resource on your ESXi host before activating the filter. While processing the virtual machine read I/Os, the filter creates a virtual machine cache and places it on the VFFS volume.



To set up a virtual flash resource, you use flash devices that are connected to your host. To increase the capacity of your virtual flash resource, you can add more flash drives. An individual flash drive must be exclusively allocated to a virtual flash resource and cannot be shared with any other vSphere service, such as vSAN or VMFS. See [Set Up Virtual Flash Resource](#).

Deploying and Managing I/O Filters in vSphere

You can run the installer provided by your vendor to install, uninstall, or upgrade I/O filters in your vSphere environment.

When you work with I/O filters, the following considerations apply:

- vCenter Server uses ESX Agent Manager (EAM) to install and uninstall I/O filters. As an administrator, never invoke EAM APIs directly for EAM agencies that are created or used by vCenter Server. All operations related to I/O filters must go through VIM APIs. If you accidentally modify an EAM agency that was created by vCenter Server, you must revert the changes. If you accidentally destroy an EAM agency that is used by I/O filters, you must call `Vim.IoFilterManager#uninstallIoFilter` to uninstall the affected I/O filters. After uninstalling, perform a fresh reinstall.

- When a new host joins the cluster that has I/O filters, the filters installed on the cluster are deployed on the host. vCenter Server registers the I/O filter storage provider for the host. Any cluster changes become visible in the VM Storage Policies interface of the vSphere Client.
- When you move a host out of a cluster or remove it from vCenter Server, the I/O filters are uninstalled from the host. vCenter Server unregisters the I/O filter storage provider.
- If you use a stateless ESXi host, it might lose its I/O filter VIBs during a reboot. vCenter Server checks the bundles installed on the host after it reboots, and pushes the I/O filter VIBs to the host if necessary.

Configure I/O Filters in the vSphere Environment

To set up data services that the I/O filters provide for your virtual machines, follow several steps.

- Create a cluster that includes at least one ESXi host.
- For information about I/O filters provided by third parties, contact your vendor or your VMware representative.

Install I/O Filters in a Cluster

If you use I/O filters provided by third parties, install the I/O filters in an ESXi host cluster.

- Required privileges: **Host.Configuration.Query patch**.
- Verify that the I/O filter solution is certified by VMware.

VMware partners create I/O filters through the vSphere APIs for I/O Filtering (VAIO) developer program.

The filter packages are distributed as solution bundle ZIP packages that can include I/O filter daemons, I/O filter libraries, CIM providers, and other associated components.

Typically, to deploy the filters, you run installers provided by vendors. Installation is performed at an ESXi cluster level. You cannot install the filters on selected hosts directly.

Run the installer that the vendor provided.

The installer deploys the appropriate I/O filter extension on vCenter Server and the filter components on all hosts within a cluster.

A storage provider, also called a VASA provider, is automatically registered for every ESXi host in the cluster. Successful auto-registration of the I/O filter storage providers triggers an event at the host level. If the storage providers fail to auto-register, the system raises alarms on the hosts.

View I/O Filters and Storage Providers

Use the vSphere Client to review I/O filters available in your environment and verify that the I/O filter providers appear as expected and are active.

When you install a third-party I/O filter, a storage provider, also called VASA provider, is automatically registered for every ESXi host in the cluster. Successful auto-registration of the I/O filter storage providers triggers an event at the host level. If the storage providers fail to auto-register, the system raises alarms on the hosts.

1. Verify that the I/O filter storage providers appear as expected and are active.
 - a) Navigate to vCenter Server.
 - b) Click the **Configure** tab, and click **Storage Providers**.
 - c) Review the storage providers for I/O filters.

When the I/O filter providers are properly registered, capabilities and data services that the filters offer populate the VM Storage Policies interface.

2. Verify that the I/O filter components are listed on your cluster and ESXi hosts.

Option	Actions
View I/O filters on a cluster	1. Navigate to the cluster.

Option	Actions
	- Click the Configure tab. - Under Configuration, click I/O Filters to review the filters installed in the cluster.
View I/O filters on a host	- Navigate to the host. - Click the Configure tab. - Under Storage, click I/O Filters to review the filters installed on the host.

Handling I/O Filter Installation Failures

Typically, all ESXi hosts in a cluster have the same set of I/O filters installed. Occasionally, failures might happen during installation.

If an I/O filter installation fails on a host, the system generates events that report the failure. In addition, an alarm on the host shows the reason for the failure. Examples of failures include the following:

- The VIB URL is not accessible from the host.
- The VIB has an invalid format.
- The VIB requires the host to be in maintenance mode for an upgrade or uninstallation.
- The VIB requires the host to reboot after the installation or uninstallation.
- Attempts to put the host in maintenance mode fail because the virtual machine cannot be evacuated from the host.
- The VIB requires manual installation or uninstallation.

vCenter Server can resolve some failures. You might have to intervene for other failures. For example, you might need to edit the VIB URL, manually evacuate or power off virtual machines, or manually install or uninstall VIBs.

Install I/O Filters on a Single ESXi Host

For troubleshooting purposes, you can download an ESXi component of the I/O filter, packaged as a VIB file, and install it on the ESXi host. Use the `esxcli` command to install the VIB file.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

- Install the VIBs by running the following command:

```
esxcli software vib install --depot path_to_VMware_vib_ZIP_file
```

Options for the `install` command allow you to perform a dry run, specify a specific VIB, bypass acceptance-level verification, and so on. Do not bypass verification on production systems. See the *ESXCLI Reference* documentation.

- Verify that the VIBs are installed on your ESXi host.

```
esxcli software vib list
```

Uninstall I/O Filters from a Cluster

You can uninstall I/O filters deployed in an ESXi host cluster.

- Required privileges: **Host.Config.Patch**.

- Uninstall the I/O filter by running the installer that your vendor provides.

During uninstallation, a third party I/O filter installer automatically places the hosts into maintenance mode.

If the uninstallation is successful, the filter and any related components are removed from the hosts.

2. Verify that the I/O filter components are properly uninstalled from your ESXi hosts. Use one of the following methods:

- Run the `esxcli software vib list` command.
- View the I/O filters in the vSphere Client. See [View I/O Filters and Storage Providers](#).

The uninstalled filter no longer appears on the list.

Upgrade I/O Filters in a Cluster

After you upgrade your ESXi hosts, use installers provided by I/O filter vendors to upgrade I/O filters deployed in the ESXi host cluster.

- Required privileges: **Host.Config.Patch**.
- Upgrade your hosts to ESXi 8.0. If you use vSphere Lifecycle Manager for the upgrade, see the *Managing Host and Cluster Lifecycle* documentation.

An upgrade consists of uninstalling the old filter components and replacing them with the new filter components. To determine whether an installation is an upgrade, vCenter Server checks the names and versions of existing filters. If the existing filter names match the names of the new filters but have different versions, the installation is considered an upgrade.

1. To upgrade the filter, run the vendor-provided installer.

During the upgrade, a third party I/O filter installer automatically places the hosts into maintenance mode.

The installer identifies any existing filter components and removes them before installing the new filter components.

2. Verify that the I/O filter components are properly upgraded in your ESXi hosts. Use one of the following methods:

- Run the `esxcli software vib list` command.
- View the I/O filters in the vSphere Client. See [View I/O Filters and Storage Providers](#).

After the upgrade, the system places the hosts back into operational mode.

Activate I/O Filter Data Services on Virtual Disks in vSphere

In vSphere, activating data services that I/O filters provide is a two-step process. You create a virtual machine policy based on data services that the I/O filters provide, and then attach this policy to a virtual machine.

For the cache I/O filters, configure the virtual flash resource on your ESXi host before activating the filter. See [Set Up Virtual Flash Resource](#).

For information about cache I/O filters, see [Using Cache I/O Filters with Flash Storage Devices on ESXi Hosts](#).

1. Define a VM policy based on I/O filter services.

Make sure that the virtual machine policy lists data services provided by the I/O filters.

See [Create a VM Storage Policy for Host-Based Data Services](#).

2. Assign the I/O filter policy to a virtual machine.

To activate data services that the I/O filter provides, associate the I/O filter policy with virtual disks. You can assign the policy when you provision the virtual machine.

See [Assign the I/O Filter Policy to Virtual Machines](#).

Assign the I/O Filter Policy to Virtual Machines

To activate data services that I/O filters provide, associate the I/O filter policy with virtual disks. You can assign the policy when you create or edit a virtual machine.

Verify that the I/O filter is installed on the ESXi host where the virtual machine runs.

You can assign the I/O filter policy during an initial deployment of a virtual machine. This topic describes how to assign the policy when you create a new virtual machine. For information about other deployment methods, see the *vSphere Virtual Machine Administration* documentation.

Note: You cannot change or assign the I/O filter policy when migrating or cloning a virtual machine.

1. Start the virtual machine provisioning process and follow the appropriate steps.
2. Assign the same storage policy to all virtual machine files and disks.
 - a) On the **Select storage** page, select a storage policy from the **VM Storage Policy** drop-down menu.
 - b) Select the datastore from the list of compatible datastores and click **Next**.
 The datastore becomes the destination storage resource for the virtual machine configuration file and all virtual disks. The policy also activates I/O filter services for the virtual disks.
3. Change the VM storage policy for the virtual disk.
 Use this option to enable I/O filters just for your virtual disks.
 - a) On the **Customize hardware** page, expand the **New hard disk** pane.
 - b) From the **VM storage policy** drop-down menu, select the storage policy to assign to the virtual disk.
 - c) Optional: Change the storage location of the virtual disk.
 Use this option to store the virtual disk on a datastore other than the datastore where the VM configuration file resides.
4. Complete the virtual machine provisioning process.

After you create the virtual machine, the **Summary** tab displays the assigned storage policies and their compliance status. You can later change the virtual policy assignment. See [Change Storage Policy Assignment for Virtual Machine Files and Disks](#).

Migrating Virtual Machines with I/O Filters

When you migrate a virtual machine with I/O filters, specific considerations apply.

If you use Storage vMotion to migrate a virtual machine with I/O filters, a destination datastore must be connected to hosts with compatible I/O filters installed.

You might need to migrate a virtual machine with I/O filters across different types of datastores, for example between VMFS and Virtual Volumes. If you do so, make sure that the VM storage policy includes rule sets for every type of datastore you are planning to use. For example, if you migrate your virtual machine between the VMFS and Virtual Volumes datastores, create a mixed VM storage policy that includes the following rules:

- Common Rules for the I/O filters
- Rule Set 1 for the VMFS datastore. Because Storage Policy Based Management does not offer an explicit VMFS policy, the rule set must include tag-based rules for the VMFS datastore.
- Rule Set 2 for the Virtual Volumes datastore

When Storage vMotion migrates the virtual machine, the correct rule set that corresponds to the target datastore is selected. The I/O filter rules remain unchanged.

If you do not specify rules for datastores and define only Common Rules for the I/O filters, the system applies default storage policies for the datastores.

Storage Hardware Acceleration in vSphere

The hardware acceleration functionality enables the ESXi host to integrate with compliant storage systems. The host can offload certain virtual machine and storage management operations to the storage systems. With the storage hardware assistance, your host performs these operations faster and consumes less CPU, memory, and storage fabric bandwidth.

Block storage devices, Fibre Channel and iSCSI, and NAS devices support the hardware acceleration.

Hardware Acceleration Benefits

When the hardware acceleration functionality is supported, the host can get hardware assistance and perform the following tasks faster and more efficiently.

- Migrating virtual machines with Storage vMotion
- Deploying virtual machines from templates
- Cloning virtual machines or templates
- VMFS clustered locking and metadata operations for virtual machine files
- Provisioning thick virtual disks
- Creating fault-tolerant virtual machines
- Creating and cloning thick disks on NFS datastores

Hardware Acceleration Requirements

The hardware acceleration functionality works only if you use an appropriate host and storage array combination.

Table 227: Hardware Acceleration Storage Requirements

ESXi	Block Storage Devices	NAS Devices
ESXi	Support T10 SCSI standard, or block storage plug-ins for array integration (VAAI)	Support NAS plug-ins for array integration

Note: If your SAN or NAS storage fabric uses an intermediate appliance in front of a storage system that supports hardware acceleration, the intermediate appliance must also support hardware acceleration and be properly certified. The intermediate appliance might be a storage virtualization appliance, I/O acceleration appliance, encryption appliance, and so on.

Hardware Acceleration Support Status

For each storage device and datastore, the vSphere Client display the hardware acceleration support status.

The status values are Unknown, Supported, and Not Supported. The initial value is Unknown.

For block devices, the status changes to Supported after the host successfully performs the offload operation. If the offload operation fails, the status changes to Not Supported. The status remains Unknown if the device provides partial hardware acceleration support.

With NAS, the status becomes Supported when the storage can perform at least one hardware offload operation.

When storage devices do not support or provide partial support for the host operations, your host reverts to its native methods to perform unsupported operations.

Hardware Acceleration Considerations

When you use the hardware acceleration functionality with ESXi, certain considerations apply.

Several reasons might cause a hardware-accelerated operation to fail.

For any primitive that the array does not implement, the array returns an error. The error triggers the ESXi host to attempt the operation using its native methods.

The VMFS data mover does not leverage hardware offloads and instead uses software data movement when one of the following occurs:

- The source and destination VMFS datastores have different block sizes.
- The source file type is RDM and the destination file type is non-RDM (regular file).
- The source VMDK type is eagerzeroedthick and the destination VMDK type is thin.
- The source or destination VMDK is in sparse or hosted format.
- The source virtual machine has a snapshot.
- The logical address and transfer length in the requested operation are not aligned to the minimum alignment required by the storage device. All datastores created with the vSphere Client are aligned automatically.
- The VMFS has multiple LUNs or extents, and they are on different arrays.

Hardware cloning between arrays, even within the same VMFS datastore, does not work.

vSphere Hardware Acceleration for Block Storage Devices

With hardware acceleration, your ESXi host can integrate with block storage devices, Fibre Channel, or iSCSI, and use storage array operations that include full copy, block zeroing, and hardware assisted locking.

ESXi hardware acceleration supports the following array operations:

- Full copy, also called clone blocks or copy offload. Storage arrays can make full copies of data within the array without having the host read and write the data. This operation reduces the time and network load when cloning virtual machines, provisioning from a template, or migrating with vMotion.
- Block zeroing, also called write same. Storage arrays can zero out a large number of blocks to provide newly allocated storage, free of previously written data. This operation reduces the time and network load when creating virtual machines and formatting virtual disks.
- Hardware assisted locking, also called atomic test and set (ATS). Supports discrete virtual machine locking without use of SCSI reservations. This operation allows disk locking per sector, instead of the entire LUN as with SCSI reservations.

Check with your vendor for the hardware acceleration support. Certain storage arrays require that you activate the support on the storage side.

On your host, the hardware acceleration is available by default. If your storage does not support the hardware acceleration, you can deactivate it.

In addition to hardware acceleration support, ESXi includes support for array thin provisioning. For information, see [ESXi and Array Thin Provisioning](#).

Deactivate Hardware Acceleration for Block Storage Devices

On your host, the hardware acceleration for block storage devices is activated by default. You can use the vSphere Client advanced settings to deactivate the hardware acceleration operations.

As with any advanced settings, before you deactivate the hardware acceleration, consult with the VMware support team.

1. In the vSphere Client, navigate to the ESXi host.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. To deactivate the hardware acceleration, change the value for these options to 0:
 - VMFS3.HardwareAcceleratedLocking
 - DataMover.HardwareAcceleratedMove
 - DataMover.HardwareAcceleratedInit

Managing Hardware Acceleration on Block Storage Devices

To integrate with the block storage arrays, vSphere uses the ESXi extensions called Storage APIs - Array Integration (VAAI). With this integration, vSphere can use the array hardware operations.

In the vSphere 5.x and later releases, these extensions are implemented as the T10 SCSI commands. As a result, with the devices that support the T10 SCSI standard, your ESXi host can communicate directly and does not require the VAAI plug-ins.

If the device does not support T10 SCSI or provides partial support, ESXi reverts to using the VAAI plug-ins, installed on your host. The host can also use a combination of the T10 SCSI commands and plug-ins. The VAAI plug-ins are vendor-specific and can be either VMware or partner developed. To manage the VAAI capable device, your host attaches the VAAI filter and vendor-specific VAAI plug-in to the device.

For information about whether your storage requires VAAI plug-ins or supports hardware acceleration through T10 SCSI commands, see the *VMware Compatibility Guide* or contact your storage vendor.

You can use several `esxcli` commands to query storage devices for the hardware acceleration support information. For the devices that require the VAAI plug-ins, the claim rule commands are also available. For information about `esxcli` commands, see *Getting Started with ESXCLI*.

Display Hardware Acceleration Plug-Ins and Filter

To communicate with the devices that do not support the T10 SCSI standard, your host uses a single VAAI filter and a vendor-specific VAAI plug-in. Use the `esxcli` command to view the hardware acceleration filter and plug-ins currently loaded into your system.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Run the `esxcli storage core plugin list --plugin-class=value` command.

For *value*, enter one of the following parameters:

- Type `VAAI` to display plug-ins.

The output of this command is similar to the following example:

```
#esxcli storage core plugin list --plugin-class=VAAI
Plugin name      Plugin class
VMW_VAAIP_EQL    VAAI
VMW_VAAIP_NETAPP VAAI
VMW_VAAIP_CX     VAAI
```

- Type `Filter` to display the Filter.

The output of this command is similar to the following example:

```
esxcli storage core plugin list --plugin-class=Filter
Plugin name  Plugin class
VAAI_FILTER  Filter
```

Verify Hardware Acceleration Support Status

Use the `esxcli` command to verify the hardware acceleration support status of a particular storage device.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Run the `esxcli storage core device list -d=device_ID` command.

The output shows the hardware acceleration, or VAAI, status that can be unknown, supported, or unsupported.

```
# esxcli storage core device list -d naa.XXXXXXXXXXXXX4c
naa.XXXXXXXXXXXXX4c
Display Name: XXXX Fibre Channel Disk(naa.XXXXXXXXXXXXX4c)
Size: 20480
Device Type: Direct-Access
Multipath Plugin: NMP
XXXXXXXXXXXXXXXXXX
Attached Filters: VAAI_FILTER
VAAI Status: supported
XXXXXXXXXXXXXXXXXX
```

Verify Hardware Acceleration Support Details

Use the `esxcli` command to query whether the block storage device provides the hardware acceleration support.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Run the `esxcli storage core device vaa1 status get -d=device_ID` command.

If a VAAI plug-in manages the device, the output shows the name of the plug-in attached to the device. The output also shows the support status for each T10 SCSI based primitive, if available. Output appears in the following example:

```
# esxcli storage core device vaa1 status get -d naa.XXXXXXXXXXXXX4c
naa.XXXXXXXXXXXXX4c
VAAI Plugin Name: VMW_VAAIP_SYMM
ATS Status: supported
Clone Status: supported
Zero Status: supported
Delete Status: unsupported
```

List Hardware Acceleration Claim Rules

Each block storage device managed by a VAAI plug-in needs two claim rules. One claim rule specifies the hardware acceleration filter and another specifies the hardware acceleration plug-in for the device. You can use the `esxcli` commands to list the hardware acceleration filter and plug-in claim rules.

1. To list the filter claim rules, run the `esxcli storage core claimrule list --claimrule-class=Filter` command.

In this example, the filter claim rules specify devices that the VAAI_FILTER filter claims.

```
# esxcli storage core claimrule list --claimrule-class=Filter
Rule Class  Rule  Class  Type  Plugin  Matches  XCOPY Use Array Reported Values
XCOPY Use Multiple Segments  XCOPY Max Transfer Size KiB
```

Filter	65430	runtime	vendor	VAAI_FILTER	vendor=EMC	model=SYMMETRIX	False
		False				0	
Filter	65430	file	vendor	VAAI_FILTER	vendor=EMC	model=SYMMETRIX	False
		False				0	
Filter	65431	runtime	vendor	VAAI_FILTER	vendor=DGC	model=*	False
		False				0	
Filter	65431	file	vendor	VAAI_FILTER	vendor=DGC	model=*	False
		False				0	

- To list the VAAI plug-in claim rules, run the `esxcli storage core claimrule list --claimrule-class=VAAI` command.

In this example, the VAAI claim rules specify devices that the VAAI plug-in claims.

```
esxcli storage core claimrule list --claimrule-class=VAAI
```

Rule	Class	Rule	Class	Type	Plugin	Matches	XCOPY Use Array	Reported Val-
ues	XCOPY Use	Multiple	Segments	XCOPY Max	Transfer	Size KiB		
VAAI	65430	runtime	vendor	VMW_VAAIP_SYMM	vendor=EMC	model=SYMMETRIX		
False				False		0		
VAAI	65430	file	vendor	VMW_VAAIP_SYMM	vendor=EMC	model=SYMMETRIX		
False				False		0		
VAAI	65431	runtime	vendor	VMW_VAAIP_CX	vendor=DGC	model=*		
False				False		0		
VAAI	65431	file	vendor	VMW_VAAIP_CX	vendor=DGC	model=*		
False				False		0		

Add Hardware Acceleration Claim Rules

To configure the hardware acceleration for a new array, add two claim rules, one for the VAAI filter and another for the VAAI plug-in. For the new claim rules to be active, you first define the rules and then load them into your system.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

- Define a new claim rule for the VAAI filter by running the `esxcli storage core claimrule add --claimrule-class=Filter --plugin=VAAI_FILTER` command.
- Define a new claim rule for the VAAI plug-in by running the `esxcli storage core claimrule add --claimrule-class=VAAI` command.
- Load both claim rules by running the following commands:

```
esxcli storage core claimrule load --claimrule-class=Filter
esxcli storage core claimrule load --claimrule-class=VAAI
```

- Run the VAAI filter claim rule by using the `esxcli storage core claimrule run --claimrule-class=Filter` command.

Note: Only the filter-class rules must be run. When the VAAI filter claims a device, it automatically finds the proper VAAI plug-in to attach.

Defining Hardware Acceleration Claim Rules

This example shows how to configure the hardware acceleration for IBM arrays using the `VMW_VAAIP_T10` plug-in. Use the following sequence of commands. For information about the options that the command takes, see [Add Multipathing Claim Rules](#).

```
# esxcli storage core claimrule add --claimrule-class=Filter --plugin=VAAI_FILTER
--type=vendor --vendor=IBM --autoassign
```

```
# esxcli storage core claimrule add --claimrule-class=VAAI --plugin=VMW_VAAIP_T10
--type=vendor --vendor=IBM --autoassign

# esxcli storage core claimrule load --claimrule-class=Filter

# esxcli storage core claimrule load --claimrule-class=VAAI

# esxcli storage core claimrule run --claimrule-class=Filter
```

Delete Hardware Acceleration Claim Rules

Use the `esxcli` command to delete existing hardware acceleration claim rules.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Run the following commands:

```
esxcli storage core claimrule remove -r claimrule_ID --claimrule-class=Filter
esxcli storage core claimrule remove -r claimrule_ID --claimrule-class=VAAI
```

Configure XCOPY Parameters

XCOPY is one of the VAAI primitives that is used for offloading tasks to the storage array. For example, you can use XCOPY to offload such operations as migration or cloning of virtual machines to the array instead of consuming vSphere resources to perform these tasks.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

You can use the XCOPY mechanism with all storage arrays that support the SCSI T10 based VMW_VAAIP_T10 plug-in developed by VMware. To enable the XCOPY mechanism, create a claim rule of the VAAI class.

Use the following command and enter the XCOPY options:

```
esxcli storage core claimrule add --claimrule-class=VAAI
```

For information about the options that the command takes, see [Add Multipathing Claim Rules](#).

Option	Description
-a --xcopy-use-array-values	Use array reported values for XCOPY commands.
-s --xcopy-use-multi-segs	Use multiple segments for XCOPY commands. Valid only when <code>--xcopy-use-array-values</code> is specified.
-m --xcopy-max-transfer-size	Maximum transfer size in MB for the XCOPY commands when you use a transfer size different than array reported. Valid only when <code>--xcopy-use-array-values</code> is specified.
-k --xcopy-max-transfer-size-kib	Maximum transfer size in KiB for the XCOPY commands when you use a transfer size different than array reported. Valid only if <code>--xcopy-use-array-values</code> is specified.
-e --extended-xcopy-support	Use the extended XCOPY command to optimize the data copy between datastores that are in different storage arrays.

Configuring XCOPY

- ```
esxcli storage core claimrule add -r 914 -t vendor -V XtremIO -M XtremApp -P VMW_VAAIP_T10 -c
VAAI -a -s -k 64
```
- ```
# esxcli storage core claimrule add -r 65430 -t vendor -V EMC -M SYMMETRIX -P VMW_VAAIP_SYMM -c
VAAI -a -s -m 200
```

Optimize Data Copying with Extended XCOPY Parameter

The extended XCOPY parameter optimizes the data copy between datastores across two different arrays. It allows you to migrate and clone workloads to offload XCOPY between two different arrays.

Extended XCOPY characteristics:

- Use a new extended XCOPY option in a claim rule.

```
esxcli storage core claimrule add -t vendor -V vendor -M model name -P VMW_VAAIP_T10 -c VAAI --autoassign -e
```

Example:

```
esxcli storage core claimrule add -t vendor -V DelleMC -M PowerStore -P VMW_VAAIP_T10 -c VAAI --autoassign -e
```

- Extended XCOPY is supported only on T10 plug-in.

1. Use the following command to check the extended XCOPY status.

```
esxcli storage core device vaa1 status get -d device
```

Example:

```
[root@tpp-h127:~] esxcli storage core device vaa1 status get -d naa.68ccf098007f554c61148eb1786e6749
naa.68ccf098007f554c61148eb1786e6749
  VAAI Plugin Name: VMW_VAAIP_T10
  ATS Status: supported
  Clone Status: supported
  Zero Status: supported
  Delete Status: supported
  Ex Clone Status: supported"
```

2. Enable extended XCOPY for the supported storage array.

- a) Verify if the VAAI_FILTER plug-in device is added.

```
esxcli storage core claimrule list -c Filter
```

- b) If the device is not added for VAAI_FILTER plug-in, use the following command.

```
esxcli storage core claimrule add -t vendor -V vendor -M model name -P VAAI_FILTER -c Filter --autoassign
```

- c) Load the VAAI_FILTER claim rule.

```
esxcli storage core claimrule load -c Filter
```

- d) Add the extended XCOPY claim rule with T10 plug-in.

```
esxcli storage core claimrule add -t vendor -V vendor -M model name -P VMW_VAAIP_T10 -c VAAI --autoassign -e
```

Note:

- The claim rule in step d uses ESXi default transfer size and single segment for XCOPY commands. To use array reported values of different transfer size and multiple segments, use optional parameters `-a`, `-s`, `-m` | `-k` with the above command.
- Extended XCOPY setup works across different arrays.
- Add a claim rule with optional parameters such as `-a`, `-s`, `-m` | `-k` for both source and destination devices. In this case, the minimum transfer size for source and destination arrays will be considered for data transfer.

Source Device:

```
esxcli storage core claimrule add -t vendor -V vendor -M model name -P VMW_VAAIP_T10 -c VAAI --autoassign -a -s
```

Destination Device:

```
esxcli storage core claimrule add -t vendor -V vendor -M model name -P
VMW_VAAIP_T10 -c VAAI --autoassign -a -s -e
```

e) Load the VAAI_T10 plug-in.

```
esxcli storage core claimrule load -c VAAI
```

f) Run the newly added claim rule options.

```
esxcli storage core claimrule run --claimrule-class=Filter
```

vSphere Hardware Acceleration on NAS Devices

With the hardware acceleration, ESXi hosts can integrate with NAS devices and use several hardware operations, such as Full File Clone and Fast File Clone that NAS storage provides. The hardware acceleration uses vSphere APIs for Array Integration (VAAI) to facilitate communications between the hosts and storage devices.

The VAAI NAS framework supports both versions of NFS storage, NFS 3 and NFS 4.1.

The VAAI NAS uses a set of storage primitives to offload storage operations from the host to the array. The following list shows the supported NAS operations:

Full File Clone

Supports an ability of NAS device to clone virtual disk files. This operation is similar to the VMFS block cloning, except that NAS devices clone entire files instead of file segments. Tasks that benefit from the full file clone operation include VM cloning, Storage vMotion, and deployment of VMs from templates.

When the ESXi host copies data with VAAI NAS, it does not need to read the data from the NAS and write back the data to the NAS. The host simply sends the copy command offloading it to the NAS. The copy process is done in the NAS, which reduces the load on the host.

Fast File Clone

This operation, also called array-based or native snapshots, offloads creation of virtual machine snapshots and linked clones to the array.

Reserve Space

Supports an ability of storage arrays to allocate space for a virtual disk file in the thick format.

Typically, when you create a virtual disk on an NFS datastore, the NAS server determines the allocation policy. The default allocation policy on most NAS servers is thin and does not guarantee backing storage to the file. However, the reserve space operation can instruct the NAS device to use vendor-specific mechanisms to reserve space for a virtual disk. As a result, you can create thick virtual disks on the NFS datastore if the backing NAS server supports the reserve space operation.

Extended Statistics

Supports visibility to space use on NAS devices. The operation enables you to query space utilization details for virtual disks on NFS datastores. The details include the size of a virtual disk and the space consumption of the virtual disk. This functionality is useful for thin provisioning.

With NAS storage devices, the hardware acceleration integration is implemented through vendor-specific NAS plug-ins. These plug-ins are typically created by vendors and are distributed as vendor packages. No claim rules are required for the NAS plug-ins to function.

Several tools for installing and updating NAS plug-ins are available. They include the `esxcli` commands and vSphere Lifecycle Manager. For more information, see the *VMware ESXi Upgrade and Managing Host and Cluster Lifecycle* documentation. For installation and update recommendations, see the [Knowledge Base article](#).

Note: NAS storage vendors might provide additional settings that can affect the performance and operation of VAAI. Follow the vendor's recommendations and configure the appropriate settings on both the NAS storage array and ESXi. See your storage vendor documentation for more information.

Enable NAS Native Snapshots on Virtual Machines

If your deployment includes NAS arrays that support the vSphere APIs for Array Integration (VAAI), you can use the fast file clone technology, also called native NFS snapshots, to make virtual machine snapshots. With this technology, the NFS device copies the virtual machine without having the ESXi host read and write the data. This operation might reduce the time and network load when VM snapshots are created.

- Verify that the NAS array supports the fast file clone operation with the VAAI NAS program.
- On your ESXi host, install vendor-specific NAS plug-in that supports the fast file cloning with VAAI.
- Follow the recommendations of your NAS storage vendor to configure any required settings on both the NAS array and ESXi. See your storage vendor documentation for more information.

By default, all newly created VMs support traditional ESXi snapshot technology. To use the NFS native snapshot technology, enable it for the VM.

1. In the vSphere Client, right-click the virtual machine and select **Edit Settings**.
2. Click the **VM Options** tab and expand the **Advanced** menu.
3. Click **Edit Configuration** next to Configuration Parameters.
4. Configure the `snapshot.alwaysAllowNative` parameter.

If the parameter exists, make sure that its value is set to True. If the parameter is not present, add it and set its value to True.

Name	Value
<code>snapshot.alwaysAllowNative</code>	True

Note: VMs stored on NFS datastores and configured with `snapshot.alwaysAllowNative` might encounter the `Invalid Disk Chain` error. For more information, see the [VMware Knowledge Base article 89923](#).

Storage Provisioning and Space Reclamation in vSphere

vSphere supports two models of storage provisioning, thick provisioning and thin provisioning.

Thick provisioning

It is a traditional model of the storage provisioning. With the thick provisioning, large amount of storage space is provided in advance in anticipation of future storage needs. However, the space might remain unused causing underutilization of storage capacity.

Thin provisioning


This method contrast with thick provisioning and helps you eliminate storage underutilization problems by allocating storage space in a flexible on-demand manner. With ESXi, you can use two models of thin provisioning, array-level and virtual disk-level. Thin provisioning allows you to report more virtual storage space than there is real physical capacity. This discrepancy can lead to storage over-subscription, also called over-provisioning. When you use thin provisioning, monitor actual storage usage to avoid conditions when you run out of physical storage space.

Virtual Disk Thin Provisioning with vSphere Storage

When you create a virtual machine, a certain amount of storage space on a datastore is provisioned to virtual disk files. ESXi supports thick and thin provisioning for virtual disks.

By default, ESXi offers a traditional storage provisioning method for virtual machines. With this method, you first estimate how much storage the virtual machine might need for its entire life cycle. You then provision a fixed amount of storage space to the VM virtual disk in advance, for example, 40 GB. The entire provisioned space is committed to the virtual disk. A virtual disk that immediately occupies the entire provisioned space is a thick disk.

ESXi supports thin provisioning for virtual disks. With the disk-level thin provisioning feature, you can create virtual disks in a thin format. For a thin virtual disk, ESXi provisions the entire space required for the disk's current and future activities, for example 40 GB. However, the thin disk uses only as much storage space as the disk needs for its initial operations. In this example, the thin-provisioned disk occupies only 20 GB of storage. If the disk requires more space, it can expand into its entire 40 GB of provisioned space.



About Virtual Disk Provisioning Policies

When you perform certain virtual machine management operations, you can specify a provisioning policy for the virtual disk file. The operations include creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine.

NFS datastores with Hardware Acceleration and VMFS datastores support the following disk provisioning policies. On NFS datastores that do not support Hardware Acceleration, only thin format is available.

You can use Storage vMotion or cross-host Storage vMotion to transform virtual disks from one format to another.

Thick Provision Lazy Zeroed

Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand later on first write from the virtual machine. Virtual machines do not read stale data from the physical device.

Thick Provision Eager Zeroed

A type of thick virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take longer to create virtual disks in this format than to create other types of disks. Increasing the size of an Eager Zeroed Thick virtual disk causes a significant startup time for the virtual machine.

Thin Provision

Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the virtual disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations. If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire datastore space provisioned to it. Thin provisioning is the fastest method to create a virtual disk because it creates a disk with just the header information. It does not allocate or zero out storage blocks. Storage blocks are allocated and zeroed out when they are first accessed.

Note: If a virtual disk supports clustering solutions such as Fault Tolerance, do not make the disk thin.

Create Thin Provisioned Virtual Disks

To save storage space, you can create a virtual disk in a thin provisioned format. The thin provisioned virtual disk starts small and expands as more disk space is required. You can create thin disks only on the datastores that support disk-level thin provisioning.

This procedure assumes that you are creating a new virtual machine. For information, see the *vSphere Virtual Machine Administration* documentation.

1. Create a virtual machine.
 - a) Right-click any inventory object that is a valid parent object of a virtual machine, such as a data center, folder, cluster, resource pool, or host, and select **New Virtual Machine**.
 - b) Select **Create a new virtual machine** and click **Next**.
 - c) Follow the steps required to create a virtual machine.
2. Configure the thin virtual disk.
 - a) On the Customize Hardware page, click the **Virtual Hardware** tab.
 - b) Click the **New Hard disk** triangle to expand the hard disk options.
 - c) Optional: Adjust the default disk size.

With a thin virtual disk, the disk size value shows how much space is provisioned and guaranteed to the disk. At the beginning, the virtual disk might not use the entire provisioned space. The actual storage use value can be less than the size of the virtual disk.
 - d) Select **Thin Provision** for Disk Provisioning.
3. Finish the virtual machine creation.

You created a virtual machine with a disk in the thin format.

If you created a virtual disk in the thin format, you can later inflate it to its full size.

View Virtual Machine Storage Resources

You can view how datastore storage space is allocated for your virtual machines.

1. Browse to the virtual machine.
2. Double-click the virtual machine and click the **Summary** tab.
3. Review the storage use information in the upper right area of the **Summary** tab.

Storage Usage shows how much datastore space is occupied by virtual machine files, including configuration and log files, snapshots, virtual disks, and so on. When the virtual machine is running, the used storage space also includes swap files.

For virtual machines with thin disks, the actual storage use value might be less than the size of the virtual disk.

Determine the Disk Format of a Virtual Machine

You can determine whether your virtual disk is in thick or thin format.

1. Right-click the virtual machine and select **Edit Settings**.
2. Click the **Virtual Hardware** tab.
3. Click the **Hard Disk** triangle to expand the hard disk options.

The **Type** text box shows the format of your virtual disk.

If your virtual disk is in the thin format, you can inflate it to its full size.

Inflate Thin Virtual Disks

If you created a virtual disk in the thin format, you can change the format to thick.

- Make sure that the datastore where the virtual machine resides has enough space.
- Make sure that the virtual disk is thin.
- Remove snapshots.
- Power off your virtual machine.

You use the datastore browser to inflate the thin virtual disk.

1. In the vSphere Client, navigate to the folder of the virtual disk you want to inflate.
 - a) Navigate to the virtual machine.
 - b) Click the **Datastores** tab.

The datastore that stores the virtual machine files is listed.

- c) Right-click the datastore and select **Browse Files**.

The datastore browser displays contents of the datastore.

2. Expand the virtual machine folder and browse to the virtual disk file that you want to convert.

The file has the `.vmdk` extension and is marked with the virtual disk (🗄️) icon.

3. Select the virtual disk file and click **Inflate**.

Note: The option might not be available if the virtual disk is thick or when the virtual machine is running.

The inflated virtual disk occupies the entire datastore space originally provisioned to it.

Handling Datastore Over-Subscription

Because the provisioned space for thin disks can be greater than the committed space, a datastore over-subscription can occur, which results in the total provisioned space for the virtual machine disks on the datastore being greater than the actual capacity.

Over-subscription can be possible because usually not all virtual machines with thin disks need the entire provisioned datastore space simultaneously. However, if you want to avoid over-subscribing the datastore, you can set up an alarm that notifies you when the provisioned space reaches a certain threshold.

For information on setting alarms, see the *vCenter Server and Host Management* documentation.

If your virtual machines require more space, the datastore space is allocated on a first come first served basis. When the datastore runs out of space, you can add more physical storage and increase the datastore.

See [Increase Capacity of vSphere VMFS Datastore](#).

ESXi and Array Thin Provisioning

You can use thin-provisioned storage arrays with ESXi.

The ESXi host integrates with block-based storage and performs these tasks:

- The host can recognize underlying thin-provisioned LUNs and monitor their space use to avoid running out of physical space. The LUN space might change if, for example, your VMFS datastore expands or if you use Storage vMotion to migrate virtual machines to the thin-provisioned LUN. The host warns you about breaches in physical LUN space and about out-of-space conditions.
- The host can run the automatic T10 `unmap` command from VMFS6 and VM guest operating systems to reclaim unused space from the array. VMFS5 supports a manual space reclamation method.

Note: ESXi does not support activating or deactivating of thin provisioning on a storage device.

Requirements

To use the thin provisioning reporting and space reclamation features, follow these requirements:

- Use an appropriate ESXi version.

Supported thin provisioning components	ESXi 6.5 and later
Thin provisioning	Yes
Unmap command originating from VMFS	Automatic for VMFS6
Unmap command originating from guest OS	Yes (VMFS6)

- Use storage systems that support T10-based vSphere Storage APIs - Array Integration (VAAI), including thin provisioning and space reclamation. For information, contact your storage provider and check the *VMware Compatibility Guide* documentation.

Monitoring Space Use

The thin provision integration functionality helps you to monitor the use of space on thin-provisioned LUNs and to avoid running out of space.

The following sample flow demonstrates how the ESXi host and the storage array interact to generate breach of space and out-of-space warnings for a thin-provisioned LUN. The same mechanism applies when you use Storage vMotion to migrate virtual machines to the thin-provisioned LUN.

1. Using storage-specific tools, your storage administrator provisions a thin LUN and sets a soft threshold limit that, when reached, triggers an alert. This step is vendor-specific.
2. Using the vSphere Client, you create a VMFS datastore on the thin-provisioned LUN. The datastore spans the entire logical size that the LUN reports.
3. As the space used by the datastore increases and reaches the set soft threshold, the following actions take place:
 - a. The storage array reports the breach to your host.
 - b. Your host triggers a warning alarm for the datastore.
You can contact the storage administrator to request more physical space. Alternatively, you can use Storage vMotion to evacuate your virtual machines before the LUN runs out of capacity.
4. If no space is left to allocate to the thin-provisioned LUN, the following actions take place:
 - a. The storage array reports out-of-space condition to your host.

CAUTION:

In certain cases, when a LUN becomes full, it might go offline or get unmapped from the host.

- b. The host pauses virtual machines and generates an out-of-space alarm.
You can resolve the permanent out-of-space condition by requesting more physical space from the storage administrator.

Identify Thin-Provisioned Storage Devices

Use the `esxcli` command to verify whether a particular storage device is thin-provisioned.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

Run the `esxcli storage core device list -d=device_ID` command.

The following thin provisioning status indicates that the storage device is thin-provisioned.

```
# esxcli storage core device list -d naa.XXXXXXXXXXXXX4c
naa.XXXXXXXXXXXXX4c
  Display Name: XXXX Fibre Channel Disk(naa.XXXXXXXXXXXXX4c)
  Size: 20480
  Device Type: Direct-Access
  Multipath Plugin: NMP
  -----
  Thin Provisioning Status: yes
  -----
```

An unknown status indicates that a storage device is thick.


Note: Some storage systems present all devices as thin-provisioned no matter whether the devices are thin or thick. Their thin provisioning status is always `yes`. For details, check with your storage vendor.

Storage Space Reclamation in vSphere

ESXi supports the space reclamation command, also called SCSI unmap command or NVMe dataset management deallocate command, that originates from a VMFS datastore or a VM guest operating system. The command helps thin-provisioned storage arrays to reclaim unused space from the VMFS datastore and thin virtual disks on the datastore. The VMFS6 datastore can send the space reclamation command automatically. With the VMFS5 datastore, you can manually reclaim storage space.

For information about space reclamation with the Virtual Volumes datastores, see [Reclaim Space on the vSphere Virtual Volumes Datastores](#).

You free storage space inside the VMFS datastore when you delete or migrate the VM, consolidate a snapshot, and so on. Inside the virtual machine, storage space is freed when you delete files on the thin virtual disk. These operations leave blocks of unused space on the storage array. However, when the array is not aware that the data was deleted from the blocks, the blocks remain allocated by the array until the datastore releases them. VMFS uses the SCSI unmap command to indicate to the array that the storage blocks contain deleted data, so that the array can unallocate these blocks.



The command can also originate directly from the guest operating system. Both VMFS5 and VMFS6 datastores can provide support to the unmap command that proceeds from the guest operating system. However, the level of support is limited on VMFS5.

Depending on the type of your VMFS datastore, you use different methods to configure space reclamation for the datastore and your virtual machines.

Watch the following video to learn more about how space reclamation works.

Space Reclamation on vSphere VMFS Datastores

Deleting or removing files from a VMFS datastore frees space within the file system. This free space is mapped to a storage device until the file system releases or unmaps it. ESXi supports reclamation of free space, which is also called the unmap operation.

The operation helps the storage array to reclaim unused free space. Unmapped space can be then used for other storage allocation requests and needs.

Automatic Reclamation of Free Space on VMFS6 Datastore

On VMFS6 datastores, ESXi supports the automatic asynchronous reclamation of free space. VMFS6 can run the unmap command to release free storage space in the background on thin-provisioned storage arrays that support unmap operations.

Automatic unmap processing has several advantages:

- Unmap requests are sent at a constant rate, which helps to avoid any instant load on the backing array.
- Freed regions are batched and unmapped together.
- I/O performance of other workloads is not impacted by the unmap command.

For VMFS6 datastores, you can configure the following space reclamation parameters.

Space reclamation granularity

Granularity defines the minimum size of a released space sector that underlying storage can reclaim. Storage cannot reclaim those sectors that are smaller in size than the specified granularity.

For VMFS6, reclamation granularity equals the block size. When you specify the block size as 1 MB, the granularity is also 1 MB. Storage sectors of the size smaller than 1 MB are not reclaimed.

Note: Certain storage arrays recommend an optimal unmap granularity. ESXi supports automatic unmap processing on arrays with the recommended unmap granularity of 1 MB or greater, for example, 16 MB. On the arrays with the optimal granularity of 1 MB and less, the unmap operation is supported if the granularity is a factor of 1 MB. For example, 1 MB is divisible by 512 bytes, 4 KB, 64 KB, and so on.

Space reclamation method

The method can be either priority or fixed. When the method you use is priority, you configure the priority rate. For the fixed method, you must indicate the reclamation rate in MB per second.

You can use the vSphere Client to set the priority method when you create the VMFS6 datastore. To enable the fixed method, you can later modify the space reclamation settings for the existing datastore.

Alternately, you can use the `esxcli storage vmfs reclaim config set` command to set space reclamation parameters.

Space reclamation priority rate

This parameter defines the rate at which the space reclamation operation is performed when you use the priority reclamation method. Typically, VMFS6 can send the unmap commands either in bursts or sporadically depending on the workload and configuration. For VMFS6, you can specify one of the following options.

Space Reclamation Priority	Description	Configuration
None	Deactivates the unmap operations for the datastore.	vSphere Client <code>esxcli</code> command
Low (default)	Sends the unmap command at a less frequent rate, 25–50 MB per second.	vSphere Client <code>esxcli</code> command
Medium	Sends the command at a rate twice faster than the low rate, 50–100 MB per second.	<code>esxcli</code> command

Space Reclamation Priority	Description	Configuration
High	Sends the command at a rate three times faster than the low rate, over 100 MB per second.	<code>esxcli</code> command

Space reclamation hosts

This parameter defines the number of ESXi hosts that can issue unmap commands to a VMFS6 datastore for automatic space reclamation. Use this parameter to limit the cumulative amount of space reclamation load issued to the target backing VMFS6 datastore.

After you enable space reclamation, the VMFS6 datastore can start releasing the blocks of unused space only when it has at least one open file. This condition can be fulfilled when, for example, you power on one of the VMs on the datastore.

Configure Priority Reclamation when Creating a VMFS6 Datastore

Space reclamation settings control how deleted or unmapped blocks are reclaimed on the devices that back the datastore. When you create a VMFS6 datastore in the vSphere Client, the only method for space reclamation you can specify is the priority method. You can also deactivate space reclamation when creating the datastore.

1. In the vSphere Client object navigator, browse to a host, a cluster, or a data center.
2. From the right-click menu, select **Storage > New Datastore**.
3. Follow the steps required to create a VMFS6 datastore.
4. On the **Partition configuration** page, specify the space reclamation parameters.

The parameters define granularity and the priority rate at which space reclamation operations are performed. You can also use this page to deactivate space reclamation for the datastore.

Option	Description
Block size	The block size on a VMFS datastore defines the maximum file size and the amount of space the file occupies. VMFS6 supports the block size of 1 MB.
Space reclamation granularity	Specify granularity for the unmap operation. Unmap granularity equals the block size, which is 1 MB. Storage sectors of the size smaller than 1 MB are not reclaimed.
Space reclamation priority	Select one of the following options. • Low (default). Use the priority method for space reclamation. The unmap operation is performed at a low priority rate. • None. Select this option if you do not want to use the space reclamation operations for the datastore.

Note: In the vSphere Client, the only available settings for the space reclamation priority are Low and None. To change the settings to Medium or High, use the `esxcli` command.

5. Finish the datastore creation process.

After you activate space reclamation, the VMFS6 datastore can start releasing the blocks of unused space only when it has at least one open file. This condition can be fulfilled when, for example, you power on one of the VMs on the datastore.

Configure Fixed Reclamation for the Existing VMFS6 Datastore

To activate the fixed reclamation method, modify the space reclamation settings for the datastore.

1. In the vSphere Client, navigate to the datastore.
2. Select **Edit Space Reclamation** from the right-click menu.
3. Specify the space reclamation setting.

Option	Description
Enable automatic space reclamation at fixed rate	Use the fixed method for space reclamation. Specify reclamation rate in MB per second. Minimum available value is 10 MB per second.
Disable automatic space reclamation	Deleted or unmapped blocks are not reclaimed. Select this option if you want to deactivate the space reclamation operations for the datastore.

4. Click **OK** to save the new settings.
5. Unmount and remount the datastore for the changes to take place.
 - a) [Unmount Datastores](#).
 - b) [Mount Datastores](#).
6. Repeat this procedure for all ESXi hosts that access the datastore.

The modified value for the space reclamation priority appears on the **General** page for the datastore.

Use the ESXCLI Command to Change Space Reclamation on VMFS6

You can change the default space reclamation priority, granularity, and other parameters.

1. On the ESXi host, use the following command to set space reclamation parameters.

```
esxcli storage vmfs reclaim config set
```

The command takes these options:

Option	Description
-b --reclaim-bandwidth	Space reclamation fixed bandwidth in MB per second.
-g --reclaim-granularity	Minimum granularity of automatic space reclamation in bytes.
-m --reclaim-method	Method of automatic space reclamation. Supported options: • priority • fixed
-n --reclaim-maxhosts	Maximum number of ESXi hosts that can issue unmap commands to a VMFS6 datastore for automatic space reclamation.
-p --reclaim-priority	Priority of automatic space reclamation. Supported options: • none • low • medium

Option	Description
	high
-l --volume-label	The label of the target VMFS volume.
-u --volume-uuid	The uuid of the target VMFS volume.

You can use the following examples.

- Set the reclamation method to fixed and the rate to 100 MB per second.

```
esxcli storage vmfs reclaim config set --volume-label datastore_name --reclaim-method fixed -b 100
```

- Deactivate the VMFS automatic space reclamation.

```
esxcli storage vmfs reclaim config set --volume-label datastore_name --reclaim-priority none
```

Note: If you set only the `-n|--reclaim-maxhosts` parameter without any other parameters, you can skip Step 2. It takes around 30 minutes for the updated `reclaim-maxhosts` parameter value to take effect on all ESXi hosts where the VMFS6 datastore mounted.

- Unmount the VMFS6 datastore from all other ESXi hosts where the datastore is mounted, and then remount it again.
 - [Unmount Datastores.](#)
 - [Mount Datastores.](#)

This step ensures that all ESXi hosts where the VMFS6 datastore was mounted update the reclamation settings for the datastore.

Verify Settings for Automatic Space Reclamation on VMFS6

After you configure or edit space reclamation parameters for a VMFS6 datastore, you can review your settings.

- In the vSphere Client, navigate to the datastore.
- Click the **Configure** tab.
- Click **General**.
- Under **Space Reclamation**, review the setting for the space reclamation priority.

If you configured any values through the `esxcli` command, for example, Medium or High for the space reclamation priority, these values also appear in the vSphere Client.

⌵ vSphere Client
🔍 Search in all environments
⚠️ There are

📁
📁
📁
🌐

⌵
10.2.10.104 10.2.10.105 10.2.10.106

⌵
VSAN-DC

Datastore

📁 vsanDatastore

📁 Datastore
⋮ ACTIONS

Summary
Monitor
Configure
Permissions
File

Alarm Definitions

Scheduled Tasks

General

Device Backing

Connectivity and Multipathing

Hardware Acceleration

Capability sets

Properties

Name

> File system

Drive type

Capacity

Total Capacity

Provisioned Space

Free Space

Datastore Capabilities

Thin Provisioning

> Storage I/O Control

Space Reclamation

Space reclamation

Reclamation rate

Obtaining Parameters for VMFS6 Space Reclamation

You can also use the `esxcli storage vmfs reclaim config get -l=VMFS_label|-u=VMFS_uuid` command to obtain information for the space reclamation configuration.

```
# esxcli storage vmfs reclaim config get -l my_datastore
Reclaim Granularity: 1048576 Bytes
Reclaim Priority: low
Reclaim Method: priority
Reclaim Bandwidth: 26 MB/s
```

Reclaim Maximum Hosts: 128

Manually Reclaim Accumulated Storage Space on VMFS5

On VMFS datastores that do not support automatic space reclamation, such as VMFS5, you can use the `esxcli` command to reclaim unused storage space manually.

Install ESXCLI. See [Getting Started with ESXCLI](#). For troubleshooting, run `esxcli` commands in the ESXi Shell.

When you use the command, be aware that it might send many unmap requests at a time. This action can lock some of the resources during the operation. When you use the command, be aware that it might send many unmap requests at a time. This action can lock some of the resources during the operation.

1. To reclaim unused storage blocks on the thin-provisioned device, run the following command:

```
esxcli storage vmfs unmap
```

The command takes these options:

Option	Description
<code>-l --volume-label=volume_label</code>	The label of the VMFS volume to unmap. A mandatory argument. If you specify this argument, do not use <code>-u --volume-uuid=volume_uuid</code> .
<code>-u --volume-uuid=volume_uuid</code>	The UUID of the VMFS volume to unmap. A mandatory argument. If you specify this argument, do not use <code>-l --volume-label=volume_label</code> .
<code>-n --reclaim-unit=number</code>	Number of VMFS blocks to unmap per iteration. An optional argument. If it is not specified, the command uses the default value of 200.

2. To verify whether the unmap process has finished, search for unmap in the `vmkernel.log` file.

Reclaim Space on the vSphere Virtual Volumes Datastores

When your ESXi host uses the SCSI or NVMe protocol to connect to storage in the Virtual Volumes environment, the data and configuration virtual volumes are formatted with VMFS6. As a result, the space reclamation process available for VMFS can also be applied to both the data and configuration Virtual Volumes. When you reclaim space on the Virtual Volumes datastore, you can use the `esxcli storage vvol` command to manually unmap free blocks from the virtual volumes.

- For information about virtual volumes, see [Types of Virtual Volumes](#).
- For information about the space reclamation process, see [Storage Space Reclamation in vSphere](#).

Note: Starting with 8.0 Update 3:

- NVMe backed data Virtual Volumes support space reclamation requests from guest operating systems.
- Configuration Virtual Volumes support automatic space reclamation. To reclaim the space manually, follow the steps in the given procedure.

1. Obtain the storage container ID and meta vVol UUID for your configuration virtual volume formatted with VMFS.

- a) Navigate to the VM that contains virtual volumes.

For example:

```
[root@:~] cd /vmfs/volumes/vvol_datastore/VM_1
```

- b) From the path, extract the storage container ID and meta vVol UUID.

```
[root@:/vmfs/volumes/vvol:4abe7eb4bf234ee0xxxxxxxx/rfc4122.1ffee4bd-9c55-49b9-876d-xxxxxxxx
```

In this example,

Container ID is `vvol:4abe7eb4bf234ee0xxxxxxxx`

Meta vVol UUID is `rfc4122.1ffee4bd-9c55-49b9-876d-xxxxxxxx`

2. Use the following command to unmap free blocks from the configuration volume.

```
esxcli storage vvol virtualvolume unmap
```

The command takes the following options:

Option	Description
-c --container-id=<str>	Storage container ID. The option is required.
-n --reclaim-unit=<long>	Number of VMFS blocks to unmap per iteration.
-u --uuid=<str>	Meta vVol UUID. The option is required.

3. Verify that you successfully unmapped the blocks by checking the `hostd` log file.

```
grep "Total Unmapped blocks from vmfs" /var/run/log/hostd.log
```

```
1. [root@:~]esxcli storage vvol virtualvolume unmap -c vvol:4abe7eb4bf234ee0xxxxxxx -u
rfc4122.1ffee4bd-9c55-49b9-876d-xxxxxxx
```

```
2. [root@:~] grep "Total Unmapped blocks from vmfs" /var/run/log/hostd.log
```

The output looks similar to the following:

```
2023-06-01T10:27:45.836Z In(166) Hostd[1000342556]: [Originator@6876 sub=Libs
opID=esxcli-5f-1861 sid=523fac9d user=root] Unmap: Done : Total Unmapped blocks from
vmfs6 volume rfc4122.1ffee4bd-9c55-49b9-87 6d-xxxxxxx : 0 (LFB Pass)
2023-06-01T10:28:13.279Z In(166) Hostd[1000342556]: [Originator@6876 sub=Libs
opID=esxcli-5f-1861 sid=523fac9d user=root] Unmap: Done : Total Unmapped blocks from
vmfs6 volume rfc4122.1ffee4bd-9c55-49b9-87 6d-xxxxxxx : 259421 (SFB Pass)
```

Space Reclamation Requests from Guest Operating Systems

ESXi supports the unmap commands issued directly from a guest operating system to reclaim storage space. The level of support and requirements depend on the type of datastore where your virtual machine resides.

Inside a virtual machine, storage space is freed when, for example, you delete files on the thin virtual disk. The guest operating system notifies VMFS about freed space by sending the unmap command. The unmap command sent from the guest operating system releases space within the VMFS datastore. The command then proceeds to the array, so that the array can reclaim the freed blocks of space.

Space Reclamation for VMFS6 Virtual Machines

VMFS6 generally supports automatic space reclamation requests that generate from the guest operating systems, and passes these requests to the array. Many guest operating systems can send the unmap command and do not require any additional configuration. The guest operating systems that do not support the automatic unmaps might require user intervention. For information about guest operating systems that support the automatic space reclamation on VMFS6, contact your vendor.

Generally, the guest operating systems send the unmap commands based on the unmap granularity they advertise. For details, see documentation provided with your guest operating system.

The following considerations apply when you use space reclamation with VMFS6:

- VMFS6 processes the unmap request from the guest OS only when the space to reclaim equals 1 MB or is a multiple of 1 MB. If the space is less than 1 MB or is not aligned to 1 MB, the unmap requests are not processed.
- For VMs with snapshots in the default SEsparse format, VMFS6 supports the automatic space reclamation only on ESXi hosts version 6.7 or later.

Space reclamation affects only the top snapshot and works when the VM is powered on.

Space Reclamation for VMFS5 Virtual Machines

Typically, the unmap command that generates from the guest operation system on VMFS5 cannot be passed directly to the array. You must run the `esxcli storage vmfs unmap` command to trigger unmaps for the array.

However, for a limited number of the guest operating systems, VMFS5 supports the automatic space reclamation requests.

To send the unmap requests from the guest operating system to the array, the virtual machine must meet the following prerequisites:

- The virtual disk must be thin-provisioned.
- Virtual machine hardware must be of version 11 (ESXi 6.0) or later.
- The advanced setting `EnableBlockDelete` must be set to 1.
- The guest operating system must be able to identify the virtual disk as thin.

Using vmkfstools in vSphere

`vmkfstools` is one of the ESXi Shell commands for managing VMFS volumes, storage devices, and virtual disks. You can perform many storage operations using the `vmkfstools` command. For example, you can create and manage VMFS datastores on a physical partition, or manipulate virtual disk files, stored on VMFS or NFS datastores.

Note: After you make a change using the `vmkfstools`, the vSphere Client might not be updated immediately. Use a refresh or rescan operation from the client.

For more information on the ESXi Shell, see *Getting Started with ESXCLI*.

vmkfstools Command Syntax

Generally, you do not need to log in as the root user to run the `vmkfstools` commands. However, some commands, such as the file system commands, might require the root user login.

The `vmkfstools` command supports the following command syntax:

```
vmkfstoolsoptions target.
```

Target specifies a partition, device, or path to apply the command option to.

Table 228: vmkfstools Command Arguments

Argument	Description
options	One or more command-line options and associated arguments that you use to specify the activity for <code>vmkfstools</code> to perform. For example, selecting the disk format when creating a new virtual disk. After entering the option, specify a target on which to perform the operation. Target can indicate a partition, device, or path.
partition	Specifies disk partitions. This argument uses a <code>disk_ID:P</code> format, where <code>disk_ID</code> is the device ID returned by the storage array and <code>P</code> is an integer that represents the partition number. The partition digit must be greater than zero (0) and must correspond to a valid VMFS partition.
device	Specifies devices or logical volumes. This argument uses a path name in the ESXi device file system. The path name begins with <code>/vmfs/devices</code> , which is the mount point of the device file system. Use the following formats when you specify different types of devices: • <code>/vmfs/devices/disks</code> for local or SAN-based disks.

Argument	Description
	• <code>/vmfs/devices/lvm</code> for ESXi logical volumes. • <code>/vmfs/devices/generic</code> for generic SCSI devices.
path	Specifies a VMFS file system or file. This argument is an absolute or relative path that names a directory symbolic link, a raw device mapping, or a file under <code>/vmfs</code>. • To specify a VMFS file system, use this format: <code>/vmfs/volumes/file_system_UUID</code> or <code>/vmfs/volumes/file_system_label</code> • To specify a file on a VMFS datastore, use this format: <code>/vmfs/volumes/file_system_label file_system_UUID/[dir]/myDisk.vmdk</code> If the current working directory is the parent directory of <code>myDisk.vmdk</code>, do not enter the entire path.

The vmkfstools Command Options

The `vmkfstools` command has several options. Some of the options are suggested for advanced users only.

The long and single-letter forms of the options are equivalent. For example, the following commands are identical.

```
vmkfstools --createfs vmfs6 --blocksize 1m disk_ID:P
vmkfstools -C vmfs6 -b 1m disk_ID:P
```

-v Suboption

The `-v` suboption indicates the verbosity level of the command output.

The format for this suboption is as follows:

```
-v --verbose number
```

You specify the *number* value as an integer from 1 through 10.

You can specify the `-v` suboption with any `vmkfstools` option. If the output of the option is not suitable for use with the `-v` suboption, `vmkfstools` ignores `-v`.

Note: Because you can include the `-v` suboption in any `vmkfstools` command line, `-v` is not included as a suboption in the option descriptions.

File System Options of vSphere vmkfstools Command

File system options allow you to create and manage VMFS datastores. These options do not apply to NFS. You can perform many of these tasks through the vSphere Client.

Listing Attributes of a VMFS Datastore

Use the `vmkfstools` command to list attributes of a VMFS datastore.

```
-P|--queryfs
-h|--humanreadable
```

When you use this option on any file or directory that resides on a VMFS datastore, the option lists the attributes of the specified datastore. The listed attributes typically include the file system label, the number of extents for the datastore, the UUID, and a list of the devices where each extent resides.

Note: If any device backing VMFS file system goes offline, the number of extents and available space change accordingly.

You can specify the `-h|--humanreadable` suboption with the `-P` option. If you do so, `vmkfstools` lists the capacity of the volume in a more readable form.

Example of Listing VMFS Attributes

```
~ vmkfstools -P -h /vmfs/volumes/my_vmfs
VMFS-5.81 (Raw Major Version: 14) file system spanning 1 partitions.
File system label (if any): my_vmfs
Mode: public
Capacity 99.8 GB, 97.5 GB available, file block size 1 MB, max supported file size 62.9 TB
UUID: 571fe2fb-ec4b8d6c-d375-XXXXXXXXXXXX
Partitions spanned (on "lvm"):
    eui.3863316131XXXXXX:1
Is Native Snapshot Capable: YES
```

Creating a VMFS Datastore or a Scratch Partition

Use the `vmkfstools` command to create a VMFS datastore or a scratch partition.

```
-C|--createfs [vmfs5|vmfs6|vfat]
```

This option creates the VMFS datastore on the specified SCSI or NVMe partition, such as `disk_ID:P`. The partition becomes the head partition of the datastore. For VMFS5 and VMFS6, the only available block size is 1 MB.

You can specify the following suboptions with the `-C` option.

- `-S|--setfsname` - Define the volume label of the VMFS datastore you are creating. Use this suboption only with the `-C` option. The label you specify can be up to 128 characters long and cannot contain any leading or trailing blank spaces.

Note: vCenter Server supports the 80 character limit for all its entities. If a datastore name exceeds this limit, the name gets shortened when you add this datastore to vCenter Server.

After you define a volume label, you can use it whenever you specify the VMFS datastore for the `vmkfstools` command. The volume label appears in listings generated for the `ls -l` command and as a symbolic link to the VMFS volume under the `/vmfs/volumes` directory.

To change the VMFS volume label, use the `ln -sf` command. Use the following as an example:

```
ln -sf /vmfs/volumes/UUID /vmfs/volumes/datastore
```

`datastore` is the new volume label to use for the `UUID` VMFS.

Note: If your host is registered with vCenter Server, any changes you make to the VMFS volume label get overwritten by vCenter Server. This operation guarantees that the VMFS label is consistent across all vCenter Server hosts.

- `-Y|--unmapGranularity #[bBsSkKmMgGtT]` - This suboption applies to VMFS6 only. Define granularity for the unmap operation. The default granularity is 1 MB. As with the block size, enter the unit type.
- `-O|--unmapPriority <none|low|medium|high>` - This suboption applies to VMFS6 only. Define the priority for the unmap operation.

Example for Creating a VMFS File System

This example illustrates creating a VMFS6 datastore named `my_vmfs` on the `naa.ID:1` partition.

```
~ vmkfstools -C vmfs6 -S my_vmfs /vmfs/devices/disks/naa.ID:1
```

Adding an Extent to a VMFS Datastore

Use the `vmkfstools` command to add an extent to a VMFS datastore.

When you add an extent, you span the VMFS datastore from the head partition across the partition specified by `span_partition`.

```
-Z|--spanfs span_partitionhead_partition
```

You must specify the full path name for the head and span partitions, for example `/vmfs/devices/disks/disk_ID:1`. Each time you use this option, you add an extent to the VMFS datastore, so that the datastore spans multiple partitions.

CAUTION: When you run this option, you lose all data that previously existed on the SCSI or NVMe device you specified in `span_partition`.

Example for Extending a VMFS Datastore

In this example, you extend the existing head partition of the VMFS datastore over a new partition.

```
~ vmkfstools -Z /vmfs/devices/disks/naa.disk_ID_2:1 /vmfs/devices/disks/naa.disk_ID_1:1
```

The extended datastore spans two partitions, `naa.disk_ID_1:1` and `naa.disk_ID_2:1`. In this example, `naa.disk_ID_1:1` is the name of the head partition.

Expanding a VMFS Datastore

Instead of adding an extent to a VMFS datastore, you can increase the size of an existing datastore. Use the `vmkfstools -G` command.

You might increase the datastore size after the underlying storage had its capacity increased.

The command uses the following option:

```
-G|--growfs device:device
```

This option expands the VMFS datastore or its specific extent. For example,

```
vmkfstools --growfs /vmfs/devices/disks/disk_ID:1 /vmfs/devices/disks/disk_ID:1
```

Virtual Disk Options

Virtual disk options allow you to set up, migrate, and manage virtual disks stored on your datastores. You can also perform most of these tasks through the vSphere Client.

Supported Disk Formats

When you create or clone a virtual disk, you can use the `-d|--diskformat` suboption to specify the format for the disk.

Choose from the following formats:

- `zeroedthick` (default) – Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation, but is zeroed out on demand on first write from the virtual machine. The virtual machine does not read stale data from disk.
- `eagerzeroedthick` – Space required for the virtual disk is allocated at creation time. In contrast to `zeroedthick` format, the data remaining on the physical device is zeroed out during creation. It might take much longer to create disks in this format than to create other types of disks.
- `thin` – Thin-provisioned virtual disk. Unlike with the `thick` format, space required for the virtual disk is not allocated during creation, but is supplied, zeroed out, on demand.
- `rdm: device` – Virtual compatibility mode raw disk mapping.
- `rdmp: device` – Physical compatibility mode (pass-through) raw disk mapping.
- `2gbsparse` – A sparse disk with the maximum extent size of 2 GB. You can use disks in this format with hosted VMware products, such as VMware Fusion. However, you cannot power on the sparse disk on an ESXi host unless you first re-import the disk with `vmkfstools` in a compatible format, such as `thick` or `thin`.

Disk Formats on NFS Datastores

The only disk formats you can use for NFS are `thin`, `thick`, `zeroedthick`, and `2gbsparse`.

`Thick`, `zeroedthick`, and `thin` formats usually behave the same because the NFS server and not the ESXi host determines the allocation policy. The default allocation policy on most NFS servers is `thin`. However, on NFS servers that support Storage APIs - Array Integration, you can create virtual disks in `zeroedthick` format. The reserve space operation enables NFS servers to allocate and guarantee space.

For more information on array integration APIs, see [Storage Hardware Acceleration in vSphere](#).

Creating a Virtual Disk

Use the `vmkfstools` command to create a virtual disk.

```
-c|--createvirtualdisk size[bB|sS|kK|mM|gG]
-d|--diskformat [thin|zeroedthick|eagerzeroedthick]
-W|--objecttype [file|vsan|vvol]
--policyFile fileName
```

This option creates a virtual disk at the specified path on a datastore. Specify the size of the virtual disk. When you enter the value for *size*, you can indicate the unit type by adding a suffix of `k` (kilobytes), `m` (megabytes), or `g` (gigabytes). The unit type is not case-sensitive. `vmkfstools` interprets either `k` or `K` to mean kilobytes. If you do not specify a unit type, `vmkfstools` defaults to bytes.

You can specify the following suboptions with the `-c` option.

- `-d|--diskformat` specifies disk formats.
- `-W|--objecttype` specifies whether the virtual disk is a file on a VMFS or NFS datastore, or an object on a vSAN or Virtual Volumes datastore.
- `--policyFile fileName` specifies VM storage policy for the disk.

Example for Creating a Virtual Disk

This example shows how to create a two-gigabyte virtual disk file named `disk.vmdk`. You create the disk on the VMFS datastore named `myVMFS`. The disk file represents an empty virtual disk that virtual machines can access.

```
vmkfstools -c 2048m /vmfs/volumes/myVMFS/disk.vmdk
```

Initializing a Virtual Disk

Use the `vmkfstools` command to initialize a virtual disk.

```
-w|--writezeros
```

This option cleans the virtual disk by writing zeros over all its data. Depending on the size of your virtual disk and the I/O bandwidth to the device hosting the virtual disk, completing this command might take a long time.

Note: This command is not supported on NFS datastores.

CAUTION: When you use this command, you lose any existing data on the virtual disk.

Inflating a Thin Virtual Disk

Use the `vmkfstools` command to inflate a thin virtual disk.

```
-j|--inflatedisk
```

This option converts a `thin` virtual disk to `eagerzeroedthick`, preserving all existing data. The option allocates and zeroes out any blocks that are not already allocated.

Converting a Zeroedthick Virtual Disk to an Eagerzeroedthick Disk

Use the `vmkfstools` command to convert any zeroedthick virtual disk to an eagerzeroedthick disk.

```
-k|--eagerzero
```

While performing the conversion, this option preserves any data on the virtual disk.

Follow this example:

```
vmkfstools --eagerzero /vmfs/volumes/myVMFS/VMName/disk.vmdk
```

Removing Zeroed Blocks

Use the `vmkfstools` command to remove zeroed blocks.

```
-K|--punchzero
```

This option deallocates all zeroed out blocks and leaves only those blocks that were allocated previously and contain valid data. The resulting virtual disk is in thin format.

Deleting a Virtual Disk

Use the `vmkfstools` command to delete a virtual disk file at the specified path on the VMFS volume.

Use the following option:

```
-U|--deletevirtualdisk
```

Renaming a Virtual Disk

Use the `vmkfstools` command to rename a virtual disk file at the specified path on the VMFS volume.

You must specify the original filename or file path *oldName* and the new filename or file path *newName*.

```
-E|--renamevirtualdisk oldName newName
```

Cloning or Converting a Virtual Disk or RDM

Use the `vmkfstools` command to create a copy of a virtual disk or raw disk you specify.

A non-root user cannot clone a virtual disk or an RDM. You must specify the original filename or file path *oldName* and the new filename or file path *newName*.

```
-i|--clonevirtualdisk oldName newName
-d|--diskformat [thin|zeroedthick|eagerzeroedthick|rdm:device|rdmp:device|2gbsparse]
-W|--objecttype [file|vsan|vvol]
--policyFile fileName
-N|--avoidnativeclone
```

Use the following suboptions to change corresponding parameters for the copy you create.

- `-d|--diskformat` specifies disk formats.

- `-W|--objecttype` specifies whether the virtual disk is a file on a VMFS or NFS datastore, or an object on a vSAN or Virtual Volumes datastore.
- `--policyFile fileName` specifies VM storage policy for the disk.

By default, ESXi uses its native methods to perform the cloning operations. If your array supports the cloning technologies, you can off-load the operations to the array. To avoid the ESXi native cloning, specify the `-N|--avoidnativeclone` option.

Example for Cloning or Converting a Virtual Disk

This example illustrates cloning the contents of a master virtual disk from the `templates` repository to a virtual disk file named `myOS.vmdk` on the `myVMFS` file system.

```
vmkfstools -i /vmfs/volumes/myVMFS/templates/gold-master.vmdk /vmfs/volumes/myVMFS/myOS.vmdk
```

You can configure a virtual machine to use this virtual disk by adding lines to the virtual machine configuration file, as in the following example:

```
scsi0:0.present = TRUE
scsi0:0.fileName = /vmfs/volumes/myVMFS/myOS.vmdk
```

If you want to convert the format of the disk, use the `-d|--diskformat` suboption.

This suboption is useful when you import virtual disks in a format not compatible with ESXi, for example 2gbparse format. After you convert the disk, you can attach this disk to a new virtual machine you create in ESXi.

For example:

```
vmkfstools -i /vmfs/volumes/myVMFS/templates/gold-master.vmdk /vmfs/volumes/myVMFS/myOS.vmdk -d thin
```

Extending a Virtual Disk

After you create a virtual machine, you can use the `vmkfstools` command to extend the size of a disk allocated to the virtual machine.

```
-X|--extendvirtualdisk newSize[bBsSkKmMgGtT]
```

Specify the `newSize` parameter adding an appropriate unit suffix. The unit type is not case-sensitive. `vmkfstools` interprets either `k` or `K` to mean kilobytes. If you do not specify the unit type, `vmkfstools` defaults to kilobytes.

The `newSize` parameter defines the entire new size, not just the increment you add to the disk.

For example, to extend a 4-g virtual disk by 1 g, enter: `vmkfstools -X 5g disk name`.

You can extend the virtual disk to the `eagerzeroedthick` format by using the `-d eagerzeroedthick` option.

When you use the `-X` option, the following considerations apply:

- Do not extend the base disk of a virtual machine that has snapshots associated with it. If you do, you can no longer commit the snapshot or revert the base disk to its original size.
- After you extend the disk, you might need to update the file system on the disk. As a result, the guest operating system recognizes the new size of the disk and can use it.

Upgrading Virtual Disks

This option converts the specified virtual disk file from ESX Server 2 formats to the ESXi format.

Use this option to convert virtual disks of type `LEGACYSPARSE`, `LEGACYPLAIN`, `LEGACYVMFS`, `LEGACYVMFS_SPARSE`, and `LEGACYVMFS_RDM`.

```
-M|--migratevirtualdisk
```

Creating a Virtual Compatibility Mode Raw Device Mapping

Use the `vmkfstools` command to create a Raw Device Mapping (RDM) file on a VMFS volume and map a raw LUN to this file. After this mapping is established, you can access the LUN as you would a normal VMFS virtual disk. The file length of the mapping is the same as the size of the raw LUN it points to.

```
-r|--createrdm device
```

When specifying the *device* parameter, use the following format:

```
/vmfs/devices/disks/disk_ID:P
```

Example for Creating a Virtual Compatibility Mode RDM

In this example, you create an RDM file named *my_rdm.vmdk* and map the *disk_ID* raw disk to that file.

```
vmkfstools -r /vmfs/devices/disks/disk_ID my_rdm.vmdk
```

You can configure a virtual machine to use the *my_rdm.vmdk* mapping file by adding the following lines to the virtual machine configuration file:

```
scsi0:0.present = TRUE
scsi0:0.fileName = /vmfs/volumes/myVMFS/my_rdm.vmdk
```

Creating a Physical Compatibility Mode Raw Device Mapping

Use the `vmkfstools` command to map a pass-through raw device to a file on a VMFS volume. With the mapping, a virtual machine can bypass ESXi SCSI command filtering when accessing its virtual disk. This type of mapping is useful when the virtual machine must send proprietary SCSI commands, for example, when SAN-aware software runs on the virtual machine.

```
-z|--createrrdmpassthru deviceexample.vmdk
```

After you establish this type of mapping, you can use it to access the raw disk as you access any other VMFS virtual disk.

When specifying the *device* path, use the following format:

```
/vmfs/devices/disks/device_ID
```

For the *.vmdk* name, use this format. Make sure to create the datastore before using the command.

```
/vmfs/volumes/datastore_name/example.vmdk
```

For example,

```
vmkfstools -z /vmfs/devices/disks/naa.600a000000000000... /vmfs/volumes/datastore1/mydisk.vmdk
```

Listing Attributes of an RDM

Use the `vmkfstools` command to list the attributes of a raw disk mapping. The attributes help you identify the storage device to which your RDM files maps.

```
-q|--queryrdm my_rdm.vmdk
```

This option prints the name of the raw disk RDM. The option also prints other identification information, like the disk ID, for the raw disk.

Example of Listing RDM Attributes

```
# vmkfstools -q /vmfs/volumes/VMFS/my_vm/my_rdm.vmdk

Disk /vmfs/volumes/VMFS/my_vm/my_rdm.vmdk is a Passthrough Raw Device Mapping

Maps to: vml.02000000006005076801900207700000000000005323134352020
```

Displaying Virtual Disk Geometry

Use the vmkfstools command to get information about the geometry of a virtual disk.

```
-g|--geometry
```

The output is in the form: Geometry information C/H/S , where C represents the number of cylinders, H represents the number of heads, and S represents the number of sectors.

Note: When you import virtual disks from hosted VMware products to the ESXi host, you might see a disk geometry mismatch error message. A disk geometry mismatch might also trigger problems when you load a guest operating system or run a newly created virtual machine.

Checking and Repairing Virtual Disks

Use the vmkfstools command to check or repair a virtual disk if it gets corrupted.

```
-x|--fix [check|repair]
```

For example,

```
vmkfstools -x check /vmfs/volumes/my_datastore/my_disk.vmdk
```

Checking Disk Chain for Consistency

Use the vmkfstools command to check the entire snapshot chain. You can determine if any of the links in the chain are corrupted or any invalid parent-child relationships exist.

```
-e|--chainConsistent
```

Storage Device Options of vSphere vmkfstools Command

You can use the device options of the vSphere vmkfstools command to perform administrative task for physical storage devices.

Managing SCSI Reservations of LUNs

Use the vmkfstools command to reserve a SCSI LUN for exclusive use by the ESXi host. You can also release a reservation so that other hosts can access the LUN, and reset a reservation, forcing all reservations from the target to be released.

```
-L|--lock [reserve|release|lunreset|targetreset|busreset|readkeys|readresv] device
```

CAUTION: Using the `-L` option can interrupt the operations of other servers on a SAN. Use the `-L` option only when troubleshooting clustering setups.

Unless advised by VMware, never use this option on a LUN hosting a VMFS volume.

You can specify the `-L` option in several ways:

- `-L reserve` – Reserves the specified LUN. After the reservation, only the server that reserved that LUN can access it. If other servers attempt to access that LUN, a reservation error appears.
- `-L release` – Releases the reservation on the specified LUN. Other servers can access the LUN again.
- `-L lunreset` – Resets the specified LUN by clearing any reservation on the LUN and making the LUN available to all servers again. The reset does not affect any of the other LUNs on the device. If another LUN on the device is reserved, it remains reserved.
- `-L targetreset` – Resets the entire target. The reset clears any reservations on all the LUNs associated with that target and makes the LUNs available to all servers again.
- `-L busreset` – Resets all accessible targets on the bus. The reset clears any reservation on all the LUNs accessible through the bus and makes them available to all servers again.
- `-L readkeys` – Reads the reservation keys registered with a LUN. Applies to SCSI-III persistent group reservation functionality.
- `-L readresv` – Reads the reservation state on a LUN. Applies to SCSI-III persistent group reservation functionality.

When entering the *device* parameter, use the following format:

```
/vmfs/devices/disks/disk_ID:P
```

Breaking Device Locks

Use the `vmkfstools` command to break the device lock on a particular partition.

```
-B|--breaklock device
```

When entering the *device* parameter, use the following format:

```
/vmfs/devices/disks/disk_ID:P
```

You can use this command when a host fails in the middle of a datastore operation, such as expand the datastore, add an extent, or resignature. When you run this command, make sure that no other host is holding the lock.

Getting Started with Cloud Native Storage in vSphere

Cloud Native Storage is a solution that provides comprehensive data management for stateful applications in vSphere. When you use Cloud Native Storage, you can create containerized stateful applications capable of surviving restarts and outages. Stateful containers leverage storage exposed by vSphere while using such primitives as standard volume, persistent volume, and dynamic provisioning.


With Cloud Native Storage, you can create persistent container volumes independent of virtual machine and container life cycle. vSphere storage backs the volumes, and you can set a storage policy directly on the volumes. After you create the volumes, you can review them and their backing storage objects in the vSphere Client, and monitor their storage policy compliance.

vSphere Cloud Native Storage supports persistent volumes in the following Kubernetes distributions:

- Generic Kubernetes, also called vanilla, that you install from the official repositories. This *vSphere Storage* documentation covers only generic Kubernetes.
- vSphere Supervisor, formerly called vSphere with Tanzu. For more information, see the documentation.

Cloud Native Storage Concepts and Terminology in vSphere

Be familiar with several concepts essential to the Cloud Native Storage environment in vSphere.



Kubernetes Cluster

In the Cloud Native Storage environment, you can deploy a generic Kubernetes cluster in a cluster of virtual machines. On top of the Kubernetes cluster, you deploy your containerized applications. Applications can be stateful and stateless.

Note: For information on a Supervisor and TKG cluster that you can run in vSphere Supervisor environment, see the documentation.

Pod

A pod is a group of one or more containerized applications that share such resources as storage and network. Containers inside a pod are started, stopped, and replicated as a group.

Container Orchestrator

Open-source platforms, such as Kubernetes, for deployment, scaling, and management of containerized applications across clusters of hosts. The platforms provide a container-centric infrastructure.

Stateful Application

As containerized applications evolve from stateless to stateful, they require persistent storage. Unlike stateless applications that do not save data between sessions, stateful applications save data to persistent storage. The retained data is called the application's state. You can later retrieve the data and use it in the next session.

PersistentVolume

Most applications are stateful. A database is as an example of a stateful application.

Stateful applications use PersistentVolumes to store their data. A PersistentVolume is a Kubernetes volume capable of retaining its state and data. It is independent of a pod and can continue to exist even when the pod is deleted or reconfigured. In the vSphere environment, the PersistentVolume objects use vSphere virtual disks of the First Class Disk (FCD) type or vSAN file shares as their backing storage. First Class Disks are also referred to as Improved Virtual Disks (IVD) or managed virtual disks.

- Virtual disks support volumes that are mounted as ReadWriteOnce. These volumes can be used only by a single Pod in Kubernetes.

You can use the vSphere encryption technology to protect FCD virtual disks that back persistent volumes. For more information, see [Use Encryption with vSphere Cloud Native Storage](#).

- vSAN file shares support ReadWriteMany volumes that are mounted by many nodes. These volumes can be shared between multiple Pods or applications running across Kubernetes nodes or across Kubernetes clusters. For information about possible configurations with file shares, see [Using vSAN File Service to Provision vSphere File Volumes](#).

StorageClass

Kubernetes uses a StorageClass to define different tiers of storage and to describe different types of requirements for storage backing the PersistentVolume. In the vSphere environment, a storage class can be linked to a storage policy. As a vSphere administrator, you create storage policies that describe different storage requirements. The VM storage policies can be used as a part of StorageClass definition for dynamic volume provisioning.

The following sample YAML file references the `Gold` storage policy that you created earlier using the vSphere Client. The resulting persistent volume VMDK is placed on a compatible datastore that satisfies the `Gold` storage policy requirements.

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: gold-sc
  annotations:
    storageclass.kubernetes.io/is-default-class:
      "true"
provisioner: csi.vsphere.vmware.com
parameters:
  storagepolicyname: "Gold"
```

PersistentVolumeClaim

Typically, applications or pods can request persistent storage through a PersistentVolumeClaim. The PersistentVolumeClaim specifies the type and class of storage, the access mode, either ReadWriteOnce or ReadWriteMany, and other parameters for the PersistentVolume. The request can then dynamically provision the corresponding PersistentVolume object and the underlying virtual disk or vSAN file share in the vSphere environment.

Once the claim is created, the PersistentVolume is automatically bound to the claim. Pods use the claim to mount the PersistentVolume and access storage.


When you delete this claim, the corresponding PersistentVolume object and the underlying storage are deleted.

```
kind: PersistentVolumeClaim
metadata:
  name: persistent-VMDK
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 5Gi
  storageClassName: gold-sc
```

Cloud Native Storage Components

Cloud Native Storage uses several components to integrate with vSphere storage.

The following illustration shows how these components interact.



Kubernetes Cluster

Container Storage Interface (CSI) for vSphere

In the Cloud Native Storage environment, a generic Kubernetes cluster is deployed on a cluster of virtual machines, or nodes, that run in vSphere. A Kubernetes user directly interacts with the cluster when deploying stateful applications on top of it. To consume underlying infrastructure resources, the cluster requires a CSI driver.

The vSphere CSI is an out-of-tree plug-in that exposes vSphere storage to containerized workloads on container orchestrators, such as Kubernetes. The plug-in enables vSAN and other types of vSphere storage.

The vSphere CSI communicates with the CNS control plane on vCenter Server for all storage provisioning operations. The vSphere CSI supports the following functionalities:

- Dynamic provisioning of container volumes.
- The vSphere First Class Disk functionality.
- Kubernetes zones.
- Conventional and raw mounts.
- Single vCenter Server, and multiple data centers and clusters.
- Provisioning from multiple datastores or datastore clusters.

- vSAN File Service

On Kubernetes, the CSI driver is used with the out-of-tree vSphere Cloud Provider Interface (CPI). The CSI driver is shipped as a container image and must be deployed by the cluster administrator. For information, see the *vSphere Container Storage Plug-in Deployment* section of the *Getting Started with VMware vSphere Container Storage Plug-in* documentation.

The CNS server component, or the CNS control plane, resides in vCenter Server. It is an extension of vCenter Server management that implements the provisioning and life cycle operations for the container volumes.

When provisioning container volumes, it interacts with vCenter Server to create storage objects that back the volumes. The Storage Policy Based Management functionality guarantees a required level of service to the volumes.

The CNS also performs query operations that allow you to manage and monitor container volumes and their backing storage objects through vCenter Server.

Also called Improved Virtual Disk (IVD) or managed virtual disk. It is a named virtual disk unassociated with a VM. These disks reside on a vSAN, VMFS, NFS, or vVols datastore and back ReadWriteOnce container volumes.

The FCD technology allows to perform life cycle operations related to persistent volumes outside of the VM or pod life cycle. If the VM is a Kubernetes node that runs multiple container based applications and uses persistent volumes and virtual disks for many applications, CNS facilitates life cycle operations at the container and persistent volume granularity.

It is a vSAN layer that provides file shares. Currently, it supports NFSv3 and NFSv4.1 file shares. Cloud Native Storage uses vSAN file shares for persistent volumes of the ReadWriteMany type. A single ReadWriteMany volume can be mounted by multiple nodes. The volume can be shared between multiple pods or applications running across Kubernetes nodes or across Kubernetes clusters. Storage Policy Based Management is a vCenter Server service that supports provisioning of persistent volumes according to specified storage requirements. After provisioning, the service monitors compliance of the volume with the required policy characteristics.

Cloud Native Storage Server Component

First Class Disk (FCD)

vSAN File Service

Storage Policy Based Management

Using vSAN File Service to Provision vSphere File Volumes

vSAN file service offers vSAN file shares that are consumed by persistent volumes of the ReadWriteMany (RWM) type. A single RWM volume can be mounted by multiple nodes. The volume can be shared between multiple pods or applications running across Kubernetes nodes or across Kubernetes clusters in vSphere.

When a Kubernetes pod request an RWM volume, Cloud Native Storage communicates with vSAN file service to create an NFS-based file share of the requested size and storage class. Cloud Native Storage then mounts the RWM volume into the Kubernetes worker node where the pod runs. If multiple nodes are requesting access to the RWM volume, Cloud Native Storage determines that the RWM volume already exists for that particular deployment and mounts the existing volume into the nodes.

To be able to support RWM volumes, your environment must include the following items.


- vSphere 7.0 and later with vSAN

- vSAN file service enabled. For information, see *vSAN File Service* in the *Administering VMware vSAN* documentation.
- Kubernetes version 1.14 and later
- Compatible version of CSI. For information, see the *Getting Started with VMware vSphere Container Storage Plug-in* documentation.

You can use different configurations for file volumes.


Single File Volume Shared Across Applications in the Same Namespace

In this example, a single file volume is used as shared storage across different applications in the same namespace. You use a single persistent volume claim to provision the file volume.



Single File Volume Shared Across Applications and Namespaces

This example uses a single file volume as shared storage across different applications and different namespaces. For each namespace, you create a separate persistent volume claim to provision the same file volume.



vSphere Cloud Native Storage Users

The types of users involved in the process of creating and monitoring Kubernetes volumes in the vSphere Cloud Native Storage environment generally fall into two categories, a Kubernetes user and a vSphere administrator. Both types of users have access to different tools and perform different tasks.

CNS Kubernetes User

The Kubernetes user might be a Kubernetes developer and an application owner, a Kubernetes administrator, or combine functions of both. The tasks that the Kubernetes user performs in the Cloud Native Storage environment include the following:

- Deploy and manage the vSphere CSI. For information, see the [vSphere Container Storage Plug-in Deployment](#) section of the [Getting Started with VMware vSphere Container Storage Plug-in](#) documentation.
- Provision persistent volumes. For information about block volumes, see [vSphere CSI Driver - Block Volume](#). For information about file volumes, see [vSphere CSI Driver - File Volume](#).
- Perform life cycle operations for persistent volumes.
- Perform life cycle operations for storage classes.

CNS vSphere User

A CNS vSphere user, or a vSphere administrator, has access to the vSphere Client to perform the following tasks:

- Perform life cycle operations for the VM storage policies. For example, create a VM storage policy to be used for a Kubernetes storage class and communicate its name to the Kubernetes user. See [Create a Storage Policy for Kubernetes in vSphere](#).
- Use the Cloud Native Storage section of the vSphere Client to monitor health and storage policy compliance of the container volumes across the Kubernetes clusters. See [Monitor Container Volumes Across Kubernetes Clusters in vSphere](#).

Cloud Native Storage for vSphere Administrators

A vSphere administrator delivers storage resources to the Kubernetes team and creates VM storage policies that describe different storage requirements and classes of services. After the Kubernetes workloads with persistent storage are provisioned, the vSphere administrator can monitor the life cycle of the backing storage resources and their compliance to the requirements.

Requirements of vSphere Cloud Native Storage

Your Cloud Native Storage environment and virtual machines that participate in the Kubernetes cluster must meet several requirements.

Cloud Native Storage Requirements

- A compatible version of vSphere.
- A compatible version of Kubernetes.
- A Kubernetes cluster deployed on the virtual machines. For details about deploying the vSphere CSI plug-in and running the Kubernetes cluster on vSphere, see the [VMware vSphere Container Storage Plug-in Documentation](#).

Requirements for Kubernetes Cluster Virtual Machines

- Virtual machines with hardware version 15 or later. Install VMware Tools on each node virtual machine.
- Virtual machine hardware recommendations:
 - Set CPU and memory adequately based on workload requirements.
 - Use the VMware Paravirtual SCSI controller for the primary disk on the Node VM.
- All virtual machines must have access to a shared datastore, such as vSAN.
- Set the `disk.EnableUUID` parameter on each node VM. See [Configure Kubernetes Cluster in vSphere Virtual Machines](#).
- To avoid errors and unpredictable behavior, do not take snapshots of CNS node VMs.

Requirements for CNS File Volume

- Use vSphere version 7.0 or later with a compatible Kubernetes version.
- Use a compatible version of CSI. For information, see the [VMware vSphere Container Storage Plug-in Documentation](#).
- Enable and configure the vSAN file service. You must configure the necessary file service domains, IP pools, network, and so on. For information, see the *Administering VMware vSAN* documentation.
- Follow specific guidelines to configure network access from a guest OS in the Kubernetes node to a vSAN file share. See [Configuring Network Access to vSAN File Share](#).

Configuring Network Access to vSAN File Share

To be able to provision ReadWriteMany persistent volumes in your generic vSphere Kubernetes environment, configure necessary networks, switches, and routers from the Kubernetes nodes to the vSAN file service network.


Setting Up Network

When configuring the networks, follow these requirements:

- On every Kubernetes node, you can use a dedicated vNIC for the vSAN file share traffic. This option is required only if you want to use a secure data traffic path for your file volumes.

- If you use a dedicated vNIC, make sure that the traffic through the dedicated vNIC is routable to one or many vSAN file service networks.
- Make sure that only the guest OS on each Kubernetes node can directly access the vSAN file share through the file share IP address. The pods in the node cannot ping or access the vSAN file share by its IP address.
CNS CSI driver ensures that only those pods that are configured to use the CNS file volume can access the vSAN file share by creating a mount point in the guest OS.
- Avoid creating an IP address clash between the node VMs and vSAN file shares.

The following illustration is an example of the CNS network configuration with the vSAN file share service.



In the illustration, the sample networking configuration follows these guidelines.

- The configuration uses separate networks for different items in the CNS environment.

Network	Description
vSphere management network	Typically, in a generic Kubernetes cluster, every node has access to this network.
Pod or node network	Kubernetes uses this network for the node to node or pod to pod communication.
Dedicated file share network	CNS file volume data traffic uses this network.

Network	Description
vSAN file share network	Network where the vSAN file share is enabled and where file shares are available.

- Every Kubernetes node has a dedicated vNIC for the file traffic. This vNIC is separate from the vNIC used for the node to node or pod to pod communication. This configuration is used only as an example, but is not mandatory.
- Only those applications that are configured to use the CNS file share have access to vSAN file shares through the mount point in the node guest OS. For example, in the illustration, the following takes place:
 - App-1 and App-2 pods are configured to use a file volume, and have access to the file share through the mount point created by the CSI driver.
 - App-3 and App-4 are not configured with a file volume and cannot access file shares.
- The vSAN file shares are deployed as containers in a vSAN file share appliance VM on the ESXi host. A Kubernetes deployer, which is a software or service that can configure, deploy, and manage Kubernetes clusters, configures necessary routers and switches, so that the guest OS in the Kubernetes node can access the vSAN file shares.

Security Limitations

Although the dedicated vNIC prevents an unauthorized pod from accessing the file shares directly, certain security limitations exist:

- The CNS file functionality assumes that anyone who has the CNS file volume ID is an authorized user of the volume. Any user that has the CNS file volume ID can access the data stored in the volume.
- CNS file volume supports only the AUTH_SYS authentication, which is a user ID-based authentication. To protect access to the data in the CNS file volume, you must use appropriate user IDs for the containers accessing the CNS file volume.
- An unbound ReadWriteMany persistent volume referring to a CNS file volume can be bound by a persistent volume claim created by any Kubernetes user under any namespace. Make sure that only authorized users have access to Kubernetes to avoid security issues.

Configuring the CSI Driver to Access vSAN File Service Clusters

Depending on the configuration, the CSI driver can provision file volumes on one or several vSAN clusters where the file service is enabled.

You can restrict access to only specific vSAN clusters where the file service is enabled. When deploying the Kubernetes cluster, configure the CSI driver with access to specific file service vSAN clusters. As a result, the CSI driver can provision the file volumes only on those vSAN clusters.

In the default configuration, the CSI driver uses any file service vSAN cluster available in vCenter Server for the file volume provisioning. The CSI driver does not verify which file service vSAN cluster is accessible while provisioning file volumes.

vSphere Cloud Native Storage Roles and Privileges

The CNS vSphere user must have specific privileges to perform operations related to Cloud Native Storage.

You can create several roles to assign sets of permissions on the objects that participate in the Cloud Native Storage environment.

Note: These roles need to be created only for generic Kubernetes clusters. If you work in the vSphere Supervisor environment, use the Workload Storage Manager role for storage operations.

For more information about roles and permissions in vSphere, and how to create a role, see the *vSphere Security* documentation.

Role Name	Privilege Name	Description	Required On
CNS-Datastore	Datastore > Low level file operations	Allows performing read, write, delete, and rename operations in the datastore browser.	Shared datastore where persistent volumes reside.
CNS-HOST-CONFIG-STORAGE	Host > Configuration > Storage partition configuration	Allows vSAN datastore management.	Required on a vSAN cluster with vSAN file service. Required for file volume only.
CNS-VM	Virtual machine > Change Configuration > Add existing disk	Allows adding an existing virtual disk to a virtual machine.	All cluster node VMs.
	Virtual Machine > Change Configuration > Add or remove device	Allows addition or removal of any non-disk device.	
CNS-SEARCH-AND-SPBM	CNS > Searchable	Allows storage administrator to see Cloud Native Storage UI.	Root vCenter Server.
	VM storage policies > View VM storage policies	Allows viewing of defined storage policies.	
Read-only	Default role	Users with the Read Only role for an object are allowed to view the state of the object and details about the object. For example, users with this role can find the shared datastore accessible to all node VMs. For zone and topology-aware environments, all ancestors of node VMs, such as a host, cluster, and data center must have the Read-only role set for the vSphere user configured to use the CSI driver and CCM. This is required to allow reading tags and categories to prepare the nodes' topology.	All hosts where the nodes VMs reside Data center

Create a Storage Policy for Kubernetes in vSphere

The vSphere storage object that will back a Kubernetes containerized application should meet specific storage requirements. As a vSphere user, you create a VM storage policy based on the requirements given by the Kubernetes user.

- The Kubernetes user identifies the Kubernetes cluster where the stateful containerized application will be deployed.
- The Kubernetes user collects storage requirements for the containerized application and communicates them to the vSphere user.
- Required privileges: **VM storage policies > Update** and **VM storage policies > View**.

The storage policy will be associated with the virtual disk or vSAN file share that back the Kubernetes container. If you have multiple vCenter Server instances in your environment, create the VM storage policy on each instance. Use the same policy name across all instances.

If you want to create a storage policy for a datastore other than vSAN, see [Creating and Managing vSphere VM Storage Policies](#). To use thick provisioning on a VMFS datastore, make sure to select an appropriate option in the **Volume allocation** drop-down menu when you define the policy.

For information about types of datastores that Cloud Native Storage supports and which datastores support thick provisioning, see *vSphere Functionality Supported by vSphere Container Storage Plug-in* in *Getting Started with VMware vSphere Container Storage Plug-in*.

1. In the vSphere Client, open the **Create VM Storage Policy** wizard.

- Click **Menu > Policies and Profiles**.
- Under **Policies and Profiles**, click **VM Storage Policies**.
- Click **Create**.

2. Enter the policy name and description, and click **Next**.

Option	Action
vCenter Server	Select the vCenter Server instance.
Name	Enter the name of the storage policy, for example <i>Space-Efficient</i> .
Description	Enter the description of the storage policy.

3. On the **Policy structure** page under Datastore-specific rules, select **Enable rules for vSAN storage** and click **Next**.

4. On the **vSAN** page, define the policy rule set and click **Next**.

- On the **Availability** tab, define the **Site disaster tolerance** and **Failures to tolerate**.
- On the **Advanced Policy Rules** tab, define advanced policy rules, such as number of disk stripes per object and flash read cache reservation.

5. On the **Storage compatibility** page, review the list of vSAN datastores that match this policy and click **Next**.

6. On the **Review and finish** page, review the policy settings, and click **Finish**.

Edit VM Storage Policy

1 Name and description
2 Policy structure
3 vSAN
4 Storage compatibility
5 Review and finish

Review and finish

General
Name
Description
vCenter Server

Space-Efficient
sc2-rdops-vm08-dhcp-23-199.eng.vmware.com

vSAN
Availability
Site disaster tolerance
Failures to tolerate
Advanced Policy Rules
Number of disk stripes per object
IOPS limit for object
Object space reservation
Flash read cache reservation
Disable object checksum
Force provisioning

None - standard cluster
No data redundancy
1
0
Thin provisioning
0%
No
No

CANCEL
BACK
FINISH

You can now inform the Kubernetes user of the storage policy name. The VM storage policy you created will be used as a part of storage class definition for dynamic volume provisioning.

Configure Kubernetes Cluster in vSphere Virtual Machines

On each node VM, enable the `disk.EnableUUID` parameter, so that the vSphere VMs can successfully mount the virtual disks.

- Create several VMs for your Kubernetes cluster. For the VM requirements, see [Requirements of vSphere Cloud Native Storage](#).
- Required privilege: **Virtual machine > Configuration > Settings**.

Note: To avoid errors and unpredictable behavior, do not take snapshots of CNS node VMs.

Perform these steps for each of the VM nodes that participate in the cluster.

1. In the vSphere Client, right-click the virtual machine and select **Edit Settings**.
2. Click the **VM Options** tab and expand the **Advanced** menu.
3. Click **Edit Configuration** next to Configuration Parameters.
4. Configure the `disk.EnableUUID` parameter.

If the parameter exists, make sure that its value is set to True. If the parameter is not present, add it and set its value to True.


Name	Value
<code>disk.EnableUUID</code>	True

Monitor Container Volumes Across Kubernetes Clusters in vSphere

After a stateful application gets deployed in Kubernetes, the volumes and their backing vSphere storage objects become visible in the vSphere Client. You can display and monitor the volumes and troubleshoot any potential storage issues.

Note: If you experience failures on the Kubernetes CNS server, the CNS objects in the vSphere Client might not display correctly until full synchronization takes place.

1. Navigate to the vCenter Server instance, a data center, or a datastore.
2. Click the **Monitor** tab and click **Container Volumes** under **Cloud Native Storage**.
3. Observe the container volumes available in your environment and monitor their storage policy compliance status.




- Click the **SEE ALL** link in the Label column to view additional details.

The details include the name of the PersistentVolumeClaim, StorageClass, and so on, and help you map the volume to the Kubernetes objects associated with it.

- Click the link in the **Volume Name** column to review various components that back the volume and such details as placement, compliance, and storage policy.

Note: The **Virtual Objects** screen is available only when the underlying datastore is vSAN.



Use Encryption with vSphere Cloud Native Storage

You can use the vSphere encryption technology to protect FCD virtual disks that back persistent volumes.

Using encryption in your vSphere environment requires some preparation, and includes setting up a trusted connection between vCenter Server and a key provider. vCenter Server can then retrieve keys from the key provider as needed. For information about components that participate in the vSphere encryption process, see [vSphere Virtual Machine Encryption Components](#) in the *vSphere Security* documentation.

- Set up the key provider in your vSphere environment.

For information, see [Set up the Key Management Server Cluster](#).

- Encrypt all node VMs on the Kubernetes cluster.

Use the vSphere Client to perform this step.

- Navigate to a node VM.
- From the right-click menu, select **VM Policies > Edit VM Storage Policies**.
- From the **VM storage policy** drop-down menu, select **VM Encryption Policy** and click **OK**.

To expedite the encryption process of the node VMs, you can encrypt only the VM home.

3. Create encrypted persistent volumes in the Kubernetes cluster with the vSphere CSI setup.

a) Create a StorageClass that references the VM Encryption Storage Policy.

Use the following YAML file as an example.

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: encryption
provisioner: csi.vsphere.vmware.com
parameters:
  storagePolicyName: "VM Encryption Policy"
  datastore: vsanDatastore
```

b) Use the PersistentVolumeClaim to provision the persistent volume.

The PersistentVolumeClaim must include the name of the encryption storage class in the `storageClassName` field.

Migrating Container Volumes in vSphere

vSphere Cloud Native Storage supports migration of persistent container volumes between datastores.

Considerations and Limitations

When you perform the migration, keep in mind the following considerations:

- You can migrate only single block volumes.
- Make sure that original and target datastores are accessible to all nodes of the Kubernetes cluster. This applies to single zone deployments and to Kubernetes clusters deployed across multiple zones.
- If a volume is attached to a VM, the VM remains locked during volume relocation. If the volume you relocate is attached to a Kubernetes node VM, the system does not allow any other control operations on the VM. The operations include attaching or detaching other volumes to the VM, migrating the VM to a different datastore, changing any configuration of the VM, and so on.
- Cloud Native Storage does not support PMEM storage. Attempts to relocate a volume to a PMEM datastore fail.
- The target datastore must have sufficient free capacity to accommodate the migrated volume. If the free space of a datastore is less than the volume size, you cannot select the datastore for migration.
- The vSphere Client does not support automatic storage policy changes during the migration operation. However, the system doesn't prevent you from selecting a datastore with incompatible policy as a target datastore for migration.
- After the migration, you can locate the virtual disk that backs the volume in one of the following folders on a target datastore:
 - FCD directory on the datastore.
 - VM folder if the volume is attached to a VM.

For additional details, see the VMware Knowledge Base article [90607](#).

Migrate Container Volumes

You can use the vSphere Client to migrate the persistent container volumes.

- Make sure that the target datastore has sufficient free capacity to accommodate the migrated volume.

- Make sure you have permissions for the target datastore.
1. Navigate to the vCenter Server instance, a data center, a datastore, or a cluster.
 2. Click the **Monitor** tab and click **Container Volumes** under **Cloud Native Storage**.
 3. From the list of volumes, select a volume and click **Migrate**.

The screenshot displays the VMware vSphere Client interface. The top navigation bar includes the 'vSphere Client' logo and a search bar. The left sidebar shows a tree view of the environment, with 'New Cluster 2' selected. The main content area is divided into two sections. The left section contains a sidebar menu with categories like 'Issues and Alarms', 'Performance', 'Tasks and Events', 'Resource Allocation', 'vSphere Cluster Services', and 'Cloud Native Storage'. The 'Cloud Native Storage' category is expanded, showing 'Container Volumes' as the selected option. The right section displays the 'Container Volumes' page, which includes a table of volumes. The table has columns for 'Volume Name' and 'Actions'. One volume, 'vsanVolume', is listed with a checkmark in the 'Volume Name' column and a 'Migrate' button in the 'Actions' column. The bottom status bar shows 'VMware by Broadcom' on the left, a checkmark and the number '1' in the center, and the number '2723' on the right.

VMware by Broadcom

Recent Tasks Alarms

4. On the **Migrate volume** page, select the target datastore.
- The target datastore might display a warning. You can still proceed with the migration operation when the warnings include the following:
- Policy incompatible. When the storage policy of the target datastore is not the same as the storage policy of the volume.
 - Inaccessible volume. When the volume is migrated from a datastore that is accessible from a set of hosts to a datastore that is not accessible from all hosts.

Migrate volume

pv pvc pod ds1 Volume

Volume migration is an advanced operation. [Learn more](#)

Volume storage policy

vSAN Direct

Actual size of the volume

100.00 MB

Select a destination datastore for this volume to be migrated to.

	Destination datastore	Prechecks	Type
☒	vSAND_10.3	Compatible	vSAN Direct
☐	vsanDatastore (1)	Policy incompatible	vSAN Direct

☒

I acknowledge that:

- Volume migration is an advanced operation.

5. Select the **Acknowledge** check box and click **Migrate** to proceed with the migration operation.

vSphere Security

vSphere Security provides information about securing your vSphere® environment for VMware® vCenter® Server and VMware ESXi.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

To help you protect your vSphere environment, this documentation describes available security features and the measures that you can take to safeguard your environment from attack.

Table 229: vSphere Security Highlights

Topics	Content Highlights
Permissions and User Management	- Permissions model (roles, groups, objects). - Creating custom roles. - Setting permissions. - Managing global permissions.
Host Security Features	- Lockdown mode and other security profile features. - Host smart card authentication. - vSphere Authentication Proxy. - UEFI Secure Boot. - Trusted Platform Module (TPM). - VMware® vSphere Trust Authority™. - Secure ESXi Configuration and configuration sealing
Virtual Machine Encryption	- VMware vSphere® Native Key Provider™. - How does VM encryption work? - KMS setup. - Encrypting and decrypting VMs. - Troubleshooting and best practices.
Guest OS Security	- Virtual Trusted Platform Module (vTPM). - Virtualization Based Security (VBS).
Managing TLS Protocol Configuration	Changing TLS protocol configuration using a command-line utility.
Security Best Practices and Hardening	Best practices and advice from VMware security experts. - vCenter Server security - Host security - Virtual machine security - Networking security
vSphere Privileges	Complete listing of all vSphere privileges supported in this release.

Related Documentation

A companion document, *vSphere Authentication*, explains how you can use authentication services, for example, to manage authentication with vCenter Single Sign-On and to manage certificates in your vSphere environment.

In addition to these documents, VMware publishes the *vSphere Security Configuration Guide* (formerly known as the *Hardening Guide*) for each release of vSphere, accessible at <https://core.vmware.com/security>. The *vSphere Security Configuration Guide* contains guidelines on security settings that can or should be set by the customer, and security settings delivered by VMware that should be audited by the customer to ensure that they are still set to default.

What Happened to the Platform Services Controller

Beginning in vSphere 7.0, deploying a new vCenter Server or upgrading to vCenter Server 7.0 requires the use of the vCenter Server appliance, a preconfigured virtual machine optimized for running vCenter Server. The new vCenter Server contains all Platform Services Controller services, preserving the functionality and workflows, including authentication, certificate management, tags, and licensing. It is no longer necessary nor possible to deploy and use an external Platform Services Controller. All Platform Services Controller services are consolidated into vCenter Server, and deployment and administration are simplified.

As these services are now part of vCenter Server, they are no longer described as a part of Platform Services Controller. In vSphere 7.0, the *vSphere Authentication* publication replaces the *Platform Services Controller Administration* publication. The new publication contains complete information about authentication and certificate management. For information about upgrading or migrating from vSphere 6.5 and 6.7 deployments using an existing external Platform Services Controller to vSphere 7.0 using vCenter Server appliance, see the *vSphere Upgrade* documentation.

Intended Audience

This information is for experienced system administrators who are familiar with virtual machine technology and data center operations.

Certifications

VMware publishes a public list of VMware products that have completed Common Criteria certifications. To check if a particular VMware product version has been certified, see the Common Criteria Evaluation and Validation webpage at <https://www.vmware.com/security/certifications/common-criteria.html>.

Updated Information

This *vSphere Security* document is updated with each release of the product or when necessary.

This table provides the update history of the *vSphere Security* documentation.

Revision	Description
25 February 2025	
29 AUG 2024	- Minor update to Setting Timeouts for the ESXi Shell and the vSphere Client. - Updated video links.
23 AUG 2024	- Minor update to Create a Virtual Machine with a Virtual Trusted Platform Module. - Removed broken links.
07 AUG 2024	- Minor update to ESXi Log File Locations. - Minor update to vSphere Native Key Provider Overview. - Minor update to Create a Virtual Machine with a Virtual Trusted Platform Module. - Minor update to vSphere TLS Configuration. - Minor update to VM Storage Policies Privileges.
25 JUL 2024	Minor update to ManagingESXi Audit Records.
24 JUL 2024	- Minor update to Replace the Default Certificate with a Custom Certificate Using the vSphere Client. - Corrected a typo in Activate or Deactivate an ESXi Service. - Clarified how audit records are stored and transmitted in ManagingESXi Audit Records. - Minor update to Defined Privileges.

Revision	Description
25 JUN 2024	Initial release.

Security in the vSphere Environment

The components of a vSphere environment are secured out of the box by several features such as authentication, authorization, a firewall on each ESXi host, and so on. You can modify the default setup in many ways. For example, you can set permissions on vCenter Server objects, open firewall ports, or change the default certificates. You can take security measures for different vSphere objects, for example, vCenter Server systems, ESXi hosts, virtual machines, and network and storage objects.

A high-level overview of different areas of vSphere that require attention helps you plan your security strategy. You also benefit from other vSphere Security resources on the VMware website.

Securing the ESXi Hypervisor

The ESXi hypervisor is secured out of the box. You can further protect ESXi hosts by using lockdown mode and other built-in features. For consistency, you can set up a reference host and keep all hosts in sync with the host profile of the reference host. You can also protect your environment by performing scripted management, which ensures that changes apply to all hosts.

You can enhance protection of ESXi hosts that are managed by vCenter Server with the following actions. Security considerations for standalone hosts are similar, though the management tasks might differ. See the *vSphere Single Host Management - VMware Host Client* documentation.

Limit ESXi Access

By default, the ESXi Shell and the SSH services are not running and only the root user can log in to the Direct Console User Interface (DCUI). If you decide to enable ESXi or SSH access, you can set timeouts to limit the risk of unauthorized access. Users who can access the ESXi host must have permissions to manage the host. You set permissions on the host object from the vCenter Server system that manages the host.

See [Using the ESXi Shell](#).

Use Named Users and Least Privilege

By default, the root user can perform many tasks. Do not allow administrators to log in to the ESXi host using the root user account. Instead, create named administrator users from vCenter Server and assign those users the Administrator role. You can also assign those users a custom role. See [Create a vCenter Server Custom Role](#).

If you manage users directly on the host, role management options are limited. See the *vSphere Single Host Management - VMware Host Client* documentation.

Minimize the Number of Open ESXi Firewall Ports

By default, firewall ports on your ESXi host are opened only when you start a corresponding service. You can use the vSphere Client, or ESXCLI or PowerCLI commands to check and manage firewall port status.

See [Configuring the ESXi Firewall](#).

Automate ESXi Host Management

Because it is often important that different hosts in the same data center are in sync, use scripted installation or vSphere Auto Deploy to provision hosts. You can manage the hosts using scripts. Host profiles are an alternative to scripted management. You set up a reference host, export the host profile, and apply the host profile to all hosts. You can apply the host profile directly or as part of provisioning with Auto Deploy.

See [Use Scripts to Manage ESXi Host Configuration Settings](#) and see the *vCenter Server Installation and Setup* documentation for information about vSphere Auto Deploy.

Take Advantage of ESXi Lockdown Mode

In lockdown mode, ESXi hosts can be accessed only through vCenter Server by default. You can select strict lockdown mode or normal lockdown mode. You can define Exception Users to allow direct access to service accounts such as backup agents.

See [Configuring and Managing Lockdown Mode on ESXi Hosts](#).

Check VIB Package Integrity

Each vSphere Installation Bundle (VIB) package has an associated acceptance level. You can add a VIB to an ESXi host only if the VIB acceptance level is the same or better than the acceptance level of the host. You cannot add a CommunitySupported or PartnerSupported VIB to a host unless you explicitly change the acceptance level of the host.

See [Manage the Acceptance Levels of ESXi Hosts and vSphere Installation Bundles](#).

Manage ESXi Certificates

The VMware Certificate Authority (VMCA) provisions each ESXi host with a signed certificate that has VMCA as the root certificate authority by default. If your company policy requires it, you can replace the existing certificates with certificates that are signed by a third-party or an enterprise certificate authority.

See [Managing Certificates for ESXi Hosts](#).

Consider Smart Card Authentication for ESXi

ESXi supports the use of smart card authentication instead of user name and password authentication. Two-factor authentication is also supported for vCenter Server. You can configure user name and password authentication and smart card authentication at the same time.

See [Configuring and Managing Smart Card Authentication for ESXi](#).

Consider ESXi Account Lockout

Account locking is supported for access through SSH and through the vSphere Web Services SDK. By default, a maximum of five failed attempts is allowed before the account is locked. The account is unlocked after 15 minutes by default.

Note: The Direct Console Interface (DCUI) and the ESXi Shell do not support account lockout.

See [ESXi Passwords and Account Lockout](#).

Securing vCenter Server Systems and Associated Services

Authentication through vCenter Single Sign-On, and authorization through the vCenter Server permissions model, protects your vCenter Server system and associated services. You can modify the default behavior, and you can take steps to limit access to your environment.

As you protect your vSphere environment, consider that all services that are associated with the vCenter Server instances must be protected. In some environments, you might protect several vCenter Server instances.

vCenter Server Uses Encrypted Communication

By default ("out of the box"), all data communication between the vCenter Server system and the other vSphere components is encrypted. In some cases, depending on how you configure your environment, some traffic might be

unencrypted. For example, you can configure unencrypted SMTP for email alerts and unencrypted SNMP for monitoring. DNS traffic is also unencrypted. vCenter Server listens on port 80 (TCP) and port 443 (TCP). Port 443 (TCP) is the industry-standard HTTPS (secure HTTP) port and uses TLS encryption for protection. See [vSphere TLS Configuration](#). Port 80 (TCP) is the industry-standard HTTP port and does not use encryption. The purpose of port 80 is to redirect requests from port 80 to port 443, where they are secure.

Harden vCenter Server Systems

The first step in protecting your vCenter Server environment is hardening each machine on which vCenter Server or an associated service runs. Similar considerations apply to a physical machine or a virtual machine. Always install the latest security patches for your operating system and follow industry standard best practices to protect the host machine.

Learn About the vSphere Certificate Model

By default, the VMware Certificate Authority (VMCA) provisions each ESXi host and each machine in the environment with a certificate signed by VMCA. If your company policy requires it, you can change the default behavior. See the *vSphere Authentication* documentation for details.

For additional protection, explicitly remove expired or revoked certificates and failed installations.

Configure vCenter Single Sign-On

vCenter Server and associated services are protected by the vCenter Single Sign-On authentication framework. When you first install the software, you specify a password for the administrator of the vCenter Single Sign-On domain, administrator@vsphere.local by default. Only that domain is initially available as an identity source. You can add an external identity provider, such as Microsoft Active Directory Federation Services (AD FS), for federated authentication. You can add other identity sources, either Active Directory or LDAP, and set a default identity source. Users who can authenticate to one of those identity sources can view objects and perform tasks if they are authorized to do so. See the *vSphere Authentication* documentation for details.

Note: You are encouraged to use federated authentication as vSphere moves towards token-based authentication. vCenter Server continues to have local accounts, for administrative access and error recovery.

Assign vCenter Server Roles to Named Users or Groups

For better logging, associate each permission that you give on an object with a named user or group and a predefined role or custom role. The vSphere permissions model allows great flexibility through multiple ways of authorizing users or groups. See [Understanding Authorization in vSphere](#) and [Required vCenter Server Privileges for Common Tasks](#).

Restrict administrator privileges and the use of the administrator role. If possible, do not use the anonymous Administrator user.

Set up Precision Time Protocol or Network Time Protocol

Set up Precision Time Protocol (PTP) or Network Time Protocol (NTP) for each node in your environment. The vSphere certificate infrastructure requires an accurate time stamp and does not work correctly if the nodes are out of sync.

See [Synchronizing Clocks on the vSphere Network](#).

Securing Virtual Machines

To secure your virtual machines, keep the guest operating systems patched and protect your virtual environment just as you protect your physical machine. Consider deactivating unnecessary functionality, minimize the use of the virtual machine console, and follow other best practices.

Protect the Guest Operating System

To protect your guest operating system, make sure that it uses the most recent patches and, if appropriate, anti-spyware and anti-malware applications. See the documentation from your guest operating system vendor and, potentially, other information available in books or on the Internet for that operating system.

Deactivate Unnecessary Virtual Machine Functionality

Check that unnecessary functionality is deactivated to minimize potential points of attack. Many of the features that are used infrequently are deactivated by default. Remove unnecessary hardware and deactivate certain features such as host-guest filesystem (HGFS) or copy and paste between the virtual machine and a remote console.

See [Deactivate Unnecessary Functions Inside Virtual Machines](#).

Use Virtual Machine Templates and Scripted Management

Virtual machine templates enable you to set up the operating system so that it meets your requirements, and to create other virtual machines with the same settings.

If you want to change virtual machine settings after initial deployment, consider using PowerCLI scripts. For the most part, this documentation explains how to perform tasks using the vSphere Client. Consider using scripts instead of the vSphere Client to keep your environment consistent. In large environments, you can group virtual machines into folders to optimize scripting.

For information on templates, see [Use Templates to Deploy Virtual Machines](#) and the *vSphere Virtual Machine Administration* documentation. For information on PowerCLI, see the VMware PowerCLI documentation.

Minimize Use of the Virtual Machine Console

The virtual machine console provides the same function for a virtual machine that a monitor on a physical server provides. Users with access to a virtual machine console have access to virtual machine power management and to removable device connectivity controls. As a result, virtual machine console access might allow a malicious attack on a virtual machine.

Consider UEFI Secure Boot for Virtual Machines

You can configure your virtual machines to use UEFI boot. If the operating system supports secure UEFI boot, you can select that option for your virtual machines for additional security. See [Activate or Deactivate UEFI Secure Boot for a Virtual Machine](#).

Securing the Virtual Networking Layer

The virtual networking layer includes virtual network adapters, virtual switches, distributed virtual switches, and ports and port groups. ESXi relies on the virtual networking layer to support communications between virtual machines and their users. In addition, ESXi uses the virtual networking layer to communicate with iSCSI SANs, NAS storage, and so on.

vSphere includes the full array of features necessary for a secure networking infrastructure. You can secure each element of the infrastructure, such as virtual switches, distributed virtual switches, and virtual network adapters, separately. In addition, consider the following guidelines, discussed in more detail in [Securing vSphere Networking](#).

Isolate Network Traffic

Isolation of network traffic is essential to a secure ESXi environment. Different networks require different access and level of isolation. A management network isolates client traffic, command-line interface (CLI) or API traffic, and third-party software traffic from normal traffic. Ensure that the management network is accessible only by system, network, and security administrators.

See [ESXi Networking Security Recommendations](#).

Use Firewalls to Secure Virtual Network Elements

You can open and close firewall ports and secure each element in the virtual network separately. For ESXi hosts, firewall rules associate services with corresponding firewalls and can open and close the firewall according to the status of the service.

You can also open ports on vCenter Server instances explicitly.

For the list of all supported ports and protocols in VMware products, including vSphere and , see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com/>. You can search ports by VMware product, create a customized list of ports, and print or save port lists.

Consider Network Security Policies

Network security policies provide protection of traffic against MAC address impersonation and unwanted port scanning. The security policy of a standard or distributed switch is implemented in Layer 2 (Data Link Layer) of the network protocol stack. The three elements of the security policy are promiscuous mode, MAC address changes, and forged transmits.

See the *vSphere Networking* documentation for instructions.

Secure Virtual Machine Networking

The methods that you use to secure virtual machine networking depend on several factors, including:

- The guest operating system that is installed
- Whether the virtual machines operate in a trusted environment

Virtual switches and distributed virtual switches provide significant protection when used with other common security practices, such as installing firewalls.

See [Securing vSphere Networking](#).

Consider VLANs to Protect Your Environment

ESXi supports IEEE 802.1q VLANs. VLANs let you segment a physical network. You can use VLANs to further protect the virtual machine network or storage configuration. When you use VLANs, two virtual machines on the same physical network cannot send packets to or receive packets from each other unless they are on the same VLAN.

See [Securing Virtual Machines with VLANs](#).

Secure Connections to Virtualized Storage

A virtual machine stores operating system files, application files, and other data on a virtual disk. Each virtual disk appears to the virtual machine as a SCSI drive that is connected to a SCSI controller. A virtual machine is isolated from storage details and cannot access the information about the LUN where its virtual disk resides.

The Virtual Machine File System (VMFS) is a distributed file system and volume manager that presents virtual volumes to the ESXi host. You are responsible for securing the connection to storage. For example, if you are using iSCSI storage, you can set up your environment to use Challenge Handshake Authentication Protocol (CHAP). If required by company policy, you can set up mutual CHAP. Use the vSphere Client or CLIs to set up CHAP.

See [Storage Security Best Practices](#).

Evaluate the Use of Internet Protocol Security

ESXi supports Internet Protocol Security (IPSec) over IPv6. You cannot use IPSec over IPv4.

See [Using Internet Protocol Security on ESXi Hosts](#).

Securing Passwords in Your vSphere Environment

Password restrictions, password expiration, and account lockout in your vSphere environment depend on the system that the user targets, who the user is, and how policies are set.

ESXi password restrictions are determined by certain requirements. See [ESXi Passwords and Account Lockout](#).

vCenter Single Sign-On manages authentication for all users who log in to vCenter Server and other vCenter services. The password restrictions, password expiration, and account lockout depend on the domain of the user and on who the user is.

Password for the vCenter Single Sign-On Administrator

The password for the administrator@vsphere.local user, or the administrator@mydomain user if you selected a different domain during installation, does not expire and is not subject to the lockout policy. In all other regards, the password must follow the restrictions that are set in the vCenter Single Sign-On password policy. See the *vSphere Authentication* documentation for details.

If you forget the password for this user, search the VMware Knowledge Base system for information on resetting this password. The reset requires additional privileges such as root access to the vCenter Server system.

Passwords for Other Users of the vCenter Single Sign-On Domain

Passwords for other vsphere.local users, or users of the domain that you specified during installation, must follow the restrictions that are set by the vCenter Single Sign-On password policy and lockout policy. See the *vSphere Authentication* documentation for details. These passwords expire after 90 days by default. Administrators can change the expiration as part of the password policy.

If you forget your vsphere.local password, an administrator user can reset the password using the `dir-cli` command.

Passwords for Users from Other Identity Sources

Password restrictions, password expiration, and account lockout for all other users are determined by the domain (identity source) to which the user can authenticate.

vCenter Single Sign-On supports one default identity source. Users can log in to the corresponding domain with the vSphere Client with their user names. If users want to log in to a non-default domain, they can include the domain name, that is, specify `user@domain` or `domain\user`. The domain password parameters apply to each domain.

Passwords for vCenter Server Direct Console User Interface Users

The vCenter Server Appliance is a preconfigured virtual machine that is optimized for running vCenter Server and the associated services.

When you deploy vCenter Server, you specify these passwords.

- Password for the root user.
- Password for the administrator of the vCenter Single Sign-On domain, administrator@vsphere.local by default.

You can change the root user password and perform other vCenter Server local user management tasks from the vCenter Server Management Interface. See the *vCenter Server Configuration* documentation.

vCenter Server and ESXi Security Best Practices and Resources

If you follow best practices, your ESXi hosts and vCenter Server systems can be as secure as or even more secure than an environment that does not include virtualization.

This manual includes best practices for the different components of your vSphere infrastructure. This manual is only one of the sources you must use to ensure a secure environment.

vSphere Security Resources

To learn more about specific aspects of vSphere security, use the following content in this manual.

Table 230: Security Best Practices

vSphere component	Resource
ESXi host	Securing ESXi Hosts
vCenter Server system	Securing vCenter Server Systems
Virtual machine	Virtual Machine Security Best Practices
vSphere Networking	vSphere Networking Security Best Practices

VMware Security Resources on the Web

VMware security resources, including security alerts and downloads, are available on the Web.

Table 231: VMware Security Resources on the Web

Topic	Resource
Information about ESXi and vCenter Server security and operations, including secure configuration and hypervisor security.	https://www.vmware.com/resources/resource-center?product_family[]=Cloud%20Infrastructure&term[]=security
Corporate security response policy	http://www.vmware.com/support/policies/security_response.html VMware is committed to helping you maintain a secure environment. Security issues are corrected in a timely manner. The VMware Security Response Policy states our commitment to resolve possible vulnerabilities in our products.
Third-party software support policy	https://www.vmware.com/resources/compatibility/search.php VMware supports a variety of storage systems, software agents such as backup agents, system management agents, and so forth. You can find lists of agents, tools, and other software that support ESXi in the ESXi compatibility guides. The industry offers more products and configurations than VMware can test. If VMware does not list a product or configuration in a compatibility guide, Technical Support attempts to help you with any problems, but cannot guarantee that the product or configuration can be used. Always evaluate security risks for unsupported products or configurations carefully.
Compliance and security standards, and partner solutions and in-depth content about virtualization and compliance	https://github.com/vmware/vcf-security-and-compliance-guidelines
Security configuration guides (formerly known as hardening guides) for different versions of vSphere and other VMware products.	ESX Security Configuration Guide

vSphere Permissions and User Management Tasks

Authentication and authorization govern access to your vSphere environment. vCenter Single Sign-On supports authentication, which means it determines whether a user can log in to vSphere components at all. Each user must also be authorized to view or manipulate vSphere objects.

For an overview of assigning roles and permissions using the vSphere Client, watch the following video.

vCenter Server allows fine-grained control over authorization with permissions and roles. When you assign a permission to an object in the vCenter Server object hierarchy, you specify which user or group has which privileges on that object. To specify the privileges, you use roles, which are sets of privileges.

Initially, only the administrator user for the vCenter Single Sign-On domain is authorized to log in to the vCenter Server system. The default domain is `vsphere.local` and the default administrator is `administrator@vsphere.local`. You can change the default domain during installation of vSphere.

As an administrator user, you can:

1. Add an identity source in which users and groups are defined to vCenter Single Sign-On. See the *vSphere Authentication* documentation.
2. Give privileges to a user or group by selecting an object such as a virtual machine or vCenter Server system and assigning a role on that object for the user or group.

Understanding Authorization in vSphere


vSphere supports several models for determining whether a user is allowed to perform a task. Group membership in a vCenter Single Sign-On group decides what you are allowed to do. Your role on an object or your global permission determines whether you are allowed to perform other tasks.

How Do Permissions Work in vSphere

vSphere allows privileged users to give other users permissions to perform tasks. You can use global permissions, or you can use local vCenter Server permissions to authorize other users for individual vCenter Server instances.

The following figure illustrates how global and local permissions work.

Figure 110: Global Permissions and Local Permissions



In this figure:

1. You assign a global permission at the root object level with "Propagate to children" selected.
2. vCenter Server propagates the permissions to the vCenter Server 1 and vCenter Server 2 object hierarchies in the environment.

3. A local permission on the root folder on vCenter Server 2 overrides the global permission.


vCenter Server Permissions

The permission model for vCenter Server systems relies on assigning permissions to objects in the object hierarchy. Users get permissions in the following ways.

- From a specific permission for the user or from the groups that the user is a member of
- From a permission on the object or through the permission inheritance from a parent object

Each permission gives one user or group a set of privileges, that is, a role for a selected object. You can use the vSphere Client to add permissions. For example, you can right-click a virtual machine, select **Add Permission**, and complete the dialog box to assign a role to a group of users. That role gives those users the corresponding privileges on the virtual machine.

Figure 111: Adding Permissions to a Virtual Machine Using the vSphere Client



Global Permissions

Global permissions give a user or group privileges to view or manage all objects in each of the inventory hierarchies of the solutions in the deployment. That is, global permissions are applied to a global root object that spans solution inventory hierarchies. (Solutions include vCenter Server, VMware Aria Automation Orchestrator, and so on.) Global permissions also apply to global objects such as tags and content libraries. For example, consider a deployment that consists of two solutions, vCenter Server and VMware Aria Automation Orchestrator. You can use global permissions to assign a role to a group of users that has read-only privileges to all objects in both the vCenter Server and VMware Aria Automation Orchestrator object hierarchies.

Group Membership in vCenter Single Sign-On Groups

ESXi Local Host Permissions

If you are managing a standalone ESXi host that is not managed by vCenter Server system, you can assign one of the predefined roles to users. See the *vSphere Single Host Management - VMware Host Client* documentation.

For managed hosts, assign roles to the ESXi host object in the vCenter Server inventory.

Understanding the Object-Level Permission Model

You authorize a user or group to perform tasks on vCenter Server objects by using permissions on the object. From a programmatic standpoint, when a user tries to perform an operation, an API method is executed. vCenter Server checks the permissions for that method to see if the user is authorized to perform the operation. For example, when a user tries to add a host, the `AddStandaloneHost_Task` method is invoked. This method requires that the role for the user has the `Host.Inventory.AddStandaloneHost` privilege. If the check does not find this privilege, the user is denied permission to add the host.

The following concepts are important.

Permissions

Each object in the object hierarchy has associated permissions. Each permission specifies for one group or user which privileges that group or user has on the object. Permissions can propagate to child objects.

Users and Groups

On systems, you can assign privileges only to authenticated users or groups of authenticated users. Users are authenticated through. Users and groups must be defined in the identity source that uses to authenticate. Define users and groups using the tools in your identity source, for example, Active Directory.

Privileges


Privileges are fine-grained access controls. You can group those privileges into roles, which you can then map to users or groups.

Roles

Roles are sets of privileges. Roles allow you to assign permissions on an object based on a typical set of tasks that users perform. System roles, such as Administrator, are predefined on and cannot be changed. also provides some default sample roles, such as Resource Pool Administrator, that you can modify. You can create custom roles either from scratch or by cloning and modifying sample roles. See [Create a vCenter Server Custom Role](#)

The following figure illustrates how a permission is constructed from privileges and roles, and assigned to a user or group for a vSphere object.

Figure 112: vSphere Permissions



To assign permissions to an object, you follow these steps:

1. Select the object to which you want to apply the permission in the vCenter Server object hierarchy.
2. Select the group or user that should have privileges on the object.
3. Select individual privileges or a role, that is a set of privileges, that the group or user should have on the object.
By default, Propagate to children is not selected. You must select the checkbox for the group or user to have the selected role on the selected object and its child objects.

vCenter Server offers sample roles, which combine frequently used privilege sets. You can also create custom roles by combining a set of roles.

Permissions must often be defined on both a source object and a destination object. For example, if you move a virtual machine, you need privileges on that virtual machine, but also privileges on the destination data center.

See the following information.

To find out about...	See...
Creating custom roles.	Create a vCenter Server Custom Role
All privileges and the objects to which you can apply the privileges	Defined Privileges
Sets of privileges that are required on different objects for different tasks.	Required vCenter Server Privileges for Common Tasks

The permissions model for standalone ESXi hosts is simpler. See [Assigning Privileges for ESXi Hosts](#).

What Is vCenter Server User Validation

vCenter Server systems that use a directory service regularly validate users and groups against the user directory domain. Validation occurs at regular intervals specified in the vCenter Server settings. For example, assume that user Smith is assigned a role on several objects. The domain administrator changes the name to Smith2. The host concludes that Smith no longer exists and removes permissions associated with that user from the vSphere objects when the next validation occurs.

Similarly, if user Smith is removed from the domain, all permissions associated with that user are removed when the next validation occurs. If a new user Smith is added to the domain before the next validation occurs, the new user Smith replaces the old user Smith in permissions on any object.


Hierarchical Inheritance of Permissions in vSphere

When you assign a permission to an object, you can choose whether the permission propagates down the object hierarchy. You set propagation for each permission. Propagation is not universally applied. Permissions defined for a child object always override the permissions that are propagated from parent objects.

The following figure illustrates the inventory hierarchy and the paths by which permissions can propagate.

Note: Global permissions support assigning privileges across solutions from a global root object. See [Using vCenter Server Global Permissions](#).

Figure 113: vSphere Inventory Hierarchy



About this figure:

- You cannot set direct permissions on the VM, host, network, and storage folders. That is, these folders act as containers, and as such are not visible to users.
- You cannot set permissions on standard switches.

Note: To be able to set and propagate permissions to children on a vSphere Distributed Switch (VDS), the switch object must reside in a network folder created on the data center.

Most inventory objects inherit permissions from a single parent object in the hierarchy. For example, a datastore inherits permissions from either its parent datastore folder or parent data center. Virtual machines inherit permissions from both the parent virtual machine folder and the parent host, cluster, or resource pool simultaneously.

For example, you can set permissions for a distributed switch and its associated distributed port groups, by setting permissions on a parent object, such as a folder or data center. You must also select the option to propagate these permissions to child objects.

Permissions take several forms in the hierarchy.

Managed Entities

Managed entities refer to the following vSphere objects. Managed entities offer specific operations that vary depending on the entity type. Privileged users can define permissions on managed entities. See the vSphere API documentation for more information about vSphere objects, properties, and methods.

- Clusters
- Data centers
- Datastores
- Datastore clusters
- Folders
- Hosts
- Networks (except vSphere Distributed Switches)
- Distributed port groups
- Resource pools
- Templates
- Virtual machines
- vSphere vApps

Global Entities

You cannot modify permissions on entities that derive permissions from the root vCenter Server system.

- Custom fields
- Licenses
- Roles
- Statistics intervals
- Sessions

How Do Multiple Permission Settings Work in vSphere

Objects might have multiple permissions, but can have only one permission for each user or group. For example, one permission might specify that GroupAdmin has the Administrator role on an object. Another permission might specify that the GroupVMAdmin has the Virtual Machine Administrator role on the same object. However, the GroupVMAdmin group cannot have another permission for the same GroupVMAdmin on this object.

A child object inherits the permissions of its parent if the parent's propagate property is set to true. A permission that is set directly on a child object overrides the permission in the parent object. See [Example 2: Child Permissions Overriding Parent Permissions](#).

If multiple group roles are defined on the same object, and a user belongs to two or more of those groups, two situations are possible:

- No permission for the user is defined directly on the object. In that case, the user gets the union of the permissions that the groups have on the object.
- A permission for the user is defined directly on the object. In that case, the permissions for the user take precedence over all group permissions.

Example 1: Permission Inheritance from Multiple Groups


This example illustrates how an object can inherit multiple permissions from groups that are granted permission on a parent object.

In this example, two permissions are assigned on the same object for two different groups.

- PowerOnVMRole can power on virtual machines.
- SnapShotRole can take snapshots of virtual machines.
- PowerOnVMGroup is granted the PowerOnVMRole on VM Folder, with the permission set to propagate to child objects.
- SnapShotGroup is granted the SnapShotRole on VM Folder, with the permission set to propagate to child objects.
- User 1 is not assigned specific privileges.

User 1, who belongs to both the PowerOnVMGroup and the SnapShotGroup, logs in. User 1 can both power on and take snapshots of both VM A and VM B.

Figure 114: Example 1: Permission Inheritance from Multiple Groups



Example 2: Child Permissions Overriding Parent Permissions

This example illustrates how permissions that are assigned on a child object can override permissions that are assigned on a parent object. You can use this overriding behavior to restrict user access to particular areas of the inventory.


In this example, permissions are defined on two different objects for two different groups.

- PowerOnVMRole can power on virtual machines.

- SnapShotRole can take snapshots of virtual machines.
- PowerOnVMGroup is granted the PowerOnVMRole on VM Folder, with the permission set to propagate to child objects.
- SnapShotGroup is granted the SnapShotRole on VM B.

User 1, who belongs to both the PowerOnVMGroup and the SnapShotGroup, logs in. Because the SnapShotRole is assigned at a lower point in the hierarchy than the PowerOnVMRole, it overrides PowerOnVMRole on VM B. User 1 can power on VM A, but not take snapshots. User 1 can take snapshots of VM B, but not power it on.

Figure 115: Example 2: Child Permissions Overriding Parent Permissions




Example 3: User Role Overriding Group Role

This example illustrates how the role assigned directly to an individual user overrides the privileges associated with a role assigned to a group.

In this example, permissions are defined on the same object. One permission associates a group with a role, the other permission associates an individual user with a role. The user is a member of the group.

- PowerOnVMRole can power on virtual machines.
- PowerOnVMGroup is granted the PowerOnVMRole on VM Folder.
- User 1 is granted the NoAccess role on VM Folder.

User 1, who belongs to PowerOnVMGroup, logs in. The NoAccess role granted to User 1 on VM Folder overrides the role assigned to the group. User 1 has no access to VM Folder or VMs A and B. VMs A and B are not visible in the hierarchy to User 1.

Figure 116: Example 3: User Permissions Overriding Group Permissions

Managing Permissions for vCenter Server Components

A permission is set on an object in the vCenter Server object hierarchy. Each permission associates the object with a group or user and the group's or user's access role. For example, you can select a virtual machine object, add one permission that gives the **ReadOnly** role to Group 1, and add a second permission that gives the **Administrator** role to User 2.

By assigning a different role to a group of users on different objects, you control the tasks that those users can perform in your vSphere environment. For example, to allow a group to configure memory for the host, select that host and add a permission that grants a role to that group that includes the **Host > Configuration > Memory Configuration** privilege.

For conceptual information about permissions, see the discussion in [Understanding the Object-Level Permission Model](#).

You can assign permissions to objects at different levels of the hierarchy, for example, you can assign permissions to a host object or to a folder object that includes all host objects. See [Hierarchical Inheritance of Permissions in vSphere](#). You can also assign propagating permissions to a global root object to apply the permissions to all object in all solutions. See [Using vCenter Server Global Permissions](#).

Add a Permission to an Inventory Object

After you create users and groups and define roles, you must assign the users and groups and their roles to the relevant inventory objects. You can assign the same propagating permissions to multiple objects simultaneously by moving the objects into a folder and setting the permissions on the folder.

On the object whose permissions you want to modify, you must have a role that includes the **Permissions > Modify permission** privilege.

When you assign permissions, the user and the group names must match Active Directory precisely, including case. If you upgraded from earlier versions of vSphere, check for case inconsistencies if you experience problems with groups.

1. Browse to the object for which you want to assign permissions in the vSphere Client object navigator.
2. Click the **Permissions** tab.
3. Click **Add**.
4. Optional: If you have configured an external identity provider for federated authentication, the domain of that identity provider is available to select in the **Domain** drop-down menu.
5. If you select **VMware ID** from the **Domain** drop-down menu, enter the user or group name.

Note:

Enter the email address of the CSP account in the **Username** field. CSP accounts cannot be searched for in the VMwareID domain.

6. Select the user or group that will have the privileges defined by the selected role.
 - a) From the **Domain** drop-down menu, select the domain for the user or group.
 - b) Enter a name in the Search box.
The system searches user names and group names.
 - c) Select the user or group.
7. Select a role from the **Role** drop-down menu.
8. Optional: To propagate the permissions, select the **Propagate to children** check box.
The role is applied to the selected object and propagates to the child objects.
9. Click **OK**.

Change or Remove Permissions on an Inventory Object

After a user or group and role pair is set for an inventory object, you can change the role paired with the user or group or change the setting of the **Propagate to children** check box. You can also remove the permission setting.

1. Browse to the object in the vSphere Client object navigator.
2. Click the **Permissions** tab.
3. Click a row to select a permission.

Task	Steps
Change permissions	1. Click Edit. 2. Select a role for the user or group from the Role drop-down menu. 3. Toggle the Propagate to children check box to change permission inheritance. 4. Click OK.
Remove permissions	1. Click Delete. 2. Click Remove.

ChangevCenter Server User Validation Settings

vCenter Server periodically validates its user and group lists against the users and groups in the user directory. It then removes users or groups that no longer exist in the domain. You can deactivate validation or change the interval between validations. If you have domains with thousands of users or groups, or if searches take a long time to complete, consider adjusting the search settings.

These settings apply to vCenter Single Sign-On identity sources, and not an external identity source, such as Active Directory, that might be associated with vCenter Server.

Note: This procedure applies only to vCenter Server user lists. You cannot search ESXi user lists in the same way.

1. Browse to the vCenter Server system in the vSphere Client object navigator.
2. Select **Configure** and click **Settings > General**.
3. Click **Edit** and select **User directory**.
4. Change the values as needed and click **Save**.

Option	Description
User directory timeout	Timeout interval, in seconds, for searching this vCenter Server installation.
Query limit	Toggle on to set a maximum number of users and groups that vCenter Server displays.
Query limit size	Maximum number of users and groups from the selected domain that vCenter Server displays in the Select Users or Groups dialog box. If you enter 0 (zero), all users and groups appear.

Using vCenter Server Global Permissions

In vCenter Server, global permissions are applied to a global root object that spans VMware solutions. In an on-premises SDDC, global permissions might span both vCenter Server and VMware Aria Automation Orchestrator. But for any vSphere SDDC, global permissions apply to global objects such as tags and content libraries.

You can assign global permissions to users or groups, and decide on the role for each user or group. The role determines the set of privileges that the user or group has for all objects in the hierarchy. You can assign a predefined role or create custom roles. See [Using vCenter Server Roles to Assign Privileges](#).

It is important to distinguish between vCenter Server permissions and global permissions.

Table 232: Differences Between vCenter Server Permissions and Global Permissions

Permission Type	Description
vCenter Server	vCenter Server permissions apply to specific objects in the inventory hierarchy, such as hosts, virtual machines, datastores, and so on. When you assign vCenter Server permissions, you specify that a user or group has a role (set of privileges) on the object.
Global	Global permissions give a user or group privileges to view or manage all objects in each of the inventory hierarchies in your deployment. Global permissions also apply to global objects such as tags and content libraries. See vCenter Server Permissions on Tag Objects .

Permission Type	Description
	If you assign a global permission and do not select Propagate, the users or groups associated with this permission do not have access to the objects in the hierarchy. They only have access to some global functionality such as creating roles.

Add a Global Permission

You can use global permissions to give a user or group privileges for all objects in all inventory hierarchies in your deployment.

To perform this task, you must have **Permissions > Modify permission** privileges on the root object for all inventory hierarchies.

Important: Use global permissions with care. Verify that you really want to assign permissions to all objects in all inventory hierarchies.

1. Log in to the vCenter Server by using the vSphere Client.
2. Select **Administration** and click **Global Permissions** in the Access Control area.
3. Click **Add**.
4. Optional: If you have configured an external identity provider for federated authentication, the domain of that identity provider is available to select in the **Domain** drop-down menu.
5. For vSphere+ environments if you select **VMware ID** from the **Domain** drop-down menu, then enter the name of the CSP account in the **Username** field.

Note:

Enter the email address of the CSP account in the **Username** field. CSP accounts cannot be searched for in the VMwareID domain.

6. Select the user or group that will have the privileges defined by the selected role.
 - a) From the **Domain** drop-down menu, select the domain for the user or group.
 - b) Enter a name in the Search box.
The system searches user names and group names.
 - c) Select the user or group.
7. Select a role from the **Role** drop-down menu.
8. Decide whether to propagate the permissions by selecting the **Propagate to children** check box.
If you assign a global permission and do not select **Propagate to children**, the users or groups associated with this permission do not have access to the objects in the hierarchy. They only have access to some global functionality such as creating roles.
9. Click **OK**.

vCenter Server Permissions on Tag Objects

In the vCenter Server object hierarchy, tag objects are not children of vCenter Server but are created at the vCenter Server top level. In environments with multiple vCenter Server instances, tag objects are shared across vCenter Server instances. Permissions for tag objects work differently than permissions for other objects in the vCenter Server object hierarchy.

Only Global Permissions or Permissions Assigned to the Tag Object Apply

If you grant permissions to a user on a vCenter Server inventory object, such as a virtual machine, that user can perform the tasks associated with the permission. However, the user cannot perform tag operations on the object.

For example, if you grant the **Assign vSphere Tag** privilege to user Dana on host TPA, that permission does not affect whether Dana can assign tags on host TPA. Dana must have the **Assign vSphere Tag** privilege at the top level, that is, a global permission, or must have the privilege for the tag object.

Table 233: How Global Permissions and Tag Object Permissions Affect What Users Can Do

Global Permission	Tag-Level Permission	vCenter Server Object-Level Permission	Effective Permission
No tagging privileges assigned.	Dana has Assign or Unassign vSphere Tag privileges for the tag.	Dana has Delete vSphere Tag privileges on ESXi host TPA.	Dana has Assign or Unassign vSphere Tag privileges for the tag.
Dana has Assign or Unassign vSphere Tag privileges.	No privileges assigned for the tag.	Dana has Delete vSphere Tag privileges on ESXi host TPA.	Dana has Assign or Unassign vSphere Tag global privileges. That includes privileges at the tag level.
No tagging privileges assigned.	No privileges assigned for the tag.	Dana has Assign or Unassign vSphere Tag privileges on ESXi host TPA.	Dana does not have tagging privileges on any object, including host TPA.

Global Permissions Complement Tag Object Permissions

Global permissions, that is, permissions that are assigned on the top-level object, complement permissions on tag objects when the permissions on the tag objects are more restrictive. The vCenter Server permissions do not affect the tag objects.

For example, assume that you assign the **Delete vSphere Tag** privilege to user Robin at the top level by using global permissions. For the tag Production, you do not assign the **Delete vSphere Tag** privilege to Robin. In that case, Robin has the privilege for the tag Production because Robin has the global permission, which propagates from the top level. You cannot restrict privileges unless you modify the global permission.

Table 234: Global Permissions Complement Tag-Level Permissions

Global Permission	Tag-Level Permission	Effective Permission
Robin has Delete vSphere Tag privileges	Robin does not have Delete vSphere Tag privileges for the tag.	Robin has Delete vSphere Tag privileges.
No tagging privileges assigned	Robin does not have Delete vSphere Tag privileges assigned for the tag.	Robin does not have Delete vSphere Tag privileges

Tag-Level Permissions Can Extend Global Permissions

You can use tag-level permissions to extend global permissions. That means users can have both a global permission and a tag-level permission on a tag.

Note: This behavior is different from how vCenter Server privileges are inherited. In vCenter Server, permissions defined for a child object always override the permissions that are propagated from parent objects.

Table 235: Global Permissions Extend Tag-Level Permissions

Global Permission	Tag-Level Permission	Effective Permission
Lee has Assign or Unassign vSphere Tag privilege.	Lee has Delete vSphere Tag privilege.	Lee has the Assign vSphere Tag privilege and the Delete vSphere Tag privilege for the tag.
No tagging privileges assigned.	Lee has Delete vSphere Tag privilege assigned for the tag.	Lee has the Delete vSphere Tag privilege for the tag.

Using vCenter Server Roles to Assign Privileges

In vCenter Server, a role is a predefined set of privileges that defines rights to perform actions and read properties. You create permissions by assigning a role to a user or group for an object. vCenter Server provides system roles and sample roles by default. You can also create custom roles.

Assigning Permissions in vCenter Server

When you assign permissions in vCenter Server, you pair a user or group with a role, and associate that pairing with an inventory object. For example, you can use the Virtual machine user sample role to allow a user to read and change virtual machine attributes.

A single user or group can have different roles for different objects in the inventory. For example, assume that you have two resource pools in your inventory, Pool A and Pool B. You can assign group Sales the Virtual machine user sample role on Pool A, and the Read-only role on Pool B. With these assignments, the users in group Sales can turn on virtual machines in Pool A, but can only view virtual machines in Pool B.

Users can schedule tasks only if they have a role that includes privileges to perform that task at the time the task is created.

What Are the Predefined vCenter Server Roles

vCenter Server provides predefined roles, as shown in the following table.

Table 236: Predefined vCenter Server Roles

Role Type	Role Names	Description
System	Administrator, Read-only, and No access.	System roles are permanent. You cannot delete system roles nor can you edit the privileges associated with these roles. The system roles are organized as a hierarchy. Each role inherits the privileges of the previous role. For example, the Administrator role inherits the privileges of the Read-only role. See the following section for more details on system roles.
Sample	vSphere provides a number of sample roles, for example, AutoUpdateUser, Resource pool administrator, and Virtual machine user.	vSphere provides sample roles for certain frequently performed combination of tasks. You can clone, modify, or remove these roles.

Role Type	Role Names	Description
		Note: To avoid losing the predefined settings in a sample role, clone the role first and make modifications to the clone. You cannot reset the sample to its default settings.

To view the privileges associated with a role, navigate to the role in the vSphere Client (**Menu > Administration > Roles**) and click the **Privileges** tab.

To view all the vSphere privileges and descriptions, see [Defined Privileges](#).

Note: Changes to roles and privileges take effect immediately, even if the users involved are logged in. The exception is searches, where changes take effect after the user has logged out and logged back in.

vCenter Server System Roles

System roles cannot be altered or deleted.

Administrator Role

Users with the Administrator role for an object are allowed to view and perform all actions on the object. This role also includes all privileges of the Read Only role. If you have the Administrator role on an object, you can assign privileges to individual users and groups.

If you are acting in the Administrator role in vCenter Server, you can assign privileges to users and groups in the default vCenter Single Sign-On identity source. See the *vSphere Authentication* documentation for supported identity services.

By default, the administrator@vsphere.local user has the Administrator role on both vCenter Single Sign-On and vCenter Server after installation. That user can then associate other users with the Administrator role on vCenter Server.

Tip: Best practice is to create a user at the root level and assign the Administrator role to that user. After creating a named user with Administrator privileges, you can remove the root user from any permissions or change its role to No Access.

Read Only Role

Users with the Read Only role for an object are allowed to view the state of the object and details about the object. For example, users with this role can view virtual machine, host, and resource pool attributes, but cannot view the remote console for a host. All actions through the menus and toolbars are disallowed.

No Access Role

Users with the No Access role for an object cannot view or change the object in any way. New users and groups are assigned this role by default. You can change the role on an object-by-object basis.

The administrator of the vCenter Single Sign-On domain, administrator@vsphere.local by default, the root user, and vpxuser are assigned the Administrator role by default. Other users are assigned the No Access role by default.

Custom Roles in vCenter Server and ESXi

You can create custom roles for vCenter Server and all objects that it manages, or for individual hosts.

vCenter Server Custom Roles (Recommended)

ESXi Custom Roles

Create custom roles by using the role-editing facilities in the vSphere Client to create privilege sets that match your needs.

You can create custom roles for individual hosts by using a CLI or the VMware Host Client. See the *vSphere Single Host Management - VMware Host Client* documentation. Custom host roles are not accessible from vCenter Server.

If you manage ESXi hosts through vCenter Server, do not maintain custom roles in both the host and vCenter Server. Define roles at the vCenter Server level.

When you manage a host using vCenter Server, the permissions associated with that host are created through vCenter Server and stored on vCenter Server. If you connect directly to a host, only the roles that are created directly on the host are available.

Note: When you add a custom role and do not assign any privileges to it, the role is created as a Read Only role with three system-defined privileges: **System > Anonymous**, **System > View**, and **System > Read**. These privileges are not visible in the vSphere Client but are used to read certain properties of some managed objects. All the predefined roles in vCenter Server contain these three system-defined privileges. See the *vSphere Web Services API* documentation for more information.

Create a vCenter Server Custom Role

To suit the access control needs of your environment, you can create vCenter Server custom roles. You can create a role or clone an existing role.

Verify that you have Administrator privileges on the vCenter Server system where you create the role.

You can create or edit a role on a vCenter Server system that is part of the same vCenter Single Sign-On domain as other vCenter Server systems. The VMware Directory Service (vmdir) propagates the role changes that you make to all other vCenter Server systems in the group. Assignments of roles to specific users and objects are not shared across vCenter Server systems.

1. Log in to the vCenter Server by using the vSphere Client.
2. Select **Administration** and click **Roles** in the **Access Control** area.
3. Create the role.

Option	Description
To create a role	1. Click New. 2. Enter a name for the new role. 3. Select and deselect privileges for the role. Scroll the privilege categories and select all privileges or a subset of privileges for that category. You can show all, selected, or unselected categories. You can also show all, selected, or unselected privileges. See Defined Privileges for more information. 4. Click Create.
To create the role by cloning	1. Select a role, and click Clone. 2. Enter a name for the role. 3. Click OK. Note: When creating a cloned role, you cannot change privileges. To change privileges, select the cloned role and click Edit.

You can now create permissions by selecting an object and assigning the role to a user or group for that object.

Using Privilege Recorder

In vSphere, privileges are fine-grained access controls that can be grouped into roles and map them to users or groups. Privilege recorder helps you identify the minimum set of privileges required to run a vCenter Server workflow.

To run a specific set of operations, it is very difficult to determine the minimal set of privileges that are required by the user. The privileges do not have one-one mapping with the specific workflow which usually consists of multiple calls to different APIs operating on the respective object. As a result, the user either has more access or too little access to the environment. With the aim to keep the environment secure, the privilege recorder feature helps you identify the minimum set of privileges required to run a vCenter Server workflow. It allows you to monitor and query the privileges that were checked while performing an operation. Privilege recorder is implemented using a REST API.

Note: This feature is available as an API, and it supports only workflows run by a script. There is no UI support for Privilege Recorder.

Querying the ListAPI allows you to retrieve lists of privilege checks along with the corresponding sessions, users, managed objects, and operation IDs (opIDs). You can use the appropriate filters to obtain privileges for a particular workflow.

For example, assume that user A needs to create a VM. Creating a VM requires a certain set of privileges. User A must request for privileges from the system administrator. The system administrator can enable the privilege recorder and execute the create VM operation. While the privilege check is performed, the data for the privileges that were checked during the Create VM operation is stored. The data contains PrivilegeID, sessionID, OpID, and so on. In this example, this system admin will use the filters to get privileges for the create VM workflow. The system administrator can now create a role with minimum required privileges and assign it to the user.

Enable the Privilege Recorder

You enable the privilege recorder by using thevSphere Client to modify thevCenter Server configuration file, `vpxd.cfg`.

Verify that you have enough privileges to run your workflow. A user with the Administrator role is recommended.

1. Log in to thevCenter Server by using thevSphere Client.
2. Navigate to thevCenter Server instance.
3. Select **Configure > Advanced Settings**.
4. Click **Edit Settings**.
5. Add the settings.

Scroll down and in the **Name** field, enter the name of the setting, and in the **Value** field, enter the value for the specified setting.

Setting	Description
<code>config.vpxd.privCheck.bufferSize</code>	The count of privileges to be kept in memory. The default value is 0. If you do not change the default value, the privilege checks recorder does not record any data.
<code>config.vpxd.privCheck.cleanupInterval</code>	The interval on which privilege checks for unused sessions are cleaned up. The default value is 30 minutes.

6. Click **Add** and **Save**.

See [Performing Privilege Checks Operations](#) in *VMware vSphere Automation SDKs Programming Guide*.

Best Practices for vCenter Server Roles and Permissions

Follow best practices for roles and permissions to maximize the security and manageability of your vCenter Server environment.

Follow these best practices when configuring roles and permissions in your vCenter Server environment:

- Where possible, assign a role to a group rather than individual users.
- Grant permissions only on the objects where they are needed, and assign privileges only to users or groups that must have them. Use the minimum number of permissions to make it easier to understand and manage the structure of your permissions.
- If you assign a restrictive role to a group, check that the group does not contain the Administrator user or other users with administrative privileges. Otherwise, you might unintentionally restrict privileges of administrators in the parts of the inventory hierarchy where you have assigned that group the restrictive role.
- Group objects into folders to make assigning permissions easier. For example, to grant the modify permission on one set of hosts and the view permission on another set of hosts, place each set of hosts in a folder.
- Consider enabling propagation when you assign permissions to an object. Propagation ensures that new objects in the object hierarchy inherit permissions. For example, you can assign a permission to a virtual machine folder and enable propagation to ensure that the permission applies to all virtual machines in the folder.
- Use the No Access role to mask specific areas of the hierarchy. The No Access role restricts access for the users or groups with that role. However, in the case of VMs and vAPPs, there are two permission propagation chains. Assigning a propagating permission with No Access role on one of the chains, does not imply that the respective vApp or VM would have no privileges propagated to it.

Required vCenter Server Privileges for Common Tasks

Many tasks require permissions on multiple objects in the vSphere inventory. If the user who attempts to perform the task only has privileges on one object, the task cannot complete successfully.

The following table lists common tasks that require more than one privilege. You can add permissions to inventory objects by pairing a user with one of the predefined roles or with multiple privileges. If you expect that you must assign a set of privileges multiple times, create custom roles. To learn more about required privileges for common tasks, see [Privilege Recorder](#).

Refer to the *vSphere Web Services API Reference* documentation to learn how operations in the vSphere Client user interface map to API calls, and what privileges are required to perform operations. For example, the API documentation for the `AddHost_Task(addHost)` method specifies that the `Host.Inventory.AddHostToCluster` privilege is required to add a host to a cluster.

If the task that you want to perform is not in this table, the following rules explain where you must assign permissions to allow particular operations:

- Any operation that consumes storage space requires the **Datastore > Allocate Space** privilege on the target datastore, and the privilege to perform the operation itself. You must have these privileges, for example, when creating a virtual disk or taking a snapshot.
- Moving an object in the inventory hierarchy requires appropriate privileges on the object itself, the source parent object (such as a folder or cluster), and the destination parent object.
- Each host and cluster has its own implicit resource pool that contains all the resources of that host or cluster. Deploying a virtual machine directly to a host or cluster requires the **Resource > Assign Virtual Machine to Resource Pool** privilege.

Table 237: Required Privileges for Common Tasks

Task	Required Privileges	Applicable Role
Create a virtual machine	On the destination folder or data center: • Virtual machine > Edit inventory > Create new • Virtual machine > Change Configuration > Add new disk (if creating a new virtual disk)	Administrator

Task	Required Privileges	Applicable Role
	Virtual machine > Change Configuration > Add existing disk (if using an existing virtual disk) Virtual machine > Configuration > Configure Raw device (if using an RDM or SCSI pass-through device)	
	On the destination host, cluster, or resource pool: Resource > Assign virtual machine to resource pool	Resource pool administrator or Administrator
	On the destination datastore or the folder that contains the datastore: Datastore > Allocate space	Datastore Consumer or Administrator
	On the network that the virtual machine will be assigned to: Network > Assign network	Network Consumer or Administrator
Power on a virtual machine	On the data center in which the virtual machine is deployed: Virtual machine > Interaction > Power On	Virtual Machine Power User or Administrator
	On the virtual machine or folder of virtual machines: Virtual machine > Interaction > Power On	
Deploy a virtual machine from a template	On the destination folder or data center: Virtual machine > Edit inventory > Create from existing Virtual machine > Change Configuration > Add new disk	Administrator
	On a template or folder of templates: Virtual machine > Provisioning > Deploy template	Administrator
	On the destination host, cluster or resource pool: Resource > Assign virtual machine to resource pool vApp > Import	Administrator
	On the destination datastore or folder of datastores: Datastore > Allocate space	Datastore Consumer or Administrator
	On the network that the virtual machine will be assigned to: Network > Assign network	Network Consumer or Administrator
Take a virtual machine snapshot	On the virtual machine or a folder of virtual machines: Virtual machine > Snapshot management > Create snapshot	Virtual Machine Power User or Administrator
Move a virtual machine into a resource pool	On the virtual machine or folder of virtual machines: Resource > Assign virtual machine to resource pool Virtual machine > Edit inventory > Move	Administrator
	On the destination resource pool: Resource > Assign virtual machine to resource pool	Administrator
Install a guest operating system on a virtual machine	On the virtual machine or folder of virtual machines: Virtual machine > Interaction > Answer question Virtual machine > Interaction > Console interaction Virtual machine > Interaction > Device connection Virtual machine > Interaction > Power Off Virtual machine > Interaction > Power On Virtual machine > Interaction > Reset Virtual machine > Interaction > Configure CD media (if installing from a CD) Virtual machine > Interaction > Configure floppy media (if installing from a floppy disk)	Virtual Machine Power User or Administrator

Task	Required Privileges	Applicable Role
	Virtual machine > Interaction > VMware Tools install	
	On a datastore that contains the installation media ISO image: Datastore > Browse datastore (if installing from an ISO image on a datastore) On the datastore to which you upload the installation media ISO image: Datastore > Browse datastore Datastore > Low level file operations Host > Configuration > System Management	Virtual Machine Power User or Administrator
Migrate a virtual machine with vMotion	On the virtual machine or folder of virtual machines: Resource > Migrate powered on virtual machine Resource > Assign Virtual Machine to Resource Pool (if destination is a different resource pool from the source)	Resource Pool Administrator or Administrator
	On the destination host, cluster, or resource pool (if different from the source): Resource > Assign virtual machine to resource pool	Resource Pool Administrator or Administrator
Cold migrate (relocate) a virtual machine	On the virtual machine or folder of virtual machines: Resource > Migrate powered off virtual machine Resource > Assign virtual machine to resource pool (if destination is a different resource pool from the source)	Resource Pool Administrator or Administrator
	On the destination host, cluster, or resource pool (if different from the source): Resource > Assign virtual machine to resource pool	Resource Pool Administrator or Administrator
	On the destination datastore (if different from the source): Datastore > Allocate space	Datastore Consumer or Administrator
Migrate a virtual machine with Storage vMotion	On the virtual machine or folder of virtual machines: Resource > Migrate powered on virtual machine	Resource Pool Administrator or Administrator
	On the destination datastore: Datastore > Allocate space	Datastore Consumer or Administrator
On the destination cluster: Host > Inventory > Add host to cluster Host > Inventory > Modify cluster	Administrator	
Add a single host to a data center by using the vSphere Client, or add a single host to a cluster by using PowerCLI or API (leveraging the addHost API)	On the host: Host > Inventory > Add host to cluster	Administrator
	On the cluster: Host > Inventory > Modify cluster Host > Inventory > Add host to cluster	Administrator
	On the data center: Host > Inventory > Add standalone host	Administrator
Add multiple hosts to a cluster	On the cluster: Host > Inventory > Modify cluster Host > Inventory > Add host to cluster	Administrator
	On the parent data center of the cluster (with propagate): Host > Inventory > Add standalone host Host > Inventory > Move host	Administrator

Task	Required Privileges	Applicable Role
	- Host > Inventory > Modify cluster - Host > Configuration > Maintenance	
Encrypt a virtual machine		Administrator
Protect a virtual machine (if using vSphere+ to protect the virtual machine)	On the data center in which the virtual machine is deployed: vSphere Tagging > Assign or Unassign vSphere Tag	Administrator

Securing ESXi Hosts

The ESXi hypervisor architecture has many built-in security features such as CPU isolation, memory isolation, and device isolation. You can configure features such as lockdown mode, certificate replacement, and smart card authentication for enhanced security.

An ESXi host is also protected with a firewall. You can open ports for incoming and outgoing traffic as needed, but should restrict access to services and ports. Using the ESXi lockdown mode and limiting access to the ESXi Shell can further contribute to a more secure environment. ESXi hosts participate in the certificate infrastructure. Hosts are provisioned with certificates that are signed by the VMware Certificate Authority (VMCA) by default.

See the VMware white paper *Security of the VMware vSphere Hypervisor* for more information about ESXi security.

Note:

ESXi is not built upon the Linux kernel or a commodity Linux distribution. It uses its own VMware specialized and proprietary kernel and software tools, delivered as a self-contained unit, and does not contain applications and components from Linux distributions.

Starting in vSphere 8.0 Update 1, ESXi runs two reverse proxy services:

- VMware reverse proxy service, `rhttpproxy`
- Envoy

Envoy owns port 443, and all incoming ESXi requests are routed through Envoy. Starting in vSphere 8.0 Update 1, `rhttpproxy` serves as a configuration management server for Envoy.

General ESXi Security Recommendations

To secure an ESXi host against an unauthorized intrusion and misuse, VMware imposes constraints on several parameters, settings, and activities. To meet your configuration needs, you can loosen the constraints. If you do, make sure that you are working in a trusted environment and take other security measures.

What Are the ESXi Built-In Security Features

ESXi mitigates risks to your hosts as follows:

- The ESXi Shell interface and the SSH interface are deactivated by default. Keep these interfaces deactivated unless you are performing troubleshooting or support activities. For day-to-day activities, use the vSphere Client, where activity is subject to role-based access control and modern access control methods.
- Only some firewall ports are open by default. You can explicitly open firewall ports that are associated with specific services.
- By default, all ports that are not required for management access to the host are closed. Open ports if you need additional services.
- ESXi runs only services that are essential to managing its functions. The distribution is limited to the features required to run ESXi.

- By default, weak ciphers are deactivated and communications from clients are secured by SSL. The exact algorithms used for securing the channel depend on the SSL handshake. Default certificates created on ESXi use PKCS#1 SHA-256 with RSA encryption as the signature algorithm.
- An internal web service is used by ESXi to support access by Web clients. The service has been modified to run only functions that a Web client requires for administration and monitoring. As a result, ESXi is not vulnerable to web service security issues reported in broader use.
- VMware monitors all security alerts that can affect ESXi security and issues a security patch if needed. To view a list of security alerts, go to <https://blogs.vmware.com/security/2024/05/where-did-my-vmware-security-advisories-go.html>.
- Insecure services such as FTP and Telnet are not installed, and the ports for these services are closed by default.
- To protect hosts from loading drivers and applications that are not cryptographically signed, use UEFI Secure boot. Enabling Secure Boot is done at the system BIOS. No additional configuration changes are required on the ESXi host, for example, to disk partitions. See [UEFI Secure Boot for ESXi Hosts](#).
- If your ESXi host has a TPM 2.0 chip, enable and configure the chip in the system BIOS. Working together with Secure Boot, TPM 2.0 provides enhanced security and trust assurance rooted in hardware. See [Securing ESXi Hosts with Trusted Platform Module](#).
- In ESXi 8.0 and later, you can run the SSH process under a sandbox domain. The shell then has reduced privileges, and only permits access to a limited subset of commands. For more information, see the VMware knowledge base article at <https://kb.vmware.com/s/article/87386>.

Taking Further ESXi Security Measures

Consider the following recommendations when evaluating host security and administration.

Limit access to ESXi hosts

If you activate access to the Direct Console User Interface (DCUI), the ESXi Shell, or SSH, enforce strict access security policies.

The ESXi Shell has privileged access to certain parts of the host. Provide only trusted users with ESXi Shell login access.

Do not access managed ESXi hosts directly

Use the vSphere Client to administer ESXi hosts that are managed by a vCenter Server. Do not access managed hosts directly with the VMware Host Client, and do not change managed hosts from the DCUI.

If you manage hosts with a scripting interface or API, do not target the host directly. Instead, target the vCenter Server system that manages the host and specify the host name.

Use DCUI only for troubleshooting

Access the host from the DCUI or the ESXi Shell as the root user only for troubleshooting. To administer your ESXi hosts, use the vSphere Client (or the VMware Host Client), or one of the VMware CLIs or APIs. See *ESXCLI Concepts and Examples*. If you use the ESXi Shell or SSH, limit the accounts that have access and set timeouts.

Use only VMware sources to upgrade ESXi components

The host runs several third-party packages to support management interfaces or tasks that you must perform. VMware only supports upgrades to these packages that come from a VMware source. If you use a download or patch from another source, you might compromise management interface security or functions. Check third-party vendor sites and the VMware knowledge base for security alerts.

ESXi Advanced System Settings

Advanced system settings control aspects of ESXi behavior, such as logging, system resources, and security.

The following table presents some of the important ESXi advanced system settings for security. To view all the advanced system settings, consult either the vSphere Client (**Host > Configure > System > Advanced System Settings**) or the API for a given release.

Table 238: Partial List of Security Advanced System Settings

Advanced System Setting	Description	Default Value
Annotations.WelcomeMessage	Displays a welcome message in the Host Client prior to login, or in the DCUI on the default screen. In the DCUI, the welcome message replaces some text, such as the host IP address.	(Empty)
Config.Etc.issue	Displays a banner during an SSH login session.	(Empty)
Config.Etc.motd	Displays the message of the day upon SSH login. Note: To insert new lines or returns into the issues and motd configurations, you can use both the vSphere API and CLI. For example, see https://williamlam.com/2021/03/adding-a-customized-notification-banner-in-the-vsphere-ui.html and https://williamlam.com/2015/02/easily-manage-esxi-vcsa-ssh-login-banner-motd-in-vsphere-6-0.html .	(Empty)
Config.HostAgent.vmacore.soap.sessionTimeout	Sets the idle time in minutes before the system automatically logs out a VIM API. A value of 0 (zero) deactivates the idle time. This setting applies only to new sessions.	30 (minutes)
Mem.MemEagerZero	Activates zeroing the user world and the guest memory pages in the VMkernel operating systems (including the VMM process) after a virtual machine exit. The default value (0) uses lazy zeroing. A value of 1 uses eager zeroing.	0 (deactivated)
Security.AccountLockFailures	Sets the maximum number of failed login attempts before the system locks a user's account. For example, to lock the account on the fifth login failure, set this value to 4. A value of 0 (zero) deactivates account locking. For implementation reasons, some login mechanisms count unexpectedly: VIM logins (including the VMware Host Client) and ESXCLI reflect the exact number of failed logins.	5

Advanced System Setting	Description	Default Value
	- SSH connections count as a login attempt when displaying a password prompt, and undo that count on successful login. This behavior is normal for challenge and response communications. - CGI logins double-count login failures. CAUTION: Due to this problem, a user can be locked out faster than the number of failed logins when using the CGI interface.	
Security.AccountUnlockTime	Sets the number of seconds that a user is locked out. Any login attempt within the specified lock timeout restarts the lock timeout.	900 (15 minutes)
Security.PasswordHistory	Sets the number of passwords to remember for each user. This setting prevents duplicate or similar passwords.	5
Security.PasswordMaxDays	Sets the maximum number of days between password changes.	99999
Security.PasswordQualityControl	Changes the required length and the character class requirement or allow pass phrases in the <code>Pam_passwdqc</code> configuration. You can use special characters in passwords. You can have password lengths of at least 15 characters. The default setting requires three character classes and a minimum length of seven characters. If implementing the DoD Annex, you can combine the <code>similar=deny</code> option plus a minimum password length to enforce a requirement that passwords are sufficiently different. The password history setting is only enforced for passwords changed through the <code>VIM LocalAccountManager.changePassword</code> API. To change the password requires that the user have administrator permission. The <code>PasswordQualityControl</code> setting, with a <code>PasswordMaxDays</code> setting, satisfies the requirements of the DoD Annex: <pre>min=disabled,disabled,disabled,disabled,15 similar=deny</pre>	retry=3 min=disabled,disabled,disabled,7,7

Advanced System Setting	Description	Default Value
UserVars.DcuiTimeOut	Sets the idle time in seconds before the system automatically logs out the DCUI. A value of 0 (zero) deactivates the timeout.	600 (10 minutes)
UserVars.ESXiShellInteractiveTimeOut	Sets the idle time in seconds before the system automatically logs out an interactive shell. This setting takes effect for new sessions only. A value of 0 (zero) deactivates the idle time. Applies to both the DCUI and the SSH shell.	0
UserVars.ESXiShellTimeOut	Sets the time in seconds a login shell waits for login. A value of 0 (zero) deactivates the timeout. Applies to both the DCUI and the SSH shell.	0
UserVars.HostClientSessionTimeout	Sets the idle time in seconds before the system automatically logs out the Host Client. A value of 0 (zero) deactivates the idle time.	900 (15 minutes)
UserVars.HostClientWelcomeMessage	Displays a welcome message in the Host Client upon login. The message is displayed following login as a "hint".	(Empty)

Configure ESXi Hosts with Host Profiles

Host profiles allow you to set up standard configurations for your ESXi hosts and automate compliance to these configuration settings. Host profiles allow you to control many aspects of host configuration including memory, storage, networking, and so on.

Host Profiles provide an automated and centrally managed mechanism for host configuration and configuration compliance. Host Profiles can improve efficiency by reducing reliance upon repetitive, manual tasks. Host Profiles capture the configuration of a pre-configured and validated reference host, store the configuration as a managed object and use the catalog of parameters contained within to configure networking, storage, security, and other host-level parameters.

You can configure host profiles for a reference host from the vSphere Client and apply the host profile to all hosts that share the characteristics of the reference host. You can also use host profiles to monitor hosts for host configuration changes. See the *vSphere Host Profiles* documentation.

You can attach the host profile to a cluster to apply it to all hosts in the cluster.

1. Set up the reference host to specification and create a host profile.
2. Attach the profile to a host or cluster.
3. Apply the host profile of the reference host to other hosts or clusters.

Use Scripts to Manage ESXi Host Configuration Settings

In environments with many ESXi hosts, managing hosts with scripts is faster and less error prone than managing the hosts from the vSphere Client.

vSphere includes scripting languages for ESXi host management. VMware PowerCLI is a Windows PowerShell interface to the vSphere API, and includes PowerShell cmdlets for administering vSphere components. ESXCLI includes a set of commands for managing ESXi hosts and virtual machines. See <https://developer.vmware.com> for reference information and programming tips. The vSphere Administrator documentation focuses on using the vSphere Client for management.

You can also use one of the scripting interfaces to the vSphere Automation SDK such as the vSphere Automation SDK for Python.

1. Create a custom role that has limited privileges.

See [Create a vCenter Server Custom Role](#).

For example, consider creating a role that has a set of privileges for managing hosts but no privileges for managing virtual machines, storage, or networking. If the script you want to use only extracts information, you can create a role with read-only privileges for the host.

2. From the vSphere Client, create a service account and assign it the custom role.

You can create multiple custom roles with different levels of access if you want access to certain hosts to be fairly limited.

3. Write scripts to perform parameter checking or modification, and run them.

For example, you can check or set the shell interactive timeout of a host as follows:

Language	Commands
ESXCLI	<pre>esxcli <conn_options> system settings advanced get /UserVars/ESXiShellTimeout esxcli --formatter=csv --format-param=field- s="Path,Int Value" system settings advanced list grep /UserVars/ESX- iShellTimeout</pre>
PowerCLI	<pre>#List UserVars.ESXiShellInteractiveTimeout for each host Get-VMHost Select Name, @{N="UserVars.ESXiShellInteractiveTimeout";E={\$_ Get-AdvancedSetting -Name UserVars.ESXiShellIn- teractiveTimeout Select -ExpandProperty Value}} # Set UserVars.ESXiShellTimeout to 900 on all hosts Get-VMHost Foreach { Get-AdvancedSetting -Entity \$_ -Name UserVars.ESXiShellInteractiveTimeout Set-Ad- vancedSetting -Value 900 }</pre>

4. In large environments, create roles with different access privileges and group hosts into folders according to the tasks that you want to perform. You can then run scripts over different folders from different service accounts.
5. Verify that the changes happened after you run the command.

ESXi Passwords and Account Lockout

For ESXi hosts, you must use a password with predefined requirements. You can change the required length and the character class requirement or allow pass phrases using the `Security.PasswordQualityControl` advanced system setting. You can also set the number of passwords to remember for each user using the `Security.PasswordHistory` advanced system setting.

Note: The default requirements for ESXi passwords can change from one release to the next. You can check and change the default password restrictions by using the `Security.PasswordQualityControl` advanced system setting.

ESXi Passwords

ESXi enforces password requirements for access from the Direct Console User Interface, the ESXi Shell, SSH, or the VMware Host Client.

- By default, you must include a mix of at least three from the following four character classes: lowercase letters, uppercase letters, numbers, and special characters such as underscore or dash when you create a password.
- By default, password length is at least 7 characters and less than 40.
- Passwords must not contain a dictionary word or part of a dictionary word.
- Passwords must not contain the user name or parts of the user name.

Note: An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used. A dictionary word used inside a password reduces the overall password strength.

Example ESXi Passwords

The following password candidates illustrate potential passwords if the option is set as follows.

```
retry=3 min=disabled,disabled,disabled,7,7
```

With this setting, a user is prompted up to three times (`retry=3`) for a new password that is not sufficiently strong or if the password was not entered correctly twice. Passwords with one or two character classes and pass phrases are not allowed, because the first three items are deactivated. Passwords from three- and four-character classes require seven characters. See the `pam_passwdqc` man page for details on other options, such as `max`, `passphrase`, and so on.

With these settings, the following passwords are allowed.

- `xQaTEhb!`: Contains eight characters from three character classes.
- `xQaT3#A`: Contains seven characters from four character classes.

The following password candidates do not meet requirements.

- `Xqat3hi`: Begins with an uppercase character, reducing the effective number of character classes to two. The minimum number of required character classes is three.
- `xQaTEh2`: Ends with a number, reducing the effective number of character classes to two. The minimum number of required character classes is three.

ESXi Pass Phrase

Instead of a password, you can also use a pass phrase. However, pass phrases are deactivated by default. You can change the default setting and other settings by using the `Security.PasswordQualityControl` advanced system setting from the vSphere Client.

For example, you can change the option to the following.

```
retry=3 min=disabled,disabled,16,7,7
```

This example allows pass phrases of at least 16 characters and at least three words.

Changing Default Password Restrictions

You can change the default restriction on passwords or pass phrases by using the `Security.PasswordQualityControl` advanced system setting for your ESXi host. See the *vCenter Server and Host Management* documentation for information on changing ESXi advanced system settings.

You can change the default, for example, to require a minimum of 15 characters and a minimum number of four words (`passphrase=4`), as follows:

```
retry=3 min=disabled,disabled,15,7,7 passphrase=4
```

See the man page for `pam_passwdqc` for details.

Note: Not all possible combinations of password options have been tested. Perform testing after you change the default password settings.

This example sets the password complexity requirement to require eight characters from four character classes that enforce a significant password difference, a remembered history of five passwords, and a 90 day rotation policy:

```
min=disabled,disabled,disabled,disabled,8 similar=deny
```

ESXi Account Lockout Behavior

Account locking is supported for access through SSH and through the vSphere Web Services SDK. The Direct Console Interface (DCUI) and the ESXi Shell do not support account lockout. By default, a maximum of five failed attempts is allowed before the account is locked. The account is unlocked after 15 minutes by default.

Configuring Login Behavior

You can configure the login behavior for your ESXi host with the following advanced system settings:

- `Security.AccountLockFailures`. Maximum number of failed login attempts before a user's account is locked. Zero deactivates account locking.
- `Security.AccountUnlockTime`. Number of seconds that a user is locked out.
- `Security.PasswordHistory`. Number of passwords to remember for each user. Starting in vSphere 8.0 Update 1, the default is five. Zero deactivates password history.

See the *vCenter Server and Host Management* documentation for information on setting ESXi advanced options.

ESXi Cryptographic Key Generation

ESXi generates several asymmetric keys for normal operation. The Transport Layer Security (TLS) key secures communication with the ESXi host using the TLS protocol. The SSH key secures communication with the ESXi host using the SSH protocol.

Transport Layer Security Key

The Transport Layer Security (TLS) key secures communication with the host using the TLS protocol. Upon first boot, the ESXi host generates the TLS key as a 2048-bit RSA key. Currently, ESXi does not implement automatic generation of ECDSA keys for TLS. The TLS private key is not intended to be serviced by the administrator.

The TLS key resides at the following non-persistent location:

```
/etc/vmware/ssl/rui.key
```

The TLS public key (including intermediate certificate authorities) resides at the following non-persistent location as an X.509 v3 certificate :

```
/etc/vmware/ssl/rui.crt
```

When you use vCenter Server with your ESXi hosts, vCenter Server generates a CSR automatically, signs it using the VMware Certificate Authority (VMCA), and generates the certificate. When you add an ESXi host to vCenter Server, vCenter Server installs that resulting certificate on the ESXi host.

The default TLS certificate is self-signed, with a subjectAltName field matching the host name at installation. You can install a different certificate, for example, to make use of a different subjectAltName or to include a particular Certificate Authority (CA) in the verification chain. See [Replacing the Default ESXi Certificate with a Custom Certificate](#).

SSH Key

The SSH key secures communication with the ESXi host using the SSH protocol. Upon first boot, the system generates a nistp256 ECDSA key, and the SSH keys as 2048-bit RSA keys. The SSH server is deactivated by default. SSH access is intended primarily for troubleshooting purposes. The SSH keys are not intended to be serviced by the administrator. Logging in through SSH requires administrative privileges equivalent to full host control. To enable SSH access, see [Activate Access to the ESXi Shell Using the vSphere Client](#).

The SSH public keys reside at the following location:

```
/etc/ssh/ssh_host_rsa_key.pub
```

```
/etc/ssh/ssh_host_ecdsa_key.pub
```

The SSH private keys reside at the following location:

```
/etc/ssh/ssh_host_rsa_key
```

```
/etc/ssh/ssh_host_ecdsa_key
```

TLS Cryptographic Key Establishment

Configuration of TLS cryptographic key establishment is governed by choice of TLS cipher suites, which selects ECC-based key agreements using ephemeral Ecliptic Curve Diffie Hellman (ECDH) (as specified in NIST Special Publication 800-56A).

SSH Cryptographic Key Establishment

Configuration of SSH cryptographic key establishment is governed by the SSHD configuration. ESXi provides a default configuration that permits ephemeral Diffie Hellman (DH) (as specified in NIST Special Publication 800-56A) key agreement, and ephemeral Ecliptic Curve Diffie Hellman (ECHD) (as specified in NIST Special Publication 800-56A). The SSHD configuration is not intended to be serviced by the administrator.

SSH Security in ESXi

The ESXi Shell interface and the SSH interface are deactivated by default. Keep these interfaces deactivated unless you are performing troubleshooting or support activities. For regular activities, use the vSphere Client, where activity is subject to role-based access control and modern access control methods.

SSH Configuration in ESXi

The SSH configuration in ESXi uses the following settings.

Version 1 SSH protocol deactivated

VMware does not support Version 1 SSH protocol and uses Version 2 protocol exclusively. Version 2 eliminates certain security

Improved cipher strength

problems present in Version 1 and provides you with a safe way to communicate with the management interface.

SSH supports only 256-bit and 128-bit AES ciphers for your connections.

These settings are designed to provide solid protection for the data you transmit to the management interface through SSH. You cannot change these settings.

ESXi SSH Keys

SSH keys can restrict, control, and secure access to an ESXi host. An SSH key can allow a trusted user or script to log in to a host without entering a password.

You can use HTTPS PUT to copy the SSH key to the host.

Instead of generating the keys externally and uploading them, you can create the keys on the ESXi host and download them. See the VMware knowledge base article at <https://kb.vmware.com/s/article/1002866>.

Enabling SSH and adding SSH keys to the host has inherent risks. Weigh the potential risk of exposing a user name and password against the risk of intrusion by a user who has a trusted key.

Upload an SSH Key Using HTTPS PUT

You can use authorized keys to log in to a host with SSH. You can upload authorized keys with HTTPS PUT.

Authorized keys allow you to authenticate remote access to a host. When users or scripts try to access a host with SSH, the key provides authentication without a password. With authorized keys you can automate authentication, which is useful when you write scripts to perform routine tasks.

You can upload the following types of SSH keys to a host using HTTPS PUT:

- Authorized keys file for root user
- DSA key
- DSA public key
- RSA key
- RSA public key

Important: Do not modify the `/etc/ssh/sshd_config` file.

1. In your upload application, open the key file.
2. Publish the file to the following locations.

Type of key	Location
Authorized key files for the root user	<code>https://hostname_or_IP_address/host/ssh_root_authorized_keys</code> You must have full administrator privileges on the host to upload this file.
DSA keys	<code>https://hostname_or_IP_address/host/ssh_host_dsa_key</code>
DSA public keys	<code>https://hostname_or_IP_address/host/ssh_host_dsa_key_pub</code>
RSA keys	<code>https://hostname_or_IP_address/host/ssh_host_rsa_key</code>
RSA public keys	<code>https://hostname_or_IP_address/host/ssh_host_rsa_key_pub</code>

PCI and PCIe Devices and ESXi

Using the VMware DirectPath I/O feature to pass through a PCI or PCIe device to a virtual machine results in a potential security vulnerability. The vulnerability can be triggered when buggy or malicious code, such as a device driver, is running in privileged mode in the guest OS. Industry-standard hardware and firmware do not currently have sufficient error containment support to protect ESXi hosts from the vulnerability.

Use PCI or PCIe passthrough to a virtual machine only if a trusted entity owns and administers the virtual machine. You must be sure that this entity does not attempt to crash or exploit the host from the virtual machine.

Your host might be compromised in one of the following ways.

- The guest OS might generate an unrecoverable PCI or PCIe error. Such an error does not corrupt data, but can crash the ESXi host. Such errors might occur because of bugs or incompatibilities in the hardware devices that are being passed through. Other reasons for errors include problems with drivers in the guest OS.
- The guest OS might generate a Direct Memory Access (DMA) operation that causes an IOMMU page fault on the ESXi host. This operation might be the result of a DMA operation that targets an address outside the virtual machine memory. On some machines, host firmware configures IOMMU faults to report a fatal error through a non-maskable interrupt (NMI). This fatal error causes the ESXi host to crash. This problem might occur because of problems with the drivers in the guest OS.
- If the operating system on the ESXi host is not using interrupt remapping, the guest OS might inject a spurious interrupt into the ESXi host on any vector. ESXi currently uses interrupt remapping on Intel platforms where it is available. Interrupt mapping is part of the Intel VT-d feature set. ESXi does not use interrupt mapping on AMD platforms. A false interrupt can result in a crash of the ESXi host. Other ways to exploit these false interrupts might exist in theory.

Deactivate the vSphere Managed Object Browser

The managed object browser (MOB) is a vSphere utility that provides a way to explore the VMkernel object model. However, attackers can use this interface to perform malicious configuration changes or actions because it is possible to change the host configuration by using the MOB. Use the MOB only for debugging, and ensure that it is deactivated in production systems.

The MOB is deactivated by default. However, for certain tasks, for example when extracting the old certificate from a system, you have to use the MOB. You can activate and deactivate the MOB as follows.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, click **Advanced System Settings**.
4. Check the value of **Config.HostAgent.plugins.solo.enableMob**, and click **Edit** to change it as appropriate.

Do not use `vim-cmd` from the ESXi Shell.

ESXi Networking Security Recommendations

Isolation of network traffic is essential to a secure ESXi environment. Different networks require a different access and level of isolation.

Your ESXi host uses several networks. Use appropriate security measures for each network, and isolate traffic for specific applications and functions. For example, ensure that VMware vSphere® vMotion® traffic does not travel over networks where virtual machines are located. Isolation prevents snooping. Having separate networks is also recommended for performance reasons.

- vSphere infrastructure networks are used for features such as vSphere vMotion, VMware vSphere Fault Tolerance, VMware vSAN, and storage. Isolate these networks for their specific functions. It is often not necessary to route these networks outside a single physical server rack.
- A management network isolates client traffic, command-line interface (CLI) or API traffic, and third-party software traffic from other traffic. In general, the management network is accessible only by system, network, and security administrators. To secure access to the management network, use a bastion host or a virtual private network (VPN). Strictly control access within this network.
- Virtual machine traffic can flow over one or many networks. You can enhance the isolation of virtual machines by using virtual firewall solutions that set firewall rules at the virtual network controller. These settings travel with a virtual machine as it migrates from host to host within your vSphere environment.

Modifying ESXi Web Proxy Settings

When you modify Web proxy settings, you have several encryption and user security guidelines to consider.

Note: Restart the host process after making any changes to host directories or authentication mechanisms.

- Do not set up certificates that use a password or pass phrases. ESXi does not support Web proxies that use passwords or pass phrases, also known as encrypted keys. If you set up a Web proxy that requires a password or pass phrase, ESXi processes cannot start correctly.
- To support encryption for user names, passwords, and packets, SSL is activated by default for vSphere Web Services SDK connections. If you want to configure these connections so that they do not encrypt transmissions, deactivate SSL for your vSphere Web Services SDK connection by switching the connection from HTTPS to HTTP.

Consider deactivating SSL only if you created a fully trusted environment for these clients, where firewalls are in place and transmissions to and from the host are fully isolated. Deactivating SSL can improve performance, because you avoid the overhead required to perform encryption.

- To protect against misuse of ESXi services, most internal ESXi services are accessible only through port 443, the port used for HTTPS transmission. Port 443 acts as a reverse proxy for ESXi. You can see a list of services on ESXi through an HTTP welcome page, but you cannot directly access the Storage Adapters services without proper authorization.

You can change this configuration so that individual services are directly accessible through HTTP connections. Do not make this change unless you are using ESXi in a fully trusted environment.

- When you upgrade your environment, the certificate remains in place.

vSphere Auto Deploy Security Considerations

When you use vSphere Auto Deploy, pay careful attention to networking security, boot image security, and potential password exposure through host profiles to protect your environment.

Networking Security

Secure your network just as you secure the network for any other PXE-based deployment method. vSphere Auto Deploy transfers data over SSL to prevent casual interference and snooping. However, the authenticity of the client or of the Auto Deploy server is not checked during a PXE boot.

You can greatly reduce the security risk of Auto Deploy by completely isolating the network where Auto Deploy is used.

Boot Image and Host Profile Security

The boot image that the vSphere Auto Deploy server downloads to a machine can have the following components.

- The VIB packages that the image profile consists of are always included in the boot image.

- The host profile and host customization are included in the boot image if Auto Deploy rules are set up to provision the host with a host profile or host customization.
 - The administrator (root) password and user passwords that are included with host profile and host customization are hashed with SHA-512.
 - Any other passwords associated with profiles are in the clear. If you set up Active Directory by using host profiles, the passwords are not protected.
Use the vSphere Authentication Proxy to avoid exposing the Active Directory passwords. If you set up Active Directory using host profiles, the passwords are not protected.
- The host's public and private SSL key and certificate are included in the boot image.

Control Access for CIM-Based Hardware Monitoring Tools

The Common Information Model (CIM) system provides an interface for remote applications to monitor hardware resources using a set of standard APIs. To ensure that the CIM interface is secure, provide only the minimum access necessary to these remote applications. If you provision a remote application with a root or Administrator account, and if the application is compromised, the virtual environment can be compromised.

CIM is an open standard that defines a framework for agent-less, standards-based monitoring of hardware resources for ESXi hosts. This framework consists of a CIM object manager, often called a CIM broker, and a set of CIM providers.

CIM providers support management access to device drivers and underlying hardware. Hardware vendors, including server manufacturers and hardware device vendors, can write providers that monitor and manage their devices. VMware writes providers that monitor server hardware, ESXi storage infrastructure, and virtualization-specific resources. These providers run inside the ESXi host and are lightweight and focused on specific management tasks. The CIM broker takes information from all CIM providers and presents it to the outside world using standard APIs. The most common API is WS-MAN.

Do not provide root credentials to remote applications that access the CIM interface. Instead, create a less-privileged vSphere user account for these applications and use the VIM API ticket function to issue a sessionId (called a "ticket") to this less-privileged user account to authenticate to CIM. If the account has been granted permission to obtain CIM tickets, the VIM API can then supply the ticket to CIM. These tickets are then supplied as both the user ID and password to any CIM-XML API call. See the `AcquireCimServicesTicket()` method for more information.

The CIM service starts when you install a third-party CIM VIB, for example, when you run the `esxcli software vib install -n VIBname` command.

If you must activate the CIM service manually, run the following command:

```
esxcli system wbem set -e true
```

If necessary, you can deactivate wsman (WSManagement Service) so that only the CIM service is running:

```
esxcli system wbem set -W false
```

To confirm that wsman is deactivated, run the following command:

```
esxcli system wbem get
...
WSManagement PID: 0
WSManagement Service: false
```

For more information about ESXCLI commands, see *ESXCLI Documentation*. For more information about activating the CIM service, see the VMware knowledge base article at <https://kb.vmware.com/s/article/1025757>.

1. Create a non-root vSphere user account for CIM applications.

See the topic on adding vCenter Single Sign-On users in *vSphere Authentication*. The required vSphere privilege for the user account is **Host.CIM > Interaction**.

2. Use the vSphere API SDK of your choice to authenticate the user account to vCenter Server. Then call `AcquireCimServicesTicket()` to return a ticket to authenticate with ESXi as an administrator-level account using CIM-XML port 5989 or WS-Man port 433 APIs.
See *vSphere Web Services API Reference* for more information.
3. Renew the ticket every two minutes as needed.

vSphere Distributed Services Engine Security Best Practices

To maximize the security of your ESXi environment, follow best practices for vSphere Distributed Services Engine.

In vSphere 8.0 and later, vSphere Distributed Services Engine enables the offloading of infrastructure functions from the CPUs of a host or a server to data processing units (DPUs, also known as SmartNICs), thus freeing up CPU cycles to serve applications. For an introduction to vSphere Distributed Services Engine, see the *VMware ESXi Installation and Setup* documentation. For more information about vSphere Distributed Services Engine, see the *Managing Host and Cluster Lifecycle* documentation.

In general, treat security aspects of vSphere Distributed Services Engine as you do when securing your ESXi environment.

- The ESXi Shell interface and the SSH interface to vSphere Distributed Services Engine are deactivated by default. Keep these interfaces deactivated unless you are performing troubleshooting or support activities.
- For day-to-day management activities of vSphere Distributed Services Engine, use the vSphere Client, where activity is subject to role-based access control and modern access control methods.

Controlling ESXi Entropy

In ESXi 8.0 and later, the ESXi Entropy implementation supports the FIPS 140-3 and EAL4 certifications. Kernel boot options control which entropy sources to activate on an ESXi host.

You must have root access on the ESXi host.

In computing, the term “entropy” refers to random characters and data that are collected for use in cryptography, such as generating encryption keys to secure data transmitted over a network. Entropy is required by security for generating keys and communicating securely over the network. Entropy is often collected from a variety of sources on a system.

FIPS entropy handling is the default behavior if the following conditions are true.

1. The hardware supports RDSEED.
2. The `disableHwrng` VMkernel boot option isn't present or is `FALSE`.
3. The `entropySources` VMkernel boot option isn't present, is 0 (zero), or is 4.

Warning: When you configure an ESXi host with `entropySources` for external entropy only (that is, `entropySources` is set to 8), you must keep supplying the external entropy to the host using the entropy API. If the entropy becomes exhausted in the host, then the host becomes unresponsive. To recover from this situation, reboot the host. If the host is still unresponsive, you must reinstall ESXi.

Starting with ESXi 8.0 Update 1, you can configure external entropy sources in the kickstart file for scripted installation. You can configure ESXi in a highly secure environment to consume entropy from external entropy sources, such as a Hardware Security Module (HSM), and align with standards such as BSI Common criteria, EAL4, and NIST FIPS CMVP, by using the scripted installation method. For more information about configuring external entropy sources, see the *VMware ESXi Installation and Setup* documentation.

You can configure the ESXi Entropy subsystem using the following VMkernel boot options:

Table 239: ESXi Entropy VMkernel Boot Options

VMkernel Boot Option	Option Type	Description	Default Value
disableHwrng (available prior to vSphere 8.0)	Boolean	Deactivates the RDRAND and the RDSEED entropy sources when set to TRUE (overrides "entropySources").	FALSE Activates hardware random number generator entropy sources if present.
entropySources (available starting in vSphere 8.0)	Integer, Bitmask	Specifies which entropy sources to activate. 0 (default) Bitmask values: 1=interrupts 2=RDRAND 4=RDSEED 8=entropyd (EAL4 entropy handling is activated) Specifying entropySources=9 activates the interrupts and the userspace entropy sources, and deactivates the RDRAND and the RDSEED entropy sources.	0 (zero) If RDSEED is supported, the default is FIPS compliance. Otherwise the default is all entropy sources except entropyd.

Note: Before making a change to use only RDRAND, RDSEED, or both entropy sources, check your vendor documentation to ensure that your ESXi host supports those configurations. If your host does not support those configurations, vCenter Server notifies you with an alert, and the host falls back to using the interrupt and userspace entropy sources.

1. Use SSH or another remote console connection to start a session on the ESXi host.
2. Log in as root.
3. Set the desired entropy VMkernel boot options.
 - a) To deactivate the RDRAND and the RDSEED entropy sources for disableHwrng:

```
esxcli system settings kernel set -s disableHwrng -v TRUE
```

- b) To set entropySources:

```
esxcli system settings kernel set -s entropySources -v entropy_source_value
```

See the preceding table for the values that you can set for entropySources.

Managing Certificates for ESXi Hosts

The VMware Certificate Authority (VMCA) provisions each new ESXi host with a signed certificate that has VMCA as the root certificate authority by default. Provisioning happens when you add a host to vCenter Server explicitly or as part of installation or upgrade of ESXi.

You can view and manage ESXi certificates from the vSphere Client and by using the `vim.CertificateManager` API in the vSphere Web Services SDK. You cannot view or manage ESXi certificates by using certificate management CLIs that are available for managing vCenter Server certificates.

Starting in vSphere 8.0 Update 3, you can replace ESXi certificates without placing the host into maintenance mode, and without having to restart the host or individual services.

Certificates and Certificate Modes

When ESXi and vCenter Server communicate, they use TLS for almost all management traffic.

vCenter Server supports the following certificates and certificate modes for ESXi hosts.

Table 240: Certificate Modes for ESXi Hosts

Certificate Mode	Description
VMware Certificate Authority (default)	By default, the VMware Certificate Authority is used as the certificate authority (CA) for ESXi host certificates. VMCA is the root CA by default, but it can be set up as the intermediary CA to another CA. In vmca mode, you can renew and refresh certificates from the vSphere Client. Also used if VMCA is a subordinate certificate.
Custom Certificate Authority	Use this mode if you want to use only custom certificates that are signed by a third-party or enterprise CA. In custom mode, you are responsible for managing the certificates. Starting in vSphere 8.0 Update 3, you can manage custom certificates from the vSphere Client. Note: Unless you change the certificate mode to Custom Certificate Authority (custom), VMCA might replace custom certificates, for example, when you select Renew in the vSphere Client.
Thumbprint Mode	vSphere 5.5 used thumbprint mode, and this mode is still available as a fallback option for vSphere 6.x. In this mode, vCenter Server checks that the certificate is formatted correctly, but does not check the validity of the certificate. Even expired certificates are accepted. Do not use this mode unless you encounter problems that you cannot resolve with one of the other two modes. Some vCenter Server 6.x and later services might not work correctly in thumbprint mode.

To change the certificate mode to use a different type of certificate, see [ESXi Certificate Mode Switch Workflows](#) and [Change the ESXi Certificate Mode](#).

ESXi Certificate Expiration

You can view information about certificate expiration for certificates that are signed by VMCA or a third-party CA in the vSphere Client. You can view the information for all hosts that vCenter Server manages or for individual hosts. A yellow alarm is raised if the certificate is in the **Expiring Shortly** state (less than eight months). A red alarm is raised if the certificate is in the **Expiration Imminent** state (less than two months).

ESXi Provisioning and Certificates

When you boot an ESXi host from installation media, the host initially has an autogenerated certificate. When you add a host to the vCenter Server system, vCenter Server provisions the host with a certificate that is signed by VMCA as the root CA.

You can also use custom certificates that are signed by a third-party or an enterprise CA for ESXi hosts.

ESXi Provisioning and Certificates in Auto Deploy

The process is similar for hosts that are provisioned with Auto Deploy. However, because those hosts do not store any state, the signed certificate is stored by the Auto Deploy server in its local certificate store. The certificate is reused during subsequent boots of the ESXi hosts. An Auto Deploy server is part of any embedded deployment or vCenter Server system.

If VMCA is not available when an Auto Deploy host boots the first time, the host first attempts to connect. If the host cannot connect, it cycles through shutdown and reboot until VMCA becomes available and the host can be provisioned with a signed certificate.

You can make Auto Deploy a subordinate Certificate Authority of a third-party Certificate Authority. In this case, the generated certificates are signed with the Auto Deploy SSL key. See [Make Auto Deploy a Subordinate Certificate Authority](#).

In ESXi 8.0 and later, you can use custom certificates (certificates signed by a Certificate Authority) with Auto Deploy. When the host starts, Auto Deploy associates the custom certificate with either a MAC address or the BIOS UUID of the ESXi host. See [Use Custom Certificates with Auto Deploy](#).

Required Privileges for ESXi Certificate Management

The **Certificates > Manage Certificates** privilege is required for users to manage your ESXi host certificates.

ESXi Host Name and IP Address Changes

An ESXi host name or IP address change might affect whether vCenter Server considers a host certificate valid. How you added the ESXi host to vCenter Server affects whether manual intervention is necessary. Manual intervention means that you either reconnect the host, or you remove the host from vCenter Server and add it back.

Table 241: When Host Name or IP Address Changes Require Manual Intervention

ESXi Host added to vCenter Server using...	ESXi Host name changes	ESXi IP address changes
Host name	vCenter Server connectivity problem. Manual intervention required.	No intervention required.
IP address	No intervention required.	vCenter Server connectivity problem. Manual intervention required.

ESXi Host Upgrades and Certificates

If you upgrade an ESXi host to ESXi6.7 or later, the upgrade process replaces the self-signed (thumbprint) certificates with VMCA-signed certificates. If the ESXi host uses custom certificates, the upgrade process retains those certificates even if those certificates are expired or invalid.

The recommended upgrade workflow depends on the current certificates.

Host Provisioned with Thumbprint Certificates

If your host is currently using thumbprint certificates, it is automatically assigned VMCA certificates as part of the upgrade process.

Note: You cannot provision legacy hosts with VMCA certificates. You must upgrade those hosts to ESXi6.7 or later.

Host Provisioned with Custom Certificates

If your host is provisioned with custom certificates, usually third-party CA-signed certificates, those certificates remain in place during upgrade. Change the certificate mode to **Custom** to ensure that the certificates are not replaced accidentally during a certificate refresh later.

Note: If your environment is in VMCA mode, and you refresh the certificates from the vSphere Client, any existing certificates are replaced with certificates that are signed by VMCA.

Going forward, vCenter Server monitors the certificates and displays information, for example, about certificate expiration, in the vSphere Client.

Hosts Provisioned with Auto Deploy

Hosts that are being provisioned by Auto Deploy are always assigned new certificates when they are first booted with ESXi6.7 or later software. When you upgrade a host that is provisioned by Auto Deploy, the Auto Deploy server generates a certificate signing request (CSR) for the host and submits it to VMCA. VMCA stores the signed certificate for the host. When the Auto Deploy server provisions the host, it retrieves the certificate from VMCA and includes it as part of the provisioning process.

You can use Auto Deploy with custom certificates.

ESXi Certificate Mode Switch Workflows

By default, VMware Certificate Authority (VMCA) provisions ESXi with certificates. Use custom mode when replacing VMCA certificates with custom certificates. Use legacy thumbprint mode for debugging. If you do require a mode switch, review the potential impact before you start.

For an explanation of certificate modes, see [Certificates and Certificate Modes](#).

Using Custom ESXi Certificates

Note: When switching from using VMCA certificates to custom certificates, make sure to allow time for organizational approval and fulfillment processes when generating certificates. Also, plan accordingly so that the current certificate does not expire during the switchover.

If your company policy requires that you use a different root CA than VMCA, you can switch the certificate mode in your environment after careful planning. The workflow is as follows.

1. Switch to **custom** mode. See [Change the ESXi Certificate Mode](#).
Switching the mode enables the vSphere Client to activate the **Manage with External CA** drop-down, allowing you to generate the Certificate Signing Request.
2. Add the root certificate of the custom CA to VMware Endpoint Certificate Store (VECS).
3. Generate the Certificate Signing Request and obtain the certificates that you want to use.
You might need to wait some time for the CSR to be returned.
4. Import the custom CA certificate to the vCenter Server host.
Allow some time for the vCenter Server to distribute the custom CA certificate to the ESXi hosts.

Switching from Custom CA Mode to VMCA Mode

If you are using custom CA mode and decide that using VMCA works better in your environment, you can perform the mode switch after careful planning. The workflow is as follows.

1. Remove all hosts from the vCenter Server system.
2. On the vCenter Server system, remove the root certificate of the third-party CA from VECS.
3. Switch to **vmca** mode. See [Change the ESXi Certificate Mode](#).

4. Add the hosts to the vCenter Server system.

Note: Any other workflow for this mode switch might result in unpredictable behavior.

Retaining Thumbprint Mode Certificates During Upgrade

The switch from VMCA mode to thumbprint mode might be necessary if you encounter problems with the VMCA certificates. In thumbprint mode, the vCenter Server system checks only whether a certificate exists and is formatted correctly, and does not check whether the certificate is valid. See [Change the ESXi Certificate Mode](#) for instructions.

Switching from Thumbprint Mode to VMCA Mode

If you use thumbprint mode and you want to start using VMCA-signed certificates, the switch requires some planning. The workflow is as follows.

1. Remove all ESXi hosts from the vCenter Server system.
2. Switch to **vmca** mode. See [Change the ESXi Certificate Mode](#).
3. Add the ESXi hosts to the vCenter Server system.

Note: Any other workflow for this mode switch might result in unpredictable behavior.

Switching from Custom CA Mode to Thumbprint Mode

If you are encountering problems with your custom CA, consider switching to thumbprint mode temporarily. The switch works seamlessly if you follow the instructions in [Change the ESXi Certificate Mode](#). After the mode switch, the vCenter Server system checks only the format of the certificate and no longer checks the validity of the certificate itself.

Switching from Thumbprint Mode to Custom CA Mode

If you set your environment to thumbprint mode during troubleshooting, and you want to start using custom CA mode, you must first generate the required certificates. The workflow is as follows.

1. Remove all ESXi hosts from the vCenter Server system.
2. Add the custom CA root certificate to the TRUSTED_ROOTS store on VECS on the vCenter Server system. See [Update the vCenter Server TRUSTED_ROOTS Store \(Custom Certificates\)](#).
3. For each ESXi host:
 - a. Deploy the custom CA certificate and key.
 - b. Restart services on the host.
4. Switch to **custom** mode. See [Change the ESXi Certificate Mode](#).
5. Add the ESXi hosts to the vCenter Server system.

ESXi Certificate Default Settings

When a host is added to a vCenter Server system, vCenter Server sends a Certificate Signing Request (CSR) for the host to VMCA. Most of the default values are well suited for many situations, but company-specific information can be changed.

You can change many of the default settings using the vSphere Client. Consider changing the organization, and location information. See [Change ESXi Certificate Default Settings](#).

Table 242: ESXi CSR Settings

Parameter	Default Value	Advanced Option
Key Size	2048	N.A.

Parameter	Default Value	Advanced Option
Key Algorithm	RSA	N.A.
Certificate Signature Algorithm	sha256WithRSAEncryption	N.A.
Common Name	Name of the host if the host was added to vCenter Server by host name. IP address of the host if the host was added to vCenter Server by IP address.	N.A.
Country	US	vpzd.certmgmt.certs.cn.country
Email address	vmca@vmware.com	vpzd.certmgmt.certs.cn.email
Locality (City)	Palo Alto	vpzd.certmgmt.certs.cn.localityName
Organization Unit Name	VMware Engineering	vpzd.certmgmt.certs.cn.organizationalUnitName
Organization Name	VMware	vpzd.certmgmt.certs.cn.organizationName
State or province	California	vpzd.certmgmt.certs.cn.state
Number of days the certificate is valid.	1825	vpzd.certmgmt.certs.daysValid
Hard threshold for the certificate expiration. vCenter Server raises a red alarm when this threshold is reached.	30 days	vpzd.certmgmt.certs.hardThreshold
Poll interval for vCenter Server certificate validity checks.	5 days	vpzd.certmgmt.certs.pollIntervalDays
Soft threshold for the certificate expiration. vCenter Server raises an event when this threshold is reached.	240 days	vpzd.certmgmt.certs.softThreshold
Mode that vCenter Server users to determine whether existing certificates are replaced. Change this mode to retain custom certificates during upgrade. See ESXi Host Upgrades and Certificates .	vmca You can also specify thumbprint or custom. See Change the ESXi Certificate Mode .	vpzd.certmgmt.mode

Change ESXi Certificate Default Settings

When an ESXi host is added to a vCenter Server system, vCenter Server sends a Certificate Signing Request (CSR) for the host to VMCA. You can change some of the default settings in the CSR using the vCenter Server Advanced Settings in the vSphere Client.

See the previous table for a list of default settings. Some of the defaults cannot be changed.

1. In the vSphere Client, select the vCenter Server system that manages the hosts.
2. Click **Configure**, and click **Advanced Settings**.
3. Click **Edit Settings**.
4. Click the **Filter** icon in the Name column, and in the Filter box, enter `vpzd.certmgmt` to display only certificate management parameters.
5. Change the value of the existing parameters to follow your company policy and click **Save**.

The next time you add a host to vCenter Server, the new settings are used in the CSR that vCenter Server sends to VMCA and in the certificate that is assigned to the host.

Changes to certificate metadata only affect new certificates. If you want to change the certificates of hosts that are already managed by the vCenter Server system, you can disconnect and reconnect the hosts or renew the certificates.

View Certificate Expiration Information for ESXi Hosts

For ESXi hosts that are in VMCA mode or custom mode, you can view certificate details from the vSphere Client. The certificate information enables you to determine whether any of your certificates expire soon. You can also use this information to debug certificate problems.

You cannot view certificate status information for ESXi hosts in thumbprint mode. You can view information for multiple ESXi hosts or a single ESXi host. The multiple-host view displays only the Certificate Valid To date information.

1. Log in to the vCenter Server by using the vSphere Client.
2. Browse the inventory list and select the vCenter Server instance.
3. Get the certificate information.

Single Host or Multiple Hosts	Steps
Single	1. Browse to the ESXi host. 2. Click Configure. 3. Under System, click Certificate.
Multiple	1. Select Hosts & Clusters > Hosts. By default, the Hosts display does not include the certificate status. 2. To show or hide columns, click Manage Columns. 3. Select the Certificate Valid To check box, and scroll to the right if necessary to view the added column. The certificate information displays when the certificate expires. 4. (Optional) Deselect other columns to make it easier to see what you are interested in.

4. Review the certificate information.

The following information is available only in the single-host view.

Field	Description
Subject	The subject used during certificate generation.

Field	Description										
Issuer	The issuer of the certificate.										
Valid From	Date on which the certificate was generated.										
Valid To	Date on which the certificate expires.										
Status	Status of the certificate, one of the following. <table> <tr> <td>Good</td><td>Normal operation.</td></tr> <tr> <td>Expiring</td><td>Certificate expires soon.</td></tr> <tr> <td>Expiring shortly</td><td>Certificate is eight months or less away from expiration (Default).</td></tr> <tr> <td>Expiration imminent</td><td>Certificate is two months or less away from expiration (Default).</td></tr> <tr> <td>Expired</td><td>Certificate is not valid because it expired.</td></tr> </table>	Good	Normal operation.	Expiring	Certificate expires soon.	Expiring shortly	Certificate is eight months or less away from expiration (Default).	Expiration imminent	Certificate is two months or less away from expiration (Default).	Expired	Certificate is not valid because it expired.
Good	Normal operation.										
Expiring	Certificate expires soon.										
Expiring shortly	Certificate is eight months or less away from expiration (Default).										
Expiration imminent	Certificate is two months or less away from expiration (Default).										
Expired	Certificate is not valid because it expired.										

Note: If a host is added to vCenter Server or reconnected after a disconnect, vCenter Server renews the certificate if the status is Expired, Expiring, Expiring shortly, or Expiration imminent. The status is Expiring if the certificate is valid for less than eight months, Expiring shortly if the certificate is valid for less than two months, and Expiration imminent if the certificate is valid for less than one month.

Renew the certificates that are about to expire. See [Renew or Refresh ESXi Certificates](#).

Renew or Refresh ESXi Certificates

If you use the VMware Certificate Authority (VMCA) to assign certificates to your hosts, you can renew those certificates from the vSphere Client. If you use either VMCA certificates or custom certificates, you can refresh all certificates from the TRUSTED_ROOTS store associated with vCenter Server.

Verify the following:

- If using VMCA certificates, the certificate mode is set to **vmca**.
- If using custom certificates, the certificate mode is set to **custom**.
- The ESXi hosts are connected to the vCenter Server system.
- There is proper time synchronization between the vCenter Server system and the ESXi hosts.
- DNS resolution works between the vCenter Server system and the ESXi hosts.
- The vCenter Server system's MACHINE_SSL_CERT and Trusted_Root certificates are valid and have not expired. See the VMware knowledge base article at <https://kb.vmware.com/s/article/2111411>.
- The ESXi hosts are not in maintenance mode.

Note: If you use custom certificates, and you need to renew them, re-import the certificates.

You can use the vSphere Client to renew your VMCA certificates when they are about to expire, or if you want to provision the host with a new certificate for other reasons. If you do not renew the VMCA certificate before it expires, disconnecting the host and reconnecting it causes vCenter Server to renew the certificate. The act of re-adding the host to vCenter Server reestablishes trust, and enables vCenter Server to unconditionally issue the renewed certificate.

By default, vCenter Server renews the VMCA certificate of a host with status Expired, Expiration imminent, or Expiring shortly, and each time the host is added to the inventory or reconnected.

You cannot renew an ESXi certificate with an expiration date beyond that of the expiration date of the trusted root certificate. For example, even if the ESXivpxd.certmgmt.certs.daysValid advanced option is set to five years, and your trusted root certificate is set to expire in two years, the ESXi certificate expiration date is limited to two years.

You can use the vSphere Client to push all certificates currently in the TRUSTED_ROOTS store in the vCenter Server VECS store to the ESXi host. Use this capability if you need to refresh the trusted roots on an ESXi host. This capability exists for both VMCA and custom certificates.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under **System**, click **Certificate**.
You can view the details about the certificate of the selected host.
4. Select the appropriate option based on the type of certificate used.

Option	Description
Manage with VMCA > Renew	Retrieves a fresh signed certificate for the host from VMCA.
Manage with VMCA > Refresh CA Certificates or Manage with External CA > Refresh CA Certificates	Pushes all certificates in the TRUSTED_ROOTS store in the vCenter Server VECS store to the host.

Change the ESXi Certificate Mode

Use VMware Certificate Authority (VMCA) to provision the certificates of the ESXi hosts in your environment unless your corporate policy requires that you use custom certificates. To use custom certificates with a different root CA, edit the advanced vCenter Server setting, `vpdx.certmgmt.mode`. After the change, the hosts are no longer automatically provisioned with VMCA certificates when you refresh the certificates. You are responsible for the certificate management in your environment.

You can use the advanced vCenter Server settings to change to thumbprint mode or to custom CA mode. Use thumbprint mode only as a fallback option.

1. In the vSphere Client, select the vCenter Server system that manages the hosts.
2. Click **Configure**, and under Settings, click **Advanced Settings**.
3. Click **Edit Settings**.
4. Click the **Filter** icon in the Name column, and in the Filter box, enter `vpdx.certmgmt` to display only certificate management parameters.

Note: The available options are **vmca**, **custom**, and **thumbprint**.

5. Change the value of `vpdx.certmgmt.mode` to **custom** if you intend to manage your own certificates, and to **thumbprint** if you temporarily want to use thumbprint mode, and click **Save**.

Replacing the Default ESXi Certificate with a Custom Certificate

Your company's security policy might require that you replace the default ESXi SSL certificate with a third-party certificate authority (CA) signed certificate on all your hosts.

By default, vSphere components use the VMCA-signed certificate and key that are created during installation. If you accidentally delete the VMCA-signed certificate, remove the host from its vCenter Server system, and add it back. When you add the host, vCenter Server requests a new certificate from VMCA and provisions the host with it.

You can replace VMCA-signed certificates with certificates from a trusted certificate authority, either a commercial CA or an organizational CA, if your company policy requires it.

You can replace the default certificates with custom certificates using the vSphere Client or the CLI.

Note: You can also use the `vim.CertificateManager` and `vim.host.CertificateManager` managed objects in the vSphere Web Services SDK. See the vSphere Web Services SDK documentation.

Before you replace the certificate, you must update the TRUSTED_ROOTS store in VECS on the vCenter Server system that manages the host to ensure that the vCenter Server and the ESXi host have a trust relationship.

Note: If you are replacing SSL certificates on an ESXi host that is part of a vSAN cluster, follow the steps that are in the VMware knowledge base article at <https://kb.vmware.com/s/article/56441>.

Requirements for ESXi Certificate Signing Requests for Custom Certificates

Use a CSR with these characteristics:

- Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
- PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When keys are added to VECS, they are converted to PKCS8.
- x509 version 3
- For root certificates, the CA extension must be set to true, and the cert sign must be in the list of requirements.
- SubjectAltName must contain DNS Name=<machine_FQDN>.
- CRT format
- Contains the following Key Usages: Digital Signature, Key Encipherment
- Start time of one day before the current time.
- CN (and SubjectAltName) set to the host name (or IP address) that the ESXi host has in the vCenter Server inventory.

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 and 3072. See [Considerations When Using FIPS](#).

vSphere does not support the following certificates.

- Certificates with wildcards.
- The algorithms md2WithRSAEncryption, md5WithRSAEncryption, RSASSA-PSS, dsaWithSHA1, ecdsa_with_SHA1, and sha1WithRSAEncryption are not supported.

To generate the CSR using the vSphere Client, see [Generate a Certificate Signing Request for a Custom Certificate Using the vSphere Client](#).

For information about generating the CSR using the CLI, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2113926>.

Generate a Certificate Signing Request for a Custom Certificate Using the vSphere Client

Starting in vSphere 8.0 Update 3, you can use the vSphere Client to generate a Certificate Signing Request (CSR) for the ESXi SSL certificate and to replace the certificate once it is ready.

Change the certificate mode to **custom**. See [Change the ESXi Certificate Mode](#). Switching the mode enables the vSphere Client to activate the **Manage with External CA** drop-down, allowing you to generate the Certificate Signing Request.

Warning: Generating a Certificate Signing Request generates a new private key. Do not generate another Certificate Signing Request during the process of replacing certificates. If you do, the previously generated CSR and subsequent certificate will no longer be valid.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under **System**, click **Certificate**.
4. From the **Manage with External CA** drop-down, select either **Generate CSR Using IP** or **Generate CSR Using FQDN**.

vCenter Server identifies the option previously used to generate the certificate on the ESXi host.

5. Select either **Copy to Clipboard** or **Download**, depending on how you want to generate the certificate signing request.

You can now send the CSR to the certificate authority, or use the CSR to generate the certificate internally.

Related Links

[Replace the Default Certificate with a Custom Certificate Using the vSphere Client on page 2779](#)

Starting in vSphere 8.0 Update 3, you can replace the default VMCA-signed ESXi certificates with custom certificates from the vSphere Client.

[Replace the Default Certificate and Key from theESXi Shell on page 2780](#)

You can replace the default VMCA-signedESXi certificates from theESXi Shell.

[Replace the Default Certificate Using HTTPS PUT on page 2781](#)

You can use third-party applications to upload certificates and key. Applications that support HTTPS PUT operations work with the HTTPS interface that is included withESXi.

[Update thevCenter Server TRUSTED_ROOTS Store \(Custom Certificates\) on page 2782](#)

If you set up yourESXi hosts to use custom certificates, you must update the `TRUSTED_ROOTS` store on thevCenter Server system that manages the hosts.

Replace the Default Certificate with a Custom Certificate Using the vSphere Client

Starting in vSphere 8.0 Update 3, you can replace the default VMCA-signed ESXi certificates with custom certificates from the vSphere Client.

- Generate the certificate signing request and send it to the certificate authority. See [Generate a Certificate Signing Request for a Custom Certificate Using the vSphere Client](#).
- When the certificate authority returns the certificate, store it on the ESXi hosts.
- Make sure that the ESXi certificate mode is set to **custom**. See [Change the ESXi Certificate Mode](#).
- Update the trusted root store. See [Update thevCenter Server TRUSTED_ROOTS Store \(Custom Certificates\)](#).

When importing the custom certificate, make sure that you:

- Add your entire CA certificate chain before proceeding with the replacement.
- Ensure that you provide the correct CA certificate for your environment. The import and replace process does not perform checking on the certificate that you use.
- Ensure that there are no SHA1 hashes in the certificate chain. SHA1 is not supported.
- Add the root CA to VECS before proceeding. If not, the host disconnects immediately after the certificate replacement.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under **System**, click **Certificate**.
4. From the **Manage with External CA** drop-down, select **Import and replace**.
5. Select the replace option.

Option	Description
Replace with external CA certificate where CSR generated by ESXi (private key embedded)	Use this option if you generated the CSR on ESXi, in which case, the private key is stored on ESXi.
Replace with external CA certificate where CSR generated from a certificate authority (requires private key)	Use this option if you sent the CSR to a third-party certificate authority, and received back the certificate and private key.

6. Click **Next**.
7. Browse for the certificate, or certificate and private key.
8. Review the information, then click **Import and Replace**.

The custom certificate replaces the existing certificate.

Related Links

[Generate a Certificate Signing Request for a Custom Certificate Using the vSphere Client on page 2778](#)

Starting in vSphere 8.0 Update 3, you can use the vSphere Client to generate a Certificate Signing Request (CSR) for the ESXi SSL certificate and to replace the certificate once it is ready.

[Replace the Default Certificate and Key from theESXi Shell on page 2780](#)

You can replace the default VMCA-signedESXi certificates from theESXi Shell.

[Replace the Default Certificate Using HTTPS PUT on page 2781](#)

You can use third-party applications to upload certificates and key. Applications that support HTTPS PUT operations work with the HTTPS interface that is included withESXi.

[Update thevCenter Server TRUSTED_ROOTS Store \(Custom Certificates\) on page 2782](#)

If you set up yourESXi hosts to use custom certificates, you must update the `TRUSTED_ROOTS` store on thevCenter Server system that manages the hosts.

Replace the Default Certificate and Key from theESXi Shell

You can replace the default VMCA-signedESXi certificates from theESXi Shell.

- If you want to use third-party CA-signed certificates, generate the certificate request, send it to the certificate authority, and store the certificates on each ESXi host.
- If necessary, enable the ESXi Shell or enable SSH traffic from the vSphere Client.
- All file transfers and other communications occur over a secure HTTPS session. The user who is used to authenticate the session must have the privilege **Host > Config > AdvancedConfig** on the host.

Note: Before you replace the certificates, update thevCenter Server `TRUSTED_ROOTS` store. See [Update thevCenter Server TRUSTED_ROOTS Store \(Custom Certificates\)](#).

1. Log in to theESXi Shell, either directly from the DCUI or from an SSH client, as a user with administrator privileges.
2. In the directory `/etc/vmware/ssl`, rename the existing certificates using the following commands.

```
mv rui.crt orig.rui.crt
mv rui.key orig.rui.key
```

3. Copy the certificates that you want to use to `/etc/vmware/ssl`.
4. Rename the new certificate and key to `rui.crt` and `rui.key`.
5. Restart the host after you install the new certificate.

Alternatively, you can put the host into maintenance mode, install the new certificate, use the Direct Console User Interface (DCUI) to restart the management agents, and set the host to exit maintenance mode.

Related Links

[Generate a Certificate Signing Request for a Custom Certificate Using the vSphere Client on page 2778](#)

Starting in vSphere 8.0 Update 3, you can use the vSphere Client to generate a Certificate Signing Request (CSR) for the ESXi SSL certificate and to replace the certificate once it is ready.

[Replace the Default Certificate with a Custom Certificate Using the vSphere Client on page 2779](#)

Starting in vSphere 8.0 Update 3, you can replace the default VMCA-signed ESXi certificates with custom certificates from the vSphere Client.

[Replace the Default Certificate Using HTTPS PUT on page 2781](#)

You can use third-party applications to upload certificates and key. Applications that support HTTPS PUT operations work with the HTTPS interface that is included with ESXi.

[Update the vCenter Server TRUSTED_ROOTS Store \(Custom Certificates\) on page 2782](#)

If you set up your ESXi hosts to use custom certificates, you must update the `TRUSTED_ROOTS` store on the vCenter Server system that manages the hosts.

Replace the Default Certificate Using HTTPS PUT

You can use third-party applications to upload certificates and key. Applications that support HTTPS PUT operations work with the HTTPS interface that is included with ESXi.

- If you want to use third-party CA-signed certificates, generate the certificate request, send it to the certificate authority, and store the certificates on each ESXi host.
- If necessary, enable the ESXi Shell or enable SSH traffic from the vSphere Client.
- All file transfers and other communications occur over a secure HTTPS session. The user who is used to authenticate the session must have the privilege **Host > Config > AdvancedConfig** on the host.

Note: Before you replace the certificates, update the vCenter Server `TRUSTED_ROOTS` store. See [Update the vCenter Server TRUSTED_ROOTS Store \(Custom Certificates\)](#).

1. Back up the existing certificates.
2. Set up basic access authentication, in which you supply a Base64 encoded username and password, separated with a single colon (:). For more information, see https://en.wikipedia.org/wiki/Basic_access_authentication.
3. In your upload application, process each file as follows:
 - a) Open the file.
 - b) Publish the file to one of these locations.

Option	Description
Certificates	<code>https://hostname/host/ssl_cert</code>
Keys	<code>https://hostname/host/ssl_key</code>

The `/host/ssl_cert` and the `host/ssl_key` locations link to the certificate files in `/etc/vmware/ssl`.

4. Restart the host.

Alternatively, you can put the host into maintenance mode, install the new certificate, use the Direct Console User Interface (DCUI) to restart the management agents, and set the host to exit maintenance mode.

Related Links

[Generate a Certificate Signing Request for a Custom Certificate Using the vSphere Client on page 2778](#)

Starting in vSphere 8.0 Update 3, you can use the vSphere Client to generate a Certificate Signing Request (CSR) for the ESXi SSL certificate and to replace the certificate once it is ready.

[Replace the Default Certificate with a Custom Certificate Using the vSphere Client on page 2779](#)

Starting in vSphere 8.0 Update 3, you can replace the default VMCA-signed ESXi certificates with custom certificates from the vSphere Client.

[Replace the Default Certificate and Key from the ESXi Shell on page 2780](#)

You can replace the default VMCA-signed ESXi certificates from the ESXi Shell.

[Update the vCenter Server TRUSTED_ROOTS Store \(Custom Certificates\) on page 2782](#)

If you set up your ESXi hosts to use custom certificates, you must update the `TRUSTED_ROOTS` store on the vCenter Server system that manages the hosts.

Update thevCenter Server TRUSTED_ROOTS Store (Custom Certificates)

If you set up your ESXi hosts to use custom certificates, you must update the `TRUSTED_ROOTS` store on thevCenter Server system that manages the hosts.

Replace the certificates on each host with custom certificates.

Note: This step is not required if thevCenter Server system is also running with custom certificates issued by the same CA as those installed on theESXi hosts.

1. To update the vCenter Server `TRUSTED_ROOTS` store using vSphere Client, see [Add a Trusted Root Certificate to the Certificate Store Using the vSphere Client](#).
2. To update the vCenter Server `TRUSTED_ROOTS` store using using command line interface, log in to thevCenter Server shell of thevCenter Server system that manages theESXi hosts.
3. To add the new certificates to the `TRUSTED_ROOTS` store, run `dir-cli`, for example:

```
/usr/lib/vmware-vmafd/bin/dir-cli trustedcert publish --certpath_to_RootCA
```
4. When prompted, provide the Single Sign-On Administrator credentials.
5. If your custom certificates are issued by an intermediate CA, you must also add the intermediate CA to the `TRUSTED_ROOTS` store on thevCenter Server, for example:

```
/usr/lib/vmware-vmafd/bin/dir-cli trustedcert publish --certpath_to_intermediateCA
```

Set certificate mode to Custom. If the certificate mode is VMCA, the default, and you perform a certificate refresh, your custom certificates are replaced with VMCA-signed certificates. See [Change the ESXi Certificate Mode](#).

Related Links

[Generate a Certificate Signing Request for a Custom Certificate Using the vSphere Client on page 2778](#)

Starting in vSphere 8.0 Update 3, you can use the vSphere Client to generate a Certificate Signing Request (CSR) for the ESXi SSL certificate and to replace the certificate once it is ready.

[Replace the Default Certificate with a Custom Certificate Using the vSphere Client on page 2779](#)

Starting in vSphere 8.0 Update 3, you can replace the default VMCA-signed ESXi certificates with custom certificates from the vSphere Client.

[Replace the Default Certificate and Key from theESXi Shell on page 2780](#)

You can replace the default VMCA-signedESXi certificates from theESXi Shell.

[Replace the Default Certificate Using HTTPS PUT on page 2781](#)

You can use third-party applications to upload certificates and key. Applications that support HTTPS PUT operations work with the HTTPS interface that is included withESXi.

Make Auto Deploy a Subordinate Certificate Authority

By default, the Auto Deploy server provisions each host with certificates that are signed by the VMware Certificate Authority (VMCA). You can set up the Auto Deploy server to provision all hosts with custom certificates that are not signed by VMCA. In that scenario, the Auto Deploy server becomes a subordinate certificate authority of your third-party certificate authority (CA).

- Request a certificate from your CA. The certificate must meet these requirements.
 - Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
 - PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When keys are added to VECS, they are converted to PKCS8.
 - x509 version 3
 - For root certificates, the CA extension must be set to true, and the cert sign must be in the list of requirements.
 - SubjectAltName must contain DNS Name=<machine_FQDN>.
 - CRT format
 - Contains the following Key Usages: Digital Signature, Key Encipherment
 - Start time of one day before the current time.
 - CN (and SubjectAltName) set to the host name (or IP address) that the ESXi host has in the vCenter Server inventory.

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 and 3072. See [Considerations When Using FIPS](#).

- Name the certificate and the key files `rbd-ca.crt` and `rbd-ca.key`.

1. Back up the default ESXi certificates.

The certificates are in the `/etc/vmware-rbd/ssl/` directory.

2. Stop the vSphere Authentication Proxy service.

Tool	Steps
vCenter Server Management Interface	- In a Web browser, go to the vCenter Server Management Interface, <code>https://vcenter-IP-address-or-FQDN:5480</code>. - Log in as root. The default root password is the password that you set while deploying the vCenter Server. - Click Services, and click VMware vSphere Authentication Proxy. - Click Stop.
CLI	<code>service-control --stop vmcam</code>

3. On the system where the Auto Deploy service runs, replace `rbd-ca.crt` and `rbd-ca.key` in `/etc/vmware-rbd/ssl/` with your custom certificate and key files.

4. On the system where the Auto Deploy service runs, run the following command to update the TRUSTED_ROOTS store inside the VMware Endpoint Certificate Store (VECS) to use your new certificates.

```
/usr/lib/vmware-vmafd/bin/dir-cli trustedcert publish --cert /etc/vmware-rbd/ssl/rbd-ca.crt
/usr/lib/vmware-vmafd/bin/vecs-cli force-refresh
```

5. Create a `castore.pem` file that contains what is in the TRUSTED_ROOTS store and place the file in the `/etc/vmware-rbd/ssl/` directory.

In custom mode, you are responsible for maintaining this file.

6. Change the ESXi certificate mode for the vCenter Server system to `custom`.

See [Change the ESXi Certificate Mode](#).

7. Restart the vCenter Server service and start the Auto Deploy service.

The next time you provision a host that is set up to use Auto Deploy, the Auto Deploy server generates a certificate. The Auto Deploy server uses the root certificate that you added to the TRUSTED_ROOTS store.

Note: If you encounter problems with Auto Deploy after certificate replacement, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2000988>.

Use Custom Certificates with Auto Deploy

In vSphere 8.0 and later, you can set up the Auto Deploy server to provision ESXi hosts with custom certificates that are signed by a third-party certificate authority (CA) or your own internal CA. By default, the Auto Deploy server provisions ESXi hosts with certificates that are signed by the VMware Certificate Authority (VMCA).

Ensure that you have the following:

- Request a certificate from your certificate authority. The certificate must meet these requirements.
 - Key size: 2048 bits (minimum) to 8192 bits (maximum) (PEM encoded)
 - PEM format. VMware supports PKCS8 and PKCS1 (RSA keys). When keys are added to VECS, they are converted to PKCS8.
 - x509 version 3
 - CRT format
 - CA extension set to true
 - Key usage of Certificate Signing
 - Start time of one day before the current time

Note: vSphere's FIPS certificate only validates RSA key sizes of 2048 and 3072. See [Considerations When Using FIPS](#).

- ESXi host MAC address or BIOS UUID. Evaluate which approach makes best sense for your environment. The BIOS UUID is more stable and less subject to change than the MAC address. If you change network adapters in an ESXi host, the MAC address changes. However, the MAC address might be more familiar to work with, and easier to obtain than the BIOS UUID.
- At least PowerCLI version 12.6.0. For more information on Auto Deploy PowerCLI cmdlets, see the Auto Deploy PowerCLI Cmdlet Overview topic in the *VMware ESXi Installation and Setup* documentation.

Ensure that you have the following privileges:

- Add custom certificate: **Autodeploy > Rule > Create**
- Get custom certificate information: **System > Read**

Prior to vSphere 8.0, your options for managing certificates with Auto Deploy include:

- Using vCenter Server and the built-in VMware Certificate Authority (the default).
- Making Auto Deploy a subordinate CA of a third-party CA. In this case, the Auto Deploy SSL key signs the certificates.

In vSphere 8.0 and later, you can upload custom certificates to Auto Deploy that are signed by either a third-party CA or your own internal CA. Auto Deploy associates the custom certificate with either the MAC address or BIOS UUID of the ESXi host. Each time an Auto Deploy host starts, Auto Deploy checks for a custom certificate. If Auto Deploy finds a custom certificate, it uses that certificate instead of generating one through the VMCA.

The high-level steps for this task include:

1. Generating the custom certificate request for either a third-party CA or for your own internal CA.
2. Obtaining the signed custom certificate (key and certificate) and storing it locally.
3. If you are using a third-party CA, and if not previously done, ensuring the root certificate of your CA is uploaded to the TRUSTED_ROOTS store on the vCenter Server.
4. Uploading the custom certificate to Auto Deploy and associating the certificate with either the MAC address or BIOS UUID of an ESXi host.
5. Booting the ESXi host.

When you assign a custom certificate to an ESXi host, Auto Deploy pushes the certificate to the host on its next boot from Auto Deploy.

Be aware of the following considerations when using custom certificates and Auto Deploy.

- You must use the PowerCLI `Add-CustomCertificate`, `Remove-CustomCertificate`, and `List-CustomCertificate` cmdlets to manage custom certificates used with Auto Deploy. The capability to manage custom certificates is not available in the vSphere Client.
- To refresh a custom certificate used for Auto Deploy, you must run the `Add-CustomCertificate` cmdlet again.
- Be sure to examine your custom certificate for potential errors. Auto Deploy verifies only that the custom certificate complies with the X.509 certificate standards and that the expiration threshold of the certificate is set to at least 240 days. Auto Deploy does not perform any other certificate validation or checking. To change the certificate threshold, you can run the `Set-DeployOption -Key certificate-refresh-threshold` cmdlet.
- If you later remove a custom certificate from an ESXi host using the `Remove-CustomCertificate` cmdlet, you must restart the host for the change to take effect.

For more information about custom certificates and Auto Deploy, see the *VMware ESXi Installation and Setup* documentation.

1. Generate the certificate request.

- a) Using the requirements listed previously for the certificate request, create a configuration (.cfg) file.
- b) To generate a CSR file and a key file, run the `openssl req` command, passing in the configuration (.cfg) file.

For example:

```
openssl req -new -config custom_cert.cfg -days 4200 -sha256 -keyout rui.key -out rui.csr
```

In this command:

- `-new` generates a new certificate request.
- `-config custom_cert.cfg` specifies your custom .cfg file.
- `-days 4200` specifies 4200 days to certify the certificate for.
- `-sha256` specifies the message digest to sign the request with.
- `-keyout rui.key` specifies the file to write the newly created private key to.
- `-out rui.csr` specifies the output file to write to.

2. Either send the certificate request to your third-party CA, or, if you sign your own certificates, run the `openssl x509 -req` command to generate your custom certificate from your `rui.csr` file.

For example:

```
openssl x509 -req -in rui.csr -CA "/etc/vmware-rbd/ssl/rbd-ca.crt" -CAkey \
"/etc/vmware-rbd/ssl/rbd-ca.key" -extfile \
openssl.cfg -extensions x509 -CAserial "/etc/vmware-rbd/ssl/rbd-ca.srl" -days \
4200 -sha256 -out signed_rui.crt
```

In this command:

- `-in rui.csr` specifies the input file.
- `-CA "/etc/vmware-rbd/ssl/rbd-ca.crt"` specifies the directory to use for server certificate verification.
- `-CAkey "/etc/vmware-rbd/ssl/rbd-ca.key"` sets the CA private key to sign a certificate with.
- `-extfile openssl.cfg` specifies an additional, optional configuration file to read certificate extensions from.
- `-extensions x509` specifies to use x509 certificate extensions.
- `-CAserial "/etc/vmware-rbd/ssl/rbd-ca.srl"` uses the serial number in `rbd-ca.srl` to sign a certificate.
- `-days 4200` specifies 4200 days to certify the certificate for.
- `-sha256` specifies the message digest to sign the request with.
- `-out signed_rui.crt` specifies the output file to write to.

3. Optional: If you haven't previously uploaded the certificate of your signing certificate authority to the TRUSTED_ROOTS store inside the VMware Endpoint Certificate Store (VECS), perform the following steps on the vCenter Server where the Auto Deploy service runs.

- a) Using a tool such as WinSCP, copy the certificate to the vCenter Server.
- b) Log in to the vCenter Server using SSH and run the following command.

```
/usr/lib/vmware-vmafd/bin/dir-cli trustedcert publish --cert path_to_ca_certificate
```

4. Obtain either the ESXi host MAC address or BIOS UUID.

5. Perform the following steps to add the custom certificate to Auto Deploy.

- a) To connect to the vCenter Server, run the `Connect-VIServer` cmdlet.

```
Connect-VIServer -server VC_ip_address -User administrator_user -Password 'password'
```

- b) Optional: To view existing custom certificates, run the `Get-CustomCertificates` cmdlet.

The first time you add custom certificates, you don't see any certificates returned by this cmdlet.

- c) To associate the custom certificate with the ESXi host, run the `Add-CustomCertificate` cmdlet.

```
Add-CustomCertificate -HostID [MAC_Address | BIOS_UUID] -Certificate "path_to_custom_cert" -Key  
"path_to_custom_cert_key"
```

You can specify either the MAC address or the BIOS UUID of the host. Auto Deploy uploads the custom certificate to the host.

- d) To verify that the certificate was uploaded, run the `Get-CustomCertificates` cmdlet.

You see output similar to the following:

```
Name:      CustomHostCert-1
CertificateId: 1
HostId:      02:08:b0:8e:18:a2
ExpirationTime: 1 2/28/2033 10:45:50 AM
TimeCreated:  9/29/2022 7:40:28 AM
LastModified:  9/29/2022 7:40:28 AM
AssociatedHostName:
```

The `AssociatedHostName` is blank for now. After you start the host, the output reflects the name of the ESXi host associated with the custom certificate.

6. Start the ESXi host.

7. To verify that the custom certificate is associated with the vCenter Server, run the `Get-CustomCertificates` cmdlet again.

You see output resembling the following.

```
Name:      CustomHostCert-1
CertificateId: 1
HostId:      02:08:b0:8e:18:a2
ExpirationTime: 1 2/28/2033 10:45:50 AM
TimeCreated:  9/29/2022 7:40:28 AM
LastModified:  9/29/2022 7:40:28 AM
AssociatedHostName: host1.example.com
```

Now the `AssociatedHostName` contains the name of the ESXi host.

Restore ESXi Certificate and Key Files When Certificate Replacement Fails

When replacing a certificate on an ESXi host fails, the system creates certificate `.bak` files, which you can use to recover to the previous state.

The host certificate and key are located in `/etc/vmware/ssl/rui.crt` and `/etc/vmware/ssl/rui.key`.

When you replace a host certificate by using either the vSphere Client or the vSphere Web Services SDK `vim.CertificateManager` managed object, and the replacement fails, the system creates `.bak` files for the previous key and certificate files.

When certificate replacement fails, you can restore previous certificates by copying over the `.bak` files to the current certificate and key files.

Customizing ESXi Host Security

You can customize many of the essential security settings for your ESXi host through the Firewall, Services, and Security Profile panels available in the vSphere Client. The Security Profile is especially useful for single host management. If you are managing multiple hosts, consider using one of the VMware CLIs or SDKs and automating the customization.

Configuring the ESXi Firewall

ESXi includes a firewall that is enabled by default. At installation time, the ESXi firewall is configured to block incoming and outgoing traffic, except traffic for services that are enabled in the security profile of the host. You manage the firewall using the vSphere Client, the CLI, and the API.

As you open ports on the firewall, consider that unrestricted access to services running on an ESXi host can expose a host to outside attacks and unauthorized access. Reduce the risk by configuring the ESXi firewall to enable access only from authorized networks.

Note: The firewall also allows Internet Control Message Protocol (ICMP) pings and communication with DHCP and DNS (UDP only) clients.

You can manage ESXi firewall ports as follows:

- Use **Configure > Firewall** for each host in the vSphere Client. See [Manage ESXi Firewall Settings](#).
- Use ESXCLI commands from the command line or in scripts. See [Using ESXCLI Firewall Commands to Configure ESXi Behavior](#).
- Use a custom VIB if the port you want to open is not included in the security profile.

To install the custom VIB, you have to change the acceptance level of the ESXi host to `CommunitySupported`.

Note: If you engage VMware Technical Support to investigate a problem on an ESXi host with a `CommunitySupported` VIB installed, VMware Support can request you to uninstall this VIB. Such a request is a troubleshooting step to determine if that VIB is related to the problem being investigated.

The behavior of the NFS Client rule set (`nfsClient`) is different from other rule sets. When the NFS Client rule set is enabled, all outbound TCP ports are open for the destination hosts in the list of allowed IP addresses. See [NFS Client Firewall Behavior](#) for more information.

Manage ESXi Firewall Settings

You can configure incoming and outgoing firewall connections for a service or a management agent from the vSphere Client or at the command line.

This task describes how to use the vSphere Client to configure ESXi firewall settings. You can use the ESXi Shell or ESXCLI commands to configure ESXi at the command line to automate the firewall configuration. See [Using ESXCLI Firewall Commands to Configure ESXi Behavior](#) for examples of using ESXCLI to manipulate firewalls and firewall rules.

Note: If different services have overlapping port rules, enabling one service might implicitly activate other services. You can specify which IP addresses are allowed to access each service on the host to avoid this problem.

1. Log in to the vCenter Server by using the vSphere Client.
2. Browse to the host in the inventory.
3. Click **Configure**, then click **Firewall** under **System**.
You can toggle between incoming and outgoing connections by clicking **Incoming** and **Outgoing**.
4. In the Firewall section, click **Edit**.
5. Select from one of the service groups, **Ungrouped**, **Secure Shell**, and **Simple Network Management Protocol**.
6. Select the rule sets to be activated, or deselect the rule sets to be deactivated.
7. For some services, you can also manage service details by navigating to **Configure > System > Services**.
For more information about starting, stopping, and restarting services, see [Activate or Deactivate an ESXi Service](#).
8. For some services, you can explicitly specify IP addresses from which connections are allowed.
See [Add Allowed IP Addresses for an ESXi Host](#).
9. Click **OK**.

Add Allowed IP Addresses for an ESXi Host

By default, the firewall for each service allows access to all IP addresses. To restrict traffic, change each service to allow traffic only from your management subnet. You can also deselect some services if your environment does not use them.

To update the Allowed IP list for a service you can use the vSphere Client, ESXCLI, or PowerCLI. This task describes how to use the vSphere Client. See the topic titled "Manage the ESXi Firewall" in the *ESXCLI Concepts and Examples* documentation for instructions on using ESXCLI.

1. Log in to the vCenter Server by using the vSphere Client.
2. Browse to the ESXi host.
3. Click **Configure**, then click **Firewall** under **System**.
You can toggle between incoming and outgoing connections by clicking **Incoming** and **Outgoing**.
4. In the Firewall section, click **Edit**.
5. Select from one of the three service groups, **Ungrouped**, **Secure Shell**, and **Simple Network Management Protocol**.
6. To display the Allowed IP Addresses section, expand a service.
7. In the Allowed IP Addresses section, deselect **Allow connections from any IP address** and enter the IP addresses of networks that are allowed to connect to the host.
Separate IP addresses with commas. You can use the following address formats:
 - 192.168.0.0/24
 - 192.168.1.2, 2001::1/64
 - fd3e:29a6:0a81:e478::/64

8. Ensure that the service itself is selected.
9. Click **OK**.
10. Verify your change in the **Allowed IP addresses** column for the service.

Incoming and Outgoing Firewall Ports for ESXi Hosts

Open and close firewall ports for each service by using either the vSphere Client or the VMware Host Client.

includes a firewall that is enabled by default. At installation time, the firewall is configured to block incoming and outgoing traffic, except traffic for services that are enabled in the host's security profile. For the list of supported ports and protocols in the firewall, see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com/>.

The VMware Ports and Protocols Tool lists port information for services that are installed by default. If you install other VIBs on your host, additional services and firewall ports might become available. The information is primarily for services that are visible in the but the VMware Ports and Protocols Tool includes some other ports as well.

NFS Client Firewall Behavior

The NFS Client firewall rule set behaves differently than other ESXi firewall rule sets. ESXi configures NFS Client settings when you mount or unmount an NFS datastore. The behavior differs for different versions of NFS.

When you add, mount, or unmount an NFS datastore, the resulting behavior depends on the version of NFS.

NFS v3 Firewall Behavior

When you add or mount an NFS v3 datastore, ESXi checks the state of the NFS Client (`nfsClient`) firewall rule set.

- If the `nfsClient` rule set is deactivated, ESXi activates the rule set and deactivates the Allow All IP Addresses policy by setting the `allowedAll` flag to `FALSE`. The IP address of the NFS server is added to the allowed list of outgoing IP addresses.
- If the `nfsClient` rule set is activated, the state of the rule set and the allowed IP address policy are not changed. The IP address of the NFS server is added to the allowed list of outgoing IP addresses.

Note: If you manually activate the `nfsClient` rule set or manually set the Allow All IP Addresses policy, either before or after you add an NFS v3 datastore to the system, your settings are overridden when the last NFS v3 datastore is unmounted. The `nfsClient` rule set is deactivated when all NFS v3 datastores are unmounted.

When you remove or unmount an NFS v3 datastore, ESXi performs one of the following actions.

- If none of the remaining NFS v3 datastores are mounted from the server of the datastore being unmounted, ESXi removes the server's IP address from the list of outgoing IP addresses.
- If no mounted NFS v3 datastores remain after the unmount operation, ESXi deactivates the `nfsClient` firewall rule set.

NFS v4.1 Firewall Behavior

When you mount the first NFS v4.1 datastore, ESXi activates the `nfs41client` rule set and sets its `allowedAll` flag to `TRUE`. This action opens port 2049 for all IP addresses. Unmounting an NFS v4.1 datastore does not affect the firewall state. That is, the first NFS v4.1 mount opens port 2049 and that port remains activated unless you close it explicitly.

Using ESXCLI Firewall Commands to Configure ESXi Behavior

If your environment includes multiple ESXi hosts, automate firewall configuration by using ESXCLI commands or the vSphere Web Services SDK.

Firewall Command Reference

You can use the ESXi Shell or ESXCLI commands to configure ESXi at the command line to automate a firewall configuration. To manipulate firewalls and firewall rules, see *Getting Started with ESXCLI* for an introduction, and *ESXCLI Concepts and Examples* for examples of using ESXCLI.

In ESXi 7.0 and later, access to the `service.xml` file, used to create custom firewall rules, is restricted. See VMware Knowledge Base article [2008226](#) for information about creating custom firewall rules using the `/etc/rc.local.d/local.sh` file.

Table 243: Firewall Commands

Command	Description
<code>esxcli network firewall get</code>	Return the status of the firewall and list the default actions.
<code>esxcli network firewall set --default-action</code>	Set to true to set the default action to pass. Set to false to set the default action to drop.
<code>esxcli network firewall set --enabled</code>	Activate or deactivate the ESXi firewall.
<code>esxcli network firewall load</code>	Load the firewall module and the rule set configuration files.
<code>esxcli network firewall refresh</code>	Refresh the firewall configuration by reading the rule set files if the firewall module is loaded.
<code>esxcli network firewall unload</code>	Destroy filters and unload the firewall module.
<code>esxcli network firewall ruleset list</code>	List rule sets information.
<code>esxcli network firewall ruleset set --allowed-all</code>	Set to true to allow all access to all IPs. Set to false to use a list of allowed IP addresses.
<code>esxcli network firewall ruleset set --enabled --ruleset-id=<string></code>	Set enabled to true to activate the specified ruleset. Set enabled to false to deactivate the specified ruleset.
<code>esxcli network firewall ruleset allowedip list</code>	List the allowed IP addresses of the specified rule set.
<code>esxcli network firewall ruleset allowedip add</code>	Allow access to the rule set from the specified IP address or range of IP addresses.
<code>esxcli network firewall ruleset allowedip remove</code>	Remove access to the rule set from the specified IP address or range of IP addresses.
<code>esxcli network firewall ruleset rule list</code>	List the rules of each ruleset in the firewall.

Activate or Deactivate an ESXi Service

You can activate or deactivate ESXi services from the vSphere Client.

Connect to vCenter Server with the vSphere Client.

An ESXi host includes several services that are running by default. If your company policy allows it, you can deactivate services from the security profile, or activate services.

Note: Activating services affects the security of your host. Do not activate a service unless strictly necessary.

After installation, certain services are running by default, while others are stopped. Sometimes, you must perform other steps before a service becomes available in the UI. For example, the NTP service is a way of getting accurate time information, but this service only works when required ports are opened in the firewall.

Available services depend on the VIBs that are installed on the ESXi host. You cannot add services without installing a VIB. Some VMware products, for example, vSphere HA, install VIBs on hosts and make services and the corresponding firewall ports available.

In a default installation, you can modify the status of the following services from the vSphere Client.

Table 244: ESXi Services in the Security Profile

Service	Default	Description
Direct Console UI	Running	The Direct Console User Interface (DCUI) service allows you to interact with an ESXi host from the local console host using text-based menus.
ESXi Shell	Stopped	The ESXi Shell is available from the Direct Console User Interface and includes a set of fully supported commands and a set of commands for troubleshooting and remediation. You must activate access to the ESXi Shell from the direct console of each system. You can activate access to the local ESXi Shell or access to the ESXi Shell with SSH.
SSH	Stopped	The SSH client service on the host that allows remote connections through the Secure Shell.
attestd	Stopped	vSphere Trust Authority Attestation Service.
dpd	Stopped	Data Protection daemon.
Load-Based Teaming Daemon	Running	Load-Based Teaming.
kmxd	Stopped	vSphere Trust Authority Key Provider Service.
Active Directory Service	Stopped	When you configure ESXi for Active Directory, this service is started.
NTP Daemon	Stopped	Network Time Protocol daemon.
PC/SC Smart Card Daemon	Stopped	When you activate the host for smart card authentication, this service starts. See Configuring and Managing Smart Card Authentication for ESXi .
CIM Server	Running	Service that can be used by Common Information Model (CIM) applications.
slpd	Stopped	Service Location Protocol daemon.
SNMP Server	Stopped	SNMP daemon. See the <i>vSphere Monitoring and Performance</i> documentation for information on configuring SNMP v1, v2, and v3.
VDTC Service	Running	vSphere Distributed Tracing Collector service.
vltd	Stopped	VCDR LWD Transport daemon.
Syslog Server	Stopped	Syslog daemon. You can activate the syslog from the Advanced System Settings in the vSphere Client. See the <i>vCenter Server Installation and Setup</i> documentation.
VMware vCenter Agent	Running	vCenter Server agent. Allows a vCenter Server to connect to an ESXi host. Specifically, vpxa is the communication conduit to the host daemon, which in turn communicates with the ESXi kernel.

Service	Default	Description
X.Org Server	Stopped	X.Org Server. This optional feature is used internally for 3D graphics for virtual machines.

1. Browse to an ESXi host in the inventory.
2. Click **Configure**, then click **Services** under **System**.
3. Select the service you want to change.
 - a) Select **Restart**, **Start**, or **Stop** for a one-time change to the host's status.
 - b) To change the status of the host across reboots, click **Edit Startup Policy** and select a policy.
 - **Start and stop with host:** The service starts shortly after the host starts, and closes shortly before the host shuts down. Much like **Start and stop with port usage**, this option means that the service regularly attempts to complete its tasks, such as contacting the specified NTP server. If the port was closed but is later opened, the client begins completing its tasks shortly thereafter.
 - **Start and stop manually:** The host preserves the user-determined service settings, regardless of whether ports are open or not. When a user starts the NTP service, that service is kept running if the host is powered on. If the service is started and the host is powered off, the service is stopped as part of the shutdown process. When the host is powered on, the service is started again, preserving the user-determined state.
 - **Start and stop with port usage:** The default setting for these services. If any port is open, the client attempts to contact the network resources for the service. If some ports are open, but the port for a particular service is closed, the attempt fails. If and when the applicable outgoing port is opened, the service begins completing its startup.

Note: These settings apply only to service settings that are configured through the UI or to applications that are created with the vSphere Web Services SDK. Configurations made through other means, such as from the ESXi Shell or with configuration files, are not affected by these settings.
4. Click **OK**.

Configuring and Managing Lockdown Mode on ESXi Hosts

To increase the security of your ESXi hosts, you can put them in lockdown mode. In lockdown mode, operations must be performed through vCenter Server by default.

You can select normal lockdown mode or strict lockdown mode, which offer different degrees of lockdown. You can also use the Exception User list. Exception users do not lose their privileges when the host enters lockdown mode. Use the Exception User list to add the accounts of third-party solutions and external applications that need to access the host directly when the host is in lockdown mode.

Lockdown Mode Behavior

In lockdown mode, some services are deactivated, and some services are accessible only to certain users.

Lockdown Mode Services Available for Different Users

When the host is running, available services depend on whether lockdown mode is activated, and on the type of lockdown mode.

- In strict and normal lockdown mode, privileged users can access the host through vCenter Server, from the vSphere Client, or by using the vSphere Web Services SDK.

- Direct Console Interface behavior differs for strict lockdown mode and normal lockdown mode.
 - In strict lockdown mode, the Direct Console User Interface (DCUI) service is deactivated.
 - In normal lockdown mode, accounts on the Exception User list can access the DCUI if they have administrator privileges. In addition, all users who are specified in the `DCUI .Access` advanced system setting can access the DCUI.
- If the ESXi Shell or SSH is activated and the host is placed in lockdown mode, accounts on the Exception Users list who have administrator privileges can use these services. For all other users, ESXi Shell or SSH access is deactivated. ESXi or SSH sessions for users who do not have administrator privileges are closed.

All access is logged for both strict and normal lockdown mode.

Table 245: Lockdown Mode Behavior

Service	Normal Mode	Normal Lockdown Mode	Strict Lockdown Mode
vSphere Web Services API	All users, based on permissions	vCenter (vpxuser) Exception users, based on permissions vCloud Director (vsiauser, if available)	vCenter (vpxuser) Exception users, based on permissions vCloud Director (vsiauser, if available)
CIM Providers	Users with administrator privileges on the host	vCenter (vpxuser) Exception users, based on permissions vCloud Director (vsiauser, if available)	vCenter (vpxuser) Exception users, based on permissions vCloud Director (vsiauser, if available)
Direct Console UI (DCUI)	Users with administrator privileges on the host, and users in the <code>DCUI .Access</code> advanced system setting	Users defined in the <code>DCUI .Access</code> advanced system setting Exception users with administrator privileges on the host	DCUI service is stopped.
ESXi Shell (if activated) and SSH (if activated)	Users with administrator privileges on the host	Users defined in the <code>DCUI .Access</code> advanced option Exception users with administrator privileges on the host	Users defined in the <code>DCUI .Access</code> advanced system setting Exception users with administrator privileges on the host

Lockdown Mode Behavior for Users Logged In to the ESXi Shell When Lockdown Mode Is Activated

Users might log in to the ESXi Shell or access the host through SSH before lockdown mode is activated. In that case, users who are on the list of Exception Users and who have administrator privileges on the host remain logged in. The session is closed for all other users. Termination applies to both normal and strict lockdown mode.

How Can You Deactivate Lockdown Mode

You can deactivate lockdown mode as follows.

From the vSphere Client

Users can deactivate both normal lockdown mode and strict lockdown mode from the vSphere Client. See [Deactivate Lockdown Mode from the vSphere Client](#).

From the Direct Console User Interface

Users who can access the Direct Console User Interface on the ESXi host can deactivate normal lockdown mode. In strict lockdown mode, the Direct Console Interface service is stopped.

See [Activate or Deactivate Normal Lockdown Mode from the Direct Console User Interface](#).

Activate Lockdown Mode from the vSphere Client

Select lockdown mode to require that all host configuration changes go through vCenter Server. vSphere supports normal lockdown mode and strict lockdown mode.

If you want to disallow all direct access to a host completely, you can select strict lockdown mode. Strict lockdown mode makes it impossible to access a host if the vCenter Server is unavailable and SSH and the ESXi Shell are deactivated. See [Lockdown Mode Behavior](#).

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Security Profile**.
4. In the Lockdown Mode panel, click **Edit**.
5. Click **Lockdown Mode** and select one of the lockdown mode options.

Option	Description
Normal	The host can be accessed through vCenter Server. Only users who are on the Exception Users list and have administrator privileges can log in to the Direct Console User Interface. If SSH or the ESXi Shell is activated, access might be possible.
Strict	The host can only be accessed through vCenter Server. If SSH or the ESXi Shell is activated, running sessions for accounts in the DCUI .Access advanced system setting and for Exception User accounts that have administrator privileges remain enabled. All other sessions are closed.

6. Click **OK**.

Deactivate Lockdown Mode from the vSphere Client

Deactivate lockdown mode to allow configuration changes from direct connections to the ESXi host. Leaving lockdown mode activated results in a more secure environment.

Users can deactivate both normal lockdown mode and strict lockdown mode from the vSphere Client.

1. Browse to a host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Security Profile**.
4. In the Lockdown Mode panel, click **Edit**.
5. Click **Lockdown Mode** and select **Disabled** to deactivate lockdown mode.
6. Click **OK**.

The system exits lockdown mode, vCenter Server displays an alarm, and an entry is added to the audit log.

Activate or Deactivate Normal Lockdown Mode from the Direct Console User Interface

You can activate and deactivate normal lockdown mode from the Direct Console User Interface (DCUI). You can activate and deactivate strict lockdown mode only from the vSphere Client.

When the host is in normal lockdown mode, the following accounts can access the Direct Console User Interface:

- Accounts in the Exception Users list who have administrator privileges on the host. The Exception Users list is meant for service accounts such as a backup agent.
- Users defined in the `DCUI.Access` advanced option for the host. This option can be used to activate access in a catastrophic failure.

User permissions are preserved when you activate lockdown mode. User permissions are restored when you deactivate lockdown mode from the Direct Console Interface.

Note: If you upgrade a host that is in lockdown mode to ESXi version 6.0 without exiting lockdown mode, and if you exit lockdown mode after the upgrade, all permissions defined before the host entered lockdown mode are lost. The system assigns the administrator role to all users who are found in the `DCUI.Access` advanced option to guarantee that the host remains accessible.

To retain permissions, deactivate lockdown mode for the host from the vSphere Client before the upgrade.

1. At the Direct Console User Interface of the host, press F2 and log in.
2. Scroll to the **Configure Lockdown Mode** setting and press Enter to toggle the current setting.
3. Press Esc until you return to the main menu of the Direct Console User Interface.

Specifying Accounts with Access Privileges in Lockdown Mode

You can specify service accounts that can access the ESXi host directly by adding them to the Exception Users list. You can specify a single user who can access the ESXi host in a catastrophic vCenter Server failure.

What Can Accounts Do When vSphere Is in Lockdown Mode

The vSphere version determines what different accounts can do by default when lockdown mode is activated, and how you can change the default behavior.

- In vSphere 5.0 and earlier, only the root user can log in to the Direct Console User Interface on an ESXi host that is in lockdown mode.
- In vSphere 5.1 and later, you can add a user to the `DCUI.Access` advanced system setting for each host. The setting is meant for a catastrophic failure of vCenter Server. Companies usually lock the password for the user with this access into a safe. A user in the `DCUI.Access` list does not need to have full administrative privileges on the host.
- In vSphere 6.0 and later, the `DCUI.Access` advanced system setting is still supported. In addition, vSphere 6.0 and later supports an Exception User list, which is for service accounts that have to log in to the host directly. Accounts with administrator privileges that are on the Exception Users list can log in to the ESXi Shell. In addition, those users can log in to a host's DCUI in normal lockdown mode and can exit lockdown mode.

You specify Exception Users from the vSphere Client.

Note: Exception users are host local users or Active Directory users with privileges defined locally for the ESXi host. Users that are members of an Active Directory group lose their permissions when the host is in lockdown mode.

Add Users to the DCUI.Access Advanced System Setting

If there is a catastrophic failure, the `DCUI.Access` advanced system setting allows you to exit lockdown mode when you cannot access the host from vCenter Server. You add users to the list by editing the Advanced Settings for the host from the vSphere Client.

Note: Users in the DCUI.Access list can change lockdown mode settings regardless of their privileges. The ability to change lockdown modes can impact the security of your host. For service accounts that need direct access to the host, consider adding users to the Exception Users list instead. Exception users can only perform tasks for which they have privileges. See Specify Lockdown Mode Exception Users later in this topic.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, click **Advanced System Settings**, and click **Edit**.
4. Filter for DCUI.
5. In the **DCUI.Access** text box, enter the local ESXi user names, separated by commas.

Note: You cannot enter Active Directory users. Only local ESXi users are supported.

By default, the root user is included. Consider removing the root user from the DCUI.Access list, and specifying a named account for better auditability.

6. Click **OK**.

Specify Lockdown Mode Exception Users

You can add users to the Exception Users list from the vSphere Client. These users do not lose their permissions when the host enters lockdown mode.

Usually these users are accounts that represent third-party solutions and external applications that need to continue to function in lockdown mode. For example, it makes sense to add service accounts such as a backup agent to the Exception Users list.

Note: The Exception Users list is meant for service accounts that perform very specific tasks, and not for administrators. Adding administrator users to the Exception Users list defeats the purpose of lockdown mode.

Exception users are host local users or Active Directory users with privileges defined locally for the ESXi host. They are not members of an Active Directory group and are not vCenter Server users. These users are allowed to perform operations on the host based on their privileges. That means, for example, that a read-only user cannot deactivate lockdown mode on a host.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Security Profile**.
4. In the Lockdown Mode panel, click **Edit**.
5. Click **Exception Users** and click the **Add User** icon to add exception users.
6. Click **OK**.

Using vSphere Installation Bundles to Perform Secure Updates

Upgrading ESXi with ESXCLI requires an understanding of vSphere Installation Bundles, image profiles, and software depots.

ESXi consists of an image profile, which describes a set of vSphere Installation Bundles (VIBs) that contain the actual software. A VIB is a signed ramdisk representing a component of the system, roughly analogous to an RPM or DEB on a Linux system. An image profile is a collection of VIBs. A software depot is a collection of VIBs and image profiles. ESXi patches and depots contain updated image profiles composed from a common set of VIBs.

You can install ESXi updates on a standalone host using the `esxcli software` commands. For more information, see the *VMware ESXi Upgrade* documentation.

Note: Normally, in a vSphere 7.0 and later environment, you use VMware vSphere® vSphere Lifecycle Manager for lifecycle management of ESXi hosts.

To list all installed VIBs and their current version, or the current image profile, you can use the following ESXCLI commands.

- `esxcli software vib list`
- `esxcli software profile get`

In general, the high-level steps to upgrade ESXi securely are:

- Putting the ESXi host in maintenance mode
- Running an `esxcli software profile update` command, which points to a URL or a ZIP file transferred to the host through SSH
- Restarting the ESXi host

Because VMware cryptographically signs VIBs, secure transfer of VIBs or the entire depot is not necessary, and the update process verifies these signatures.

Manage the Acceptance Levels of ESXi Hosts and vSphere Installation Bundles

The acceptance level of a vSphere Installation Bundle (VIB) depends on the amount of certification of that VIB. The acceptance level of the ESXi host depends on the level of the lowest VIB. If you want to allow lower-level VIBs, you can change the acceptance level of the host. You can remove CommunitySupported VIBs to be able to change the host acceptance level.

VIBs are software packages that include a signature from VMware or a VMware partner. To protect the integrity of the ESXi host, do not allow users to install unsigned (community-supported) VIBs. An unsigned VIB contains code that is not certified by, accepted by, or supported by VMware or its partners. Community-supported VIBs do not have a digital signature.

The acceptance level of the ESXi host must be the same or less restrictive than the acceptance level of any VIB you want to add to the host. For example, if the host acceptance level is VMwareAccepted, you cannot install VIBs at the PartnerSupported level. You can use ESXCLI commands to set an acceptance level for a host. To protect the security and integrity of your ESXi hosts, do not allow unsigned (CommunitySupported) VIBs to be installed on hosts in production systems.

The acceptance level for an ESXi host is displayed in the **Security Profile** in the vSphere Client.

The following acceptance levels are supported.

VMwareCertified

The VMwareCertified acceptance level has the most stringent requirements. VIBs with this level go through thorough testing fully equivalent to VMware in-house Quality Assurance testing for the same technology. Today, only I/O Vendor Program (IOVP) program drivers are published at this level. VMware takes support calls for VIBs with this acceptance level.

VMwareAccepted

VIBs with this acceptance level go through verification testing, but the tests do not fully test every function of the software. The partner runs the tests and VMware verifies the result. Today, CIM providers and PSA plug-ins are among the VIBs published at this level. VMware directs customers with support calls for VIBs with this acceptance level to contact the partner's support organization.

PartnerSupported

VIBs with the PartnerSupported acceptance level are published by a partner that VMware trusts. The partner performs all testing. VMware does not verify the results. This level is used for a new or nonmainstream technology that partners want to enable for VMware systems. Today, driver VIB technologies such as Infiniband, ATAoE, and SSD are at this level with nonstandard hardware drivers. VMware directs customers with support calls for

CommunitySupported

VIBs with this acceptance level to contact the partner's support organization.

The CommunitySupported acceptance level is for VIBs created by individuals or companies outside of VMware partner programs. VIBs at this level have not gone through any VMware-approved testing program and are not supported by VMware Technical Support or by a VMware partner.

1. Connect to each ESXi host by using SSH.
2. Verify that the acceptance level is set to VMwareCertified, VMwareAccepted, or PartnerSupported by running the following command.

```
esxcli software acceptance get
```

3. If the host acceptance level is CommunitySupported, determine whether any of the VIBs are at the CommunitySupported level by running the following commands.

```
esxcli software vib list
esxcli software vib get -n vibname
```

4. Remove any CommunitySupported VIBs by running the following command.

```
esxcli software vib remove --vibname vib
```

5. Change the acceptance level of the host by using one of the following methods.

Option	Description
CLI command	<pre>esxcli software acceptance set --level level</pre> The <code>level</code> parameter is required and specifies the acceptance level to set. Should be one of <code>VMwareCertified</code>, <code>VMwareAccepted</code>, <code>PartnerSupported</code>, or <code>CommunitySupported</code>. See <i>ESXCLI Reference</i> for more information.
vSphere Client	1. Select a host in the inventory. 2. Click Configure. 3. Under System, select Security Profile. 4. Click Edit for Host Image Profile Acceptance Level and choose the acceptance level.

The new acceptance level is in effect.

Note:

ESXi conducts integrity checks of VIBs governed by the Acceptance Level. You can use the `VMkernel.Boot.execInstalledOnly` setting to instruct ESXi to only execute binaries that originate from a valid VIB installed on the host. Combined with Secure Boot, this setting ensures that every single process ever run on an ESXi host is signed, allowed, and expected. By default, the `VMkernel.Boot.execInstalledOnly` setting is deactivated for partner compatibility in vSphere 7.0 and later. Activating this setting when possible improves security. For more information on configuring advanced options for ESXi, see the VMware knowledge base article at <https://kb.vmware.com/s/article/1038578>.

Assigning Privileges for ESXi Hosts

Usually, you give privileges to users by assigning permissions to ESXi host objects that are managed by a vCenter Server system. If you are using a standalone ESXi host, you can assign privileges directly.

Assigning Permissions to ESXi Hosts That Are Managed by vCenter Server

If your ESXi host is managed by a vCenter Server, perform management tasks through the vSphere Client.

You can select the ESXi host object in the vCenter Server object hierarchy and assign the administrator role to a limited number of users. Those users can then perform direct management on the ESXi host. See [Using vCenter Server Roles to Assign Privileges](#).

Best practice is to create at least one named user account, assign it full administrative privileges on the host, and use this account instead of the root account. Set a highly complex password for the root account and limit the use of the root account. Do not remove the root account.

Assigning Permissions to Standalone ESXi Hosts

You can add local users and define custom roles from the Management tab of the VMware Host Client. See the *vSphere Single Host Management - VMware Host Client* documentation.

For all versions of ESXi, you can see the list of predefined users in the `/etc/passwd` file.

The following roles are predefined.

Read Only	Allows a user to view objects associated with the ESXi host but not to make any changes to objects.
Administrator	Administrator role.
No Access	No access. This role is the default role. You can override the default role.

You can manage local users and groups and add local custom roles to an ESXi host using a VMware Host Client connected directly to the ESXi host. See the *vSphere Single Host Management - VMware Host Client* documentation.

In vSphere 6.0 and later, you can use ESXCLI account management commands for managing ESXi local user accounts. You can use ESXCLI permission management commands for setting or removing permissions on both Active Directory accounts (users and groups) and on ESXi local accounts (users only).

Note: If you define a user for the ESXi host by connecting to the host directly, and a user with the same name also exists in vCenter Server, those users are different. If you assign a role to the ESXi user, the vCenter Server user is not assigned the same role.

Predefined ESXi Users and Privileges

If your environment does not include a vCenter Server system, the following users are predefined.

root User	By default each ESXi host has a single root user account with the Administrator role. That root user account can be used for local administration and to connect the host to vCenter Server. Assigning root user privileges can make it easier to break into an ESXi host because the name is already known. Having a common root account also makes it harder to match actions to users. For better auditing, create individual accounts with Administrator privileges. Set a highly complex password for the root account and limit the use of the root account, for example, for use when adding a host to vCenter Server. Do not remove the root account. For more information about assigning permissions to a user for an ESXi host, see <i>vSphere Single Host Management - VMware Host Client</i> documentation. Best practice is to ensure that any account with the Administrator role on an ESXi host is assigned to a specific user with a named
------------------	--

account. Use ESXi Active Directory capabilities, which allow you to manage Active Directory credentials.

Important: You can remove the access privileges for the root user. However, you must first create another permission at the root level that has a different user assigned to the Administrator role.

vCenter Server uses vpxuser privileges when managing activities for the host.

The vCenter Server administrator can perform most of the same tasks on the host as the root user and also schedule tasks, work with templates, and so forth. However, the vCenter Server administrator cannot directly create, delete, or edit local users and groups for hosts. Only a user with Administrator privileges can perform these tasks directly on a host.

You cannot manage the vpxuser user using Active Directory.

CAUTION: Do not change the vpxuser user in any way. Do not change its password. Do not change its permissions. If you do so, you might experience problems when working with hosts through vCenter Server.

The dcui user runs on hosts and acts with Administrator rights. This user's primary purpose is to configure hosts for lockdown mode from the Direct Console User Interface (DCUI).

This user acts as an agent for the direct console and cannot be modified or used by interactive users.

vpxuser User

dcui User

Deactivating Shell Access for non-root ESXi Users

In vSphere 8.0 and later, you can deactivate shell access for non-root ESXi users, such as the predefined vpxuser and dcui users. By deactivating shell access, you can enhance security by enforcing an "API only" stance for these users.

To deactivate shell access, you can use the `esxcli system account set --id user--shell-access false` command. The corresponding API is `LocalAccountManager.updateUser`. You can also use the VMware Host Client to change the Enable Shell Access flag of ESXi local users.

Note: When you deactivate the shell access for a user with administrative access, by virtue of being denied shell access, that user cannot grant shell access to other users, or change the passwords of users that have shell access. Other permissions, such as host profiles, will still allow users such as vpxuser and dcui to change passwords of other users.

When making changes of this kind, verify that they do not break existing third-party workflows.

Using Active Directory to Manage ESXi Users

You can configure ESXi to use a directory service such as Active Directory to manage users.

Creating local user accounts on each host presents challenges with having to synchronize account names and passwords across multiple hosts. Join ESXi hosts to an Active Directory domain to eliminate the need to create and maintain local user accounts. Using Active Directory for user authentication simplifies the ESXi host configuration and reduces the risk for configuration issues that could lead to unauthorized access.

When you use Active Directory, users supply their Active Directory credentials and the domain name of the Active Directory server when adding a host to a domain.

Configure an ESXi Host to Use Active Directory

You can configure an ESXi host to use a directory service such as Active Directory to manage users and groups.

- Verify that you have an Active Directory domain. See your directory server documentation.
- Verify that the host name of ESXi is fully qualified with the domain name of the Active Directory forest.

fully qualified domain name = host_name.domain_name

When you add an ESXi host to Active Directory, the DOMAIN group **ESX Admins** is assigned full administrative access to the host if it exists. If you do not want to make full administrative access available, see VMware Knowledge Base article [1025569](#) for a workaround.

If a host is provisioned with Auto Deploy, Active Directory credentials cannot be stored on the hosts. You can use the vSphere Authentication Proxy to join the host to an Active Directory domain. Because a trust chain exists between the vSphere Authentication Proxy and the host, the Authentication Proxy can join the host to the Active Directory domain. See [Using vSphere Authentication Proxy](#).

Note: When you define user account settings in Active Directory, you can limit the computers that a user can log in to by the computer name. By default, no equivalent restrictions are set on a user account. If you set this limitation, LDAP Bind requests for the user account fail with the message `LDAP binding not successful`, even if the request is from a listed computer. You can avoid this problem by adding the netBIOS name for the Active Directory server to the list of computers that the user account can log in to.

1. Synchronize the time between ESXi and the directory service system.
See [Synchronize ESXi Clocks with a Network Time Server](#) or the VMware Knowledge Base for information about how to synchronize ESXi time with a Microsoft Domain Controller.
2. Ensure that the DNS servers that you configured for the host can resolve the host names for the Active Directory controllers.
 - a) Browse to the host in the vSphere Client inventory.
 - b) Click **Configure**.
 - c) Under Networking, click **TCP/IP configuration**.
 - d) Under TCP/IP Stack: Default, click **DNS** and verify that the host name and DNS server information for the host are correct.

Join the host to a directory service domain. See [Add an ESXi Host to a Directory Service Domain](#). For hosts that are provisioned with Auto Deploy, set up the vSphere Authentication Proxy. See [Using vSphere Authentication Proxy](#). You can configure permissions so that users and groups from the joined Active Directory domain can access the vCenter Server components. For information about managing permissions, see [Add a Permission to an Inventory Object](#).

Add an ESXi Host to a Directory Service Domain

To have your ESXi host use a directory service, you must join the host to the directory service domain.

You can enter the domain name in one of two ways:

- `name.tld` (for example, `domain.com`): The account is created under the default container.
- `name.tld/container/path` (for example, `domain.com/OU1/OU2`): The account is created under a particular organizational unit (OU).

To use the vSphere Authentication Proxy service, see [Using vSphere Authentication Proxy](#).

1. Browse to a host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Authentication Services**.
4. Click **Join Domain**.
5. Enter a domain.

Use the form `name.tld` or `name.tld/container/path`.

6. Enter the user name and password of a directory service user who has permissions to join the host to the domain, and click **OK**.
7. Optional: If you intend to use an authentication proxy, enter the proxy server IP address.
8. Click **OK** to close the Directory Services Configuration dialog box.

You can configure permissions so that users and groups from the joined Active Directory domain can access the vCenter Server components. For information about managing permissions, see [Add a Permission to an Inventory Object](#).

View Directory Service Settings for an ESXi Host

You can view the type of directory server, if any, that the ESXi host uses to authenticate users and the directory server settings.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Authentication Services**.

The Authentication Services page displays the directory service and domain settings.

You can configure permissions so that users and groups from the joined Active Directory domain can access the vCenter Server components. For information about managing permissions, see [Add a Permission to an Inventory Object](#).

Using vSphere Authentication Proxy

You can add ESXi hosts to an Active Directory domain by using vSphere Authentication Proxy instead of adding the hosts explicitly to the Active Directory domain.

You only have to set up the host so it knows about the domain name of the Active Directory server and about the IP address of vSphere Authentication Proxy. When vSphere Authentication Proxy is enabled, it automatically adds hosts that are being provisioned with Auto Deploy to the Active Directory domain. You can also use vSphere Authentication Proxy with hosts that are not provisioned by using Auto Deploy.

Auto Deploy

If you are provisioning hosts with Auto Deploy, you can set up a reference host that points to Authentication Proxy. You then set up a rule that applies the reference host's profile to any ESXi host that is provisioned with Auto Deploy. vSphere Authentication Proxy stores the IP addresses of all hosts that Auto Deploy provisions using PXE in its access control list. When the host boots, it contacts vSphere Authentication Proxy, and vSphere Authentication Proxy joins those hosts, which are already in its access control list, to the Active Directory domain.

Even if you use vSphere Authentication Proxy in an environment that uses certificates that are provisioned by VMCA or third-party certificates, the process works seamlessly if you follow the instructions for using custom certificates with Auto Deploy.

You can set up other hosts to use vSphere Authentication Proxy if you want to make it possible for the host to join the domain without using Active Directory credentials. That means you do not need to transmit Active Directory credentials to the host, and you do not save Active Directory credentials in the host profile.

In that case, you add the host's IP address to the vSphere Authentication Proxy access control list, and vSphere Authentication Proxy authorizes the host based on its IP address by default. You can enable client authentication to have vSphere Authentication Proxy check the host's certificate.

Other ESXi Hosts

Note: IPv6 is supported across vSphere. It is not supported on vSphere with Tanzu.

Start the vSphere Authentication Proxy Service

The vSphere Authentication Proxy service is available on each vCenter Server system. By default, the service is not running. If you want to use vSphere Authentication Proxy in your environment, you can start the service from the vCenter Server Management Interface or from the command line.

Verify that you are using vCenter Server 6.5 or later. In earlier versions of vSphere, vSphere Authentication Proxy is installed separately. See the documentation for the earlier version of the product for instructions.

The vSphere Authentication Proxy service binds to an IPv4 address for communication with vCenter Server, and does not support IPv6. The vCenter Server instance can be on a host machine in an IPv4-only or IPv4/IPv6 mixed-mode network environment. However, when you specify the address of vSphere Authentication Proxy, you must specify an IPv4 address.

1. Start the VMware vSphere Authentication Proxy service.

Option	Description
vCenter Server Management Interface	1. In a Web browser, go to the vCenter Server Management Interface, https://vcenter-IP-address-or-FQDN:5480. 2. Log in as root. The default root password is the password that you set while deploying the vCenter Server. 3. Click Services, and click the VMware vSphere Authentication Proxy service. 4. Click Start. 5. After the service has started, click Set Startup Type and click Automatic to make the startup automatic.
CLI	<code>service-control --start vmcam</code>

2. Confirm that the service started successfully.

You can now set the vSphere Authentication Proxy domain. After that, vSphere Authentication Proxy handles all hosts that are provisioned with Auto Deploy, and you can explicitly add hosts to vSphere Authentication Proxy.

Add a Domain to vSphere Authentication Proxy Using the vSphere Client

You can add a domain to vSphere Authentication Proxy from the vSphere Client.

You can add a domain to vSphere Authentication Proxy only after you enable the proxy. After you add the domain, vSphere Authentication Proxy adds all hosts that you provision with Auto Deploy to that domain. For other hosts, you can also use vSphere Authentication Proxy if you do not want to give those hosts domain privileges.

1. Connect to a system with the .
2. Select the , and click **Configure**.
3. Click **Authentication Proxy** and click **Edit**.
4. Enter the name of the domain that vSphere Authentication Proxy will add hosts to, and the name and password of a user who has Active Directory privileges to add hosts to the domain.
5. Click **Save**.

Add a Domain to vSphere Authentication Proxy Using the camconfig Command

You can add a domain to vSphere Authentication using the `camconfig` command.

You can add a domain to vSphere Authentication Proxy only after you enable the proxy. After you add the domain, vSphere Authentication Proxy adds all hosts that you provision with Auto Deploy to that domain. For other hosts, you can also use vSphere Authentication Proxy if you do not want to give those hosts domain privileges.

1. Log in to the system as a user with administrator privileges.
2. Run the command to enable access to the Bash shell.

```
shell
```

3. Go to the `/usr/lib/vmware-vmcam/bin/` directory where the **camconfig** script is located.
4. To add the domain and user Active Directory credentials to the Authentication Proxy configuration, run the following command.

```
camconfig add-domain -d domain -u user
```

You are prompted for a password.

vSphere Authentication Proxy caches that user name and password. You can remove and recreate the user as needed. The domain must be reachable through DNS, but does not have to be a vCenter Single Sign-On identity source.

vSphere Authentication Proxy uses the user name specified by `user` to create the accounts for ESXi hosts in Active Directory. The user must have privileges to create accounts in the Active Directory domain to which you are adding the hosts. At the time of writing of this information, the Microsoft Knowledge Base article 932455 had background information for account creation privileges.

5. If you later want to remove the domain and user information from vSphere Authentication Proxy, run the following command.

```
camconfig remove-domain -d domain
```

Use vSphere Authentication Proxy to Add a Host to a Domain

The Auto Deploy server adds all hosts that it provisions to vSphere Authentication Proxy, and vSphere Authentication Proxy adds those hosts to the domain. If you want to add other hosts to a domain using vSphere Authentication Proxy, you can add those hosts to vSphere Authentication Proxy explicitly. Afterwards, the vSphere Authentication Proxy server

adds those hosts to the domain. As a result, user-supplied credentials no longer have to be transmitted to the vCenter Server system.

- If the ESXi host is using a VMCA-signed certificate, verify that the host has been added to vCenter Server. Otherwise, the Authentication Proxy service cannot trust the ESXi host.
- If the ESXi host is using a root CA-signed certificate, verify that the appropriate root CA-signed certificate has been added to the vCenter Server system. See [Managing Certificates for ESXi Hosts](#).

You can enter the domain name in one of two ways:

- `name.tld` (for example, `domain.com`): The account is created under the default container.
- `name.tld/container/path` (for example, `domain.com/OU1/OU2`): The account is created under a particular organizational unit (OU).

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under **System**, select **Authentication Services**.
4. Click **Join Domain**.
5. Enter a domain.

Use the form `name.tld`, for example `mydomain.com`, or `name.tld/container/path`, for example, `mydomain.com/organizational_unit1/organizational_unit2`.

6. Select **Using Proxy Server**.
7. Enter the IP address of the Authentication Proxy server, which is always the same as the IP address of the vCenter Server system.
8. Click **OK**.

Activate Client Authentication for vSphere Authentication Proxy

By default, vSphere Authentication Proxy adds any host if it has the IP address of that host in its access control list. For additional security, you can activate client authentication. If client authentication is activated, vSphere Authentication Proxy also checks the certificate of the host.

- Verify that the vCenter Server system trusts the host. By default, when you add a host to vCenter Server, the host is assigned a certificate that is signed by vCenter Server trusted root CA. vSphere Authentication Proxy trusts vCenter Server trusted root CA.
- If you plan on replacing ESXi certificates in your environment, perform the replacement before you activate vSphere Authentication Proxy. The certificates on the ESXi host must match that of the host's registration.

1. Log in to the system as a user with administrator privileges.
2. To activate access to the Bash shell, run the `shell` command.
3. Go to the `/usr/lib/vmware-vmcam/bin/` directory where the **camconfig** script is located.
4. To activate client authentication, run the following command.

```
camconfig ssl-cliAuth -e
```

Going forward, vSphere Authentication Proxy checks the certificate of each host that is added.

5. If you later want to deactivate client authentication again, run the following command.

```
camconfig ssl-cliAuth -n
```

Import the vSphere Authentication Proxy Certificate to ESXi Host

By default, ESXi hosts require explicit verification of the vSphere Authentication Proxy certificate. If you are using vSphere Auto Deploy, the Auto Deploy service takes care of adding the certificate to hosts that it provisions. For other hosts, you must add the certificate explicitly.

- Upload the vSphere Authentication Proxy certificate to a datastore accessible to the ESXi host. Using an SFTP application such as WinSCP, you can download the certificate from the vCenter Server host at the following location.
`/var/lib/vmware/vmcam/ssl/rui.crt`
 - Verify that the `UserVars.ActiveDirectoryVerifyCAMCertificateESXi` advanced setting is set to 1 (the default).
1. Select the ESXi host and click **Configure**.
 2. Under **System**, select **Authentication Services**.
 3. Click **Import Certificate**.
 4. Enter the certificate file path following the format `[datastore]/path/certname.crt`.
 5. Enter the IP address of vSphere Authentication Proxy server.
 6. Click **OK**.

Generate a New Certificate for vSphere Authentication Proxy

You can generate a new certificate that is provisioned by VMware Certificate Authority (VMCA), or a new certificate that includes VMCA as a subordinate certificate.

You must have root or Administrator privileges on the system on which vSphere Authentication Proxy is running.

See [Set Up vSphere Authentication Proxy to Use Custom Certificates](#) if you want to use a custom certificate that is signed by a third-party or enterprise CA.

1. Make a copy of `certool.cfg`.

```
cp /usr/lib/vmware-vmca/share/config/certool.cfg /var/lib/vmware/vmcam/ssl/vmcam.cfg
```

2. Edit the copy with some information about your organization, as in the following example.

```
Country = IE
Name = vmcam
Organization = VMware
OrgUnit = vTSU
State = Cork
Locality = Cork
Hostname = test-cam-1.test1.vmware.com
```

3. Generate the new private key in `/var/lib/vmware/vmcam/ssl/`.

```
/usr/lib/vmware-vmca/bin/certool --genkey --privkey=/var/lib/vmware/vmcam/ssl/rui.key --pubkey=/tmp/vmcam.pub --server=localhost
```

For `localhost`, supply the FQDN of the vCenter Server.

4. Generate the new certificate in `/var/lib/vmware/vmcam/ssl/` using the key and `vmcam.cfg` file that you created in Step 1 and Step 2.

```
/usr/lib/vmware-vmca/bin/certool --server=localhost --gencert --privkey=/var/lib/vmware/vmcam/ssl/rui.key --cert=/var/lib/vmware/vmcam/ssl/rui.crt --config=/var/lib/vmware/vmcam/ssl/vmcam.cfg
```

For `localhost`, supply the FQDN of the vCenter Server.

Set Up vSphere Authentication Proxy to Use Custom Certificates

Using custom certificates with vSphere Authentication Proxy consists of several steps. First you generate a CSR and send it to your CA for signing. Then you place the signed certificate and key file in a location that vSphere Authentication Proxy can access.

By default, vSphere Authentication Proxy generates a CSR during first boot and asks VMCA to sign that CSR. vSphere Authentication Proxy registers with vCenter Server using that certificate. You can use custom certificates in your environment, if you add those certificates to vCenter Server.

1. Generate a CSR for vSphere Authentication Proxy.

- a) Create a configuration file, `/var/lib/vmware/vmcam/ssl/vmcam.cfg`, as in the following example.

```
[ req ]
distinguished_name = req_distinguished_name
encrypt_key = no
prompt = no
string_mask = nombstr
req_extensions = v3_req
[ v3_req ]
basicConstraints = CA:false
keyUsage = nonRepudiation, digitalSignature, keyEncipherment
subjectAltName = DNS:vcenter1.example.com
[ req_distinguished_name ]
countryName = US
stateOrProvinceName = NY
localityName = New York
0.organizationName = Example Inc.
organizationalUnitName = IT Org
commonName = vcenter1.example.com
```

Note the following:

- **subjectAltName:** Use the format `DNS:FQDN_of_vCenter_Appliance_to_use_the_CA-signed_certificate`.
- **commonName:** Use the same FQDN of the vCenter Appliance used in `subjectAltName`.

- b) Run `openssl` to generate a CSR file and a key file, passing in the configuration file.

```
openssl req -new -nodes -out vmcam.csr -newkey rsa:2048 -keyout /var/lib/vmware/vmcam/ssl/rui.key -config /var/lib/vmware/vmcam/ssl/vmcam.cfg
```

2. Back up the `rui.crt` certificate and `rui.key` files, which are stored in the following location.

`/var/lib/vmware/vmcam/ssl/rui.crt`

3. Unregister vSphere Authentication Proxy.

- a) Go to the `/usr/lib/vmware-vmcam/bin` directory where the `camregister` script is located.
b) Run the following command.

```
camregister --unregister -aVC_address -uuser
```

user must be a vCenter Single Sign-On user that has administrator permissions on vCenter Server.

4. Stop the vSphere Authentication Proxy service.

Tool	Steps
vCenter Server Configuration Management Interface	1. In a Web browser, go to the vCenter Server Configuration Management Interface, <code>https://vcenter-IP-address-or-FQDN:5480</code> .

Tool	Steps
	- Log in as root. The default root password is the password that you set while deploying the vCenter Server. - Click Services, and click VMware vSphere Authentication Proxy. - Click Stop.
CLI	<code>service-control --stop vmcam</code>

- Replace the existing `ruil.crt` certificate and `ruil.key` files with the files that you received from your CA.
- Restart the vSphere Authentication Proxy service.
- Reregister vSphere Authentication Proxy explicitly with vCenter Server by using the new certificate and key.

```
camregister --register -a VC_address -u user -cfull_path_to_ruil.crt-kfull_path_to_ruil.key
```

Configuring and Managing Smart Card Authentication for ESXi

You can use smart card authentication to log in to the ESXi Direct Console User Interface (DCUI) by using a Personal Identity Verification (PIV), Common Access Card (CAC) or SC650 smart card instead specifying a user name and password.

A smart card is a small plastic card with an embedded integrated circuit chip. Many government agencies and large enterprises use smart card based two-factor authentication to increase the security of their systems and comply with security regulations.

When smart card authentication is enabled on an ESXi host, the DCUI prompts for a smart card and PIN combination instead of the default prompt for a user name and password.

- When you insert the smart card into the smart card reader, the ESXi host reads the credentials on it.
- The ESXi DCUI displays your login ID, and prompts for your PIN.
- After you enter your PIN, the ESXi host matches it with the PIN stored on the smart card and verifies the certificate on the smart card with Active Directory.
- After successful verification of the smart card certificate, ESXi logs you in to the DCUI.

You can switch to user name and password authentication from the DCUI by pressing F3.

The chip on the smart card locks after a few consecutive incorrect PIN entries, usually three. If a smart card is locked, only selected personnel can unlock it.

Activate Smart Card Authentication

Activate smart card authentication to prompt for smart card and PIN combination to log in to the ESXi DCUI.

- Set up the infrastructure to handle smart card authentication, such as accounts in the Active Directory domain, smart card readers, and smart cards.
- Configure ESXi to join an Active Directory domain that supports smart card authentication. For more information, see [Using Active Directory to Manage ESXi Users](#).
- Use the vSphere Client to add root certificates. See [Managing Certificates for ESXi Hosts](#).

- Browse to the host in the vSphere Client inventory.
- Click **Configure**.
- Under System, select **Authentication Services**.

You see the current smart card authentication status and a list with imported certificates.

4. In the Smart Card Authentication panel, click **Edit**.
5. In the Edit Smart Card Authentication dialog box, select the Certificates page.
6. Add trusted Certificate Authority (CA) certificates, for example, root and intermediary CA certificates.
Certificates must be in PEM format.
7. Open the Smart Card Authentication page, select the **Enable Smart Card Authentication** check box, and click **OK**.

Deactivate Smart Card Authentication

Deactivate smart card authentication to return to the default user name and password authentication for ESXi DCUI login.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Authentication Services**.
You see the current smart card authentication status and a list with imported certificates.
4. In the Smart Card Authentication panel, click **Edit**.
5. On the Smart Card Authentication page, deselect the **Enable Smart Card Authentication** check box, and click **OK**.

Authenticating With User Name and Password in Case of Connectivity Problems

If the Active Directory (AD) domain server is not reachable, you can log in to the ESXi DCUI by using user name and password authentication to perform emergency actions on the host.

In exceptional circumstances, the AD domain server is not reachable to authenticate the user credentials on the smart card because of connectivity problems, network outage, or disasters. In that case, you can log in to the ESXi DCUI by using the credentials of a local ESXi Administrator user. After logging in, you can perform diagnostics or other emergency actions. The fallback to user name and password login is logged. When the connectivity to AD is restored, smart card authentication is enabled again.

Note: Loss of network connectivity to vCenter Server does not affect smart card authentication if the Active Directory (AD) domain server is available.

Using Smart Card Authentication in Lockdown Mode

When activated, lockdown mode on the ESXi host increases the security of the host and limits access to the DCUI. Lockdown mode might cause the smart card authentication to no longer work.

In normal lockdown mode, only users on the Exception Users list with administrator privileges can access the DCUI. Exception users are host local users or Active Directory users with permissions defined locally for the ESXi host. If you want to use smart card authentication in normal lockdown mode, you must add users to the Exception Users list from the vSphere Client. These users do not lose their permissions when the host enters normal lockdown mode and can log in to the DCUI. For more information, see [Specify Lockdown Mode Exception Users](#).

In strict lockdown mode, the DCUI service is stopped. As a result, you cannot access the host by using smart card authentication.

Using the ESXi Shell

The ESXi Shell provides essential maintenance commands and is deactivated by default on ESXi hosts. You can activate local and remote access to the shell if necessary. To reduce the risk of unauthorized access, activate the ESXi Shell for troubleshooting only.

The ESXi Shell is independent of lockdown mode. Even if the host is running in lockdown mode, you can still log in to the ESXi Shell if it is activated.

The applicable services are as follows.

ESXi Shell
SSH

Activate this service to access the ESXi Shell locally.

Activate this service to access the ESXi Shell remotely by using SSH.

The root user and users with the Administrator role can access the ESXi Shell. Users who are in the Active Directory group ESX Admins are automatically assigned the Administrator role. By default, only the root user can run system commands (such as `vmware -v`) by using the ESXi Shell.

Note: Do not activate the ESXi Shell unless you actually need access.

Set Idle Timeout for the ESXi Shell Using the vSphere Client

If you enable the ESXi Shell on a host, but forget to log out of the session, the idle session remains connected indefinitely. The open connection increases the potential for someone to gain privileged access to the host. Prevent this by setting a timeout for idle sessions.

The idle timeout is the amount of time that can elapse before a user is logged out of an idle interactive session. You can control the amount of time for both local and remote (SSH) sessions from the Direct Console Interface (DCUI) or from the vSphere Client.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Advanced System Settings**.
4. Click **Edit**, select `UserVars.ESXiShellInteractiveTimeOut`, and enter the timeout setting.
A value of zero (0) deactivates the idle time.
5. Restart the ESXi Shell service and the SSH service for the timeout to take effect.
 - a) Go to **System > Services**.
 - b) One by one, select ESXi Shell and SSH, and click **Restart**.

If the session is idle, users are logged out after the timeout period elapses.

Set Availability Timeout for the ESXi Shell Using the vSphere Client

The ESXi Shell is deactivated by default. You can set an availability timeout for the ESXi Shell to increase security when you activate the shell.

The availability timeout setting is the amount of time that can elapse before you must log in after the shell is activated. After the timeout period, the service is deactivated and users are not allowed to log in.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, select **Advanced System Settings**.
4. Click **Edit**, and select `UserVars.ESXiShellTimeOut`.
5. Enter the idle timeout setting.
6. Click **OK**.
7. Restart the ESXi Shell service and the SSH service for the timeout to take effect.
 - a) Go to **System > Services**.
 - b) One by one, select ESXi Shell and SSH, and click **Restart**.

If you are logged in when the timeout period elapses, your session will persist. However, after you log out or your session is terminated, users are not allowed to log in.

Set Availability Timeout or Idle Timeout for the ESXi Shell Using the DCUI

The ESXi Shell is deactivated by default. To increase security when you activate the shell, you can set an availability timeout, an idle timeout, or both.

Activate the ESXi Shell. See [Activate Access to the ESXi Shell Using the DCUI](#).

The two types of timeout apply in different situations.

ESXi Shell Idle Timeout

If a user activates the ESXi Shell on a host, but forgets to log out of the session, the idle session remains connected indefinitely. The open connection can increase the potential for someone to gain privileged access to the host. You can prevent this situation by setting a timeout for idle sessions.

ESXi Shell Availability Timeout

The availability timeout determines how much time can elapse before you log in after you initially activate the shell. If you wait longer, the service is deactivated and you cannot log in to the ESXi Shell.

1. Log in to the ESXi Shell.
 2. From the Troubleshooting Mode Options menu, select **Modify ESXi Shell and SSH timeouts** and press Enter.
 3. Enter the idle timeout (in seconds) or the availability timeout.
 4. Press Enter and press Esc until you return to the main menu of the Direct Console User Interface.
 5. Click **OK**.
 6. Restart the ESXi Shell service and the SSH service for the timeout to take effect.
 - a) In the vSphere Client, select the host, and go to **Configure > System > Services**.
 - b) One by one, select ESXi Shell and SSH, and click **Restart**.
- If you set the idle timeout, users are logged out after the session is idle for the specified time.

- If you set the availability timeout, and you do not log in before that timeout elapses, logins become deactivated again.

Activate Access to the ESXi Shell Using the vSphere Client

ESXi Shell and SSH interfaces are deactivated by default. Keep these interfaces deactivated unless you are performing troubleshooting or support activities. For everyday activities, use the vSphere Client, where activity is subject to role-based access control and modern access control methods.

If you want to use an authorized SSH key, you can upload it. See [ESXi SSH Keys](#).

Note: Access the host by using the vSphere Client, remote command-line tools (ESXCLI and PowerCLI), and published APIs. Do not activate remote access to the host using SSH unless special circumstances require it.

1. Browse to the host in the inventory.
2. Click **Configure**, then click **Services** under System.
3. Manage the ESXi, SSH, or Direct Console UI services.
 - a) In the Services pane, select the service.
 - b) Click **Edit Startup Policy** and select the startup policy **Start and stop manually**.
 - c) To activate the service, click **Start**.

When you select **Start and stop manually**, the service does not start when you reboot the host. If you want the service to start when you reboot the host, select **Start and stop with host**.

Set the availability and idle timeouts for the ESXi Shell. See [Set Availability Timeout for the ESXi Shell Using the vSphere Client](#) and [Set Idle Timeout for the ESXi Shell Using the vSphere Client](#).

Activate Access to the ESXi Shell Using the DCUI

The Direct Console User Interface (DCUI) allows you to interact with the host locally using text-based menus. Evaluate whether the security requirements of your environment support activating the Direct Console User Interface.

You can use the Direct Console User Interface (DCUI) to activate local and remote access to the ESXi Shell. You access the Direct Console User Interface from the physical console attached to the host. After the host reboots and loads ESXi, press F2 to log in to the DCUI. Enter the credentials that you created when you installed ESXi.

Note: Changes made to the host using the Direct Console User Interface, the vSphere Client, ESXCLI, or other administrative tools are committed to permanent storage every hour or upon graceful shutdown. If the host fails before the changes are committed, they might be lost.

1. From the Direct Console User Interface, press F2 to access the System Customization menu.
2. Select **Troubleshooting Options** and press Enter.
3. From the Troubleshooting Mode Options menu, select a service to activate.
 - Enable
 - Enable SSH
4. Press Enter to activate the service.
5. Press Esc until you return to the main menu of the Direct Console User Interface.

Set the availability and idle timeouts for the ESXi Shell. See [Set Availability Timeout or Idle Timeout for the ESXi Shell Using the DCUI](#).

Log in to the ESXi Shell for Troubleshooting

Perform ESXi configuration tasks with the vSphere Client, ESXCLI, or VMware PowerCLI. Log in to the ESXi Shell (formerly Tech Support Mode or TSM) for troubleshooting purposes only.

1. Log in to the ESXi Shell using one of the following methods.
 - If you have direct access to the host, press Alt+F1 to open the login page on the machine's physical console.
 - If you are connecting to the host remotely, use SSH or another remote console connection to start a session on the host.
2. Enter a user name and password recognized by the host.

UEFI Secure Boot for ESXi Hosts


Secure boot is part of the UEFI firmware standard. With secure boot in use, a machine refuses to load any UEFI driver or app unless the operating system bootloader is cryptographically signed. In vSphere 6.5 and later, ESXi supports secure boot if it is enabled in the hardware.

How ESXi Uses UEFI Secure Boot

ESXi version 6.5 and later supports UEFI Secure Boot at each level of the boot stack.

Note: Before you use UEFI Secure Boot on a host that was upgraded, check for compatibility by following the instructions in [Run the Secure Boot Validation Script After ESXi Upgrade](#).

Figure 117: UEFI Secure Boot



With secure boot in use, the boot sequence proceeds as follows.

1. In vSphere 6.5 and later, the ESXi bootloader contains a VMware public key. The bootloader uses this key to verify the signature of the kernel and a small subset of the system that includes a secure boot VIB verifier.
2. The VIB verifier verifies every VIB package that is installed on the system.

At this point, the entire system boots with the root of trust in certificates that are part of the UEFI firmware.

Note: When you install or upgrade to vSphere 7.0 Update 2 or later, and an ESXi host has a TPM, the TPM seals the sensitive information by using a TPM policy based on PCR values for UEFI Secure Boot. This value is loaded during subsequent reboots if the policy is satisfied as true. To deactivate or activate UEFI Secure Boot in vSphere 7.0 Update 2 and later, see [Activate or Deactivate the Secure Boot Enforcement for a Secure ESXi Configuration](#).

Troubleshooting UEFI Secure Boot

If secure boot does not succeed at any level of the boot sequence, an error results.

The error message depends on the hardware vendor and on the level at which verification did not succeed.

- If you attempt to boot with a bootloader that is unsigned or has been tampered with, an error during the boot sequence results. The exact message depends on the hardware vendor. It might look like the following error, but might look different.

```
UEFI0073: Unable to boot PXE Device...because of the Secure Boot policy
```

- If the kernel has been tampered with, an error like the following results.

```
Fatal error: 39 (Secure Boot Failed)
```

- If a package (VIB or driver) has been tampered with, a purple screen with the following message appears.

```
UEFI Secure Boot failed:  
Failed to verify signatures of the following vib(s) (XX)
```

To resolve issues with secure boot, follow these steps.

1. Reboot the host with secure boot deactivated.
2. Run the secure boot verification script (see [Run the Secure Boot Validation Script After ESXi Upgrade](#)).
3. Examine the information in the `/var/log/esxupdate.log` file.

Securing ESXi Hosts with Trusted Platform Module

ESXi hosts can use Trusted Platform Modules (TPM) chips, which are secure cryptoprocessors that enhance host security by providing a trust assurance rooted in hardware as opposed to software.

What Is a TPM

TPM is an industry-wide standard for secure cryptoprocessors. Today, TPM chips are found in most computers, from laptops, to desktops, to servers. vSphere 6.7 and later supports TPM version 2.0.

A TPM 2.0 chip attests to an ESXi identity of a host. Host attestation is the process of authenticating and attesting to the state of the software on a host at a given point in time. UEFI secure boot, which ensures that only signed software is loaded at boot time, is a requirement for successful attestation. The TPM 2.0 chip records and securely stores measurements of the software modules booted in the system, which vCenter Server remotely verifies.

The high-level steps of the remote attestation process are:

1. Establish the trustworthiness of the remote TPM and create an Attestation Key (AK) on it.
When an ESXi host is added to, rebooted from, or reconnected to vCenter Server, vCenter Server requests an AK from the host. Part of the AK creation process also involves the verification of the TPM hardware itself, to ensure that a known (and trusted) vendor has produced it.
2. Retrieve the Attestation Report from the host.
vCenter Server requests that the host sends an Attestation Report, which contains a quote of Platform Configuration Registers (PCRs), signed by the TPM, and other signed host binary metadata. By checking that the information corresponds to a configuration it deems trusted, a vCenter Server identifies the platform on a previously untrusted host.

3. Verify the authenticity of the host.

vCenter Server verifies the authenticity of the signed quote, infers the software versions, and determines the trustworthiness of said software versions. If vCenter Server determines the signed quote is invalid, remote attestation fails and the host is not trusted.

What Are the vSphere Requirements to Use a TPM

To use a TPM 2.0 chip, your vCenter Server environment must meet these requirements:

- vCenter Server 6.7 or later
- ESXi 6.7 host or later with TPM 2.0 chip installed and enabled in UEFI
- UEFI Secure Boot enabled

Ensure that the TPM is configured in the BIOS of the ESXi host to use the SHA-256 hashing algorithm and the TIS/FIFO (First-In, First-Out) interface and not CRB (Command Response Buffer). For information about setting these required BIOS options, refer to the vendor documentation.

Review the TPM 2.0 chips certified by VMware at the following location:

<https://www.vmware.com/resources/compatibility/search.php>

What Happens When You Boot a Host with a TPM

When you boot an ESXi host with an installed TPM 2.0 chip, vCenter Server monitors the attestation status of the host. To view the hardware trust status, in the vSphere Client, select the vCenter Server, then the **Summary** tab under **Monitor**. The hardware trust status is one of the following:

- Green: Normal status, indicating full trust.
- Red: Attestation failed.

Note: If you add a TPM 2.0 chip to an ESXi host that vCenter Server already manages, you must first disconnect the host, then reconnect it. See the *vCenter Server and Host Management* documentation for information about disconnecting and reconnecting hosts.

With vSphere 7.0 and later, VMware® vSphere Trust Authority™ uses remote attestation capabilities for ESXi hosts. See [What Is the vSphere Trust Authority Attestation Service](#).

View ESXi Host Attestation Status

When added to an ESXi host, a Trusted Platform Module 2.0 compatible chip attests the integrity of the platform. You can view the attestation status of the host in the vSphere Client. You can also view the Intel Trusted Execution Technology (TXT) status.

1. Connect to vCenter Server by using the vSphere Client.
2. Navigate to a data center and click the **Monitor** tab.
3. Click **Security**.
4. Review the host's status in the Attestation column and read the accompanying message in the **Message** column.
5. If this host is a Trusted Host, see [View the Trusted Cluster Attestation Status](#) for more information.

For a Failed or Warning attestation status, see [Troubleshoot ESXi Host Attestation Problems](#). For Trusted Hosts, see [Troubleshoot Trusted Host Attestation Problems](#).

Troubleshoot ESXi Host Attestation Problems

When you install a Trusted Platform Module (TPM) device on an ESXi host, the host might fail to pass attestation. You can troubleshoot the potential causes of this problem.

1. View the ESXi host alarm status and accompanying error message. See [View ESXi Host Attestation Status](#).
2. If the error message is `Host secure boot was disabled`, you must re-enable secure boot to resolve the problem.
3. If the attestation status of the host is failed, check the vCenter Server `vpd.log` file for the following message:

```
No cached identity key, loading from DB
```

This message indicates that you are adding a TPM 2.0 chip to an ESXi host that vCenter Server already manages. You must first disconnect the host, then reconnect it. See *vCenter Server and Host Management* documentation for information about disconnecting and reconnecting hosts.

For more information about vCenter Server log files, including location and log rotation, see the VMware knowledge base article at <https://kb.vmware.com/s/article/1021804>.

4. For all other error messages, contact Customer Support.

ESXi Log Files

Log files are an important component of troubleshooting attacks and obtaining information about breaches. All ESXi hosts run a syslog service, which logs messages from the VMkernel and other system components to local files or to a remote host.

To increase the security of the host, take the following measures.

- Configure persistent logging to a datastore. By default, the logs on ESXi hosts are stored in the in-memory file system. Therefore, they are lost when you reboot the host, and only 24 hours of log data is stored. When you enable persistent logging, you have a dedicated activity record for the host.
- Remote logging to a central host allows you to gather log files on a central host. From that host, you can monitor all hosts with a single tool, do aggregate analysis, and search log data. This approach facilitates monitoring and reveals information about coordinated attacks on multiple hosts.
- Configure the remote secure syslog on ESXi hosts by using ESXCLI or PowerCLI, or by using an API client.
- Query the syslog configuration to make sure that the syslog server and port are valid.

See the *vSphere Monitoring and Performance* documentation for information about syslog setup, and for additional information on ESXi log files.

Configure Syslog on ESXi Hosts

You can use the vSphere Client, the VMware Host Client, or the `esxcli system syslog` command to configure the syslog service.

The syslog service receives, categorizes, and stores log messages for analyses that help you take preventive action in your environment.

ESXi Syslog Options

You can define the behavior of ESXi syslog files and transmissions by using a set of syslog options.

Apart from the base settings, such as `Syslog.global.logHost`, starting from ESXi 7.0 Update 1, a list of advanced options is available for customizations, and NIAP compliance.

Note: Always configure persistent storage before you set any of the audit record parameters or the `Syslog.global.logDir` parameter.

Note: All audit record settings, beginning with `Syslog.global.auditRecord`, take effect immediately. However, for other settings that you define by using ESXCLI, make sure to run the `esxcli system syslog reload` command to enable the changes.

Table 246: Legacy Syslog Options

Option	ESXCLI command	Description
<code>Syslog.global.logHost</code>	<code>esxcli system syslog config set --loghost=<str></code>	Defines a comma-delimited list of remote hosts and specifications for message transmissions. If the <code>loghost=<str></code> field is blank, no logs are forwarded. While no hard limit to the number of remote hosts to receive syslog messages exists, good practice is to keep the number of remote hosts to five or less. The format of a remote host specification is: <code>protocol://hostname ipv4 '['ipv6']':port[?formatter=value[&framing=value]]</code> . The protocol must be one of <code>tcp</code> , <code>udp</code> , or <code>ssl</code> . The value of a port can be any decimal number between 1 and 65535. If a port is not provided, UDP and TCP use 514; SSL uses 1514. For example: <code>ssl://hostname1:1514</code> . The formatter must be <code>RFC_3164</code> or <code>RFC_5424</code> ; <code>RFC_3164</code> is the default. The framing must be <code>non_transparent</code> or <code>octet_counting</code> ; the default is <code>non_transparent</code> . For more information, see Fine-tune Syslog on ESXi Hosts .
<code>Syslog.global.defaultRotate</code>	<code>esxcli system syslog config set --default-rotate=<long></code>	Maximum number of old log files to keep. You can set this number globally and for individual subloggers (see <code>Syslog.global.defaultSize</code>).
<code>Syslog.global.defaultSize</code>	<code>esxcli system syslog config set --default-size=<long></code>	Default size of log files, in KiB. After a file reaches the default size, the syslog service creates a new file. You can set this number globally and for individual subloggers.
<code>Syslog.global.logDir</code>	<code>esxcli system syslog config set --logdir=<str></code>	Directory where logs reside. The directory can be on mounted NFS or VMFS volumes. Only the <code>/scratch</code> directory on the local file system is persistent across reboots. Specify the directory as <code>[datastorename]path_to_file</code> , where the path is

Option	ESXCLI command	Description
		relative to the root of the volume backing the datastore. For example, the path [storage1] /systemlogs maps to the path /vmfs/volumes/storage1/systemlogs.
Syslog.global.logDirUnique	esxcli system syslog config set --logdir-unique=<bool>	Specifies the ESXi host name to be concatenated to the value of Syslog.global.logDir. It is critical that you enable this setting when multiple ESXi hosts log to a shared file system. Selecting this option creates a subdirectory with the name of the ESXi host under the directory specified by Syslog.global.LogDir . A unique directory is useful if the same NFS directory is used by multiple ESXi hosts.
Syslog.global.certificate.checkSSL	esxcli system syslog config set --check-ssl-certs=<bool>	Enforces checking of SSL certificates when transmitting messages to remote hosts.

Table 247: Syslog Options Available Starting from ESXi 7.0 Update 1

Option	ESXCLI command	Description
Syslog.global.auditRecord.storageCapacity	esxcli system auditrecords local set --size=<long>	Specifies the capacity of the audit record storage directory located on the ESXi host, in MiB. You cannot decrease the capacity of the audit record storage. You can increase the capacity before or after the audit record storage is enabled (see Syslog.global.auditRecord.storageEnable).
Syslog.global.auditRecord.remoteEnable	esxcli system auditrecords remote enable	Enables sending audit records to remote hosts. Remote hosts are specified by using the Syslog.global.logHost parameter.
Syslog.global.auditRecord.storageDirectory	esxcli system auditrecords local set --directory=<dir>	Creates an audit record storage directory and unless specified, sets /scratch/auditLog as the default location. You must not manually create an audit record storage directory and you cannot change the audit record storage directory while audit record storage is enabled (see Syslog.global.auditRecord.storageEnable).
Syslog.global.auditRecord.storageEnable	esxcli system auditrecords local enable	Enables the storage of audit records on an ESXi host. If the audit record storage directory does not exist, it is created with the capacity specified by Syslog.global.auditRecord.storageCapacity.
Syslog.global.certificate.checkCRL	esxcli system syslog config set --crl-check=<bool>	Enables checking the revocation status of all the certificates in an SSL certificate chain.

Option	ESXCLI command	Description
		Enables verification of X.509 CRLs, which are not checked by default in compliance with industry conventions. A NIAP-validated configuration requires CRL checks. Due to implementation limitations, if CRL checks are enabled, then all certificates in a certificate chain must provide a CRL link. Do not enable the <code>crl-check</code> option for installations not related to certification, because of the difficulty in properly configuring an environment that uses CRL checks.
<code>Syslog.global.certificate.strictX509Compliance</code>	<code>esxcli system syslog config set --x509-strict=<bool></code>	Enables strict compliance with X.509. Performs additional validity checks on CA root certificates during verification. These checks are generally not performed, as CA roots are inherently trusted, and might cause incompatibilities with existing, misconfigured CA roots. A NIAP-validated configuration requires even CA roots to pass validations. Do not enable the <code>x509-strict</code> option for installations not related to certification, because of the difficulty in properly configuring an environment that uses CRL checks.
<code>Syslog.global.droppedMsgs.fileRotate</code>	<code>esxcli system syslog config set --drop-log-rotate=<long></code>	Specifies the number of old dropped message log files to keep.
<code>Syslog.global.droppedMsgs.fileSize</code>	<code>esxcli system syslog config set --drop-log-size=<long></code>	Specifies the size of each dropped message log file before switching to a new one, in KiB.
<code>Syslog.global.logCheckSSLCerts</code>	<code>esxcli system syslog config set --check-ssl-certs=<bool></code>	Enforces checking of SSL certificates when transmitting messages to remote hosts. Note: Deprecated. Use <code>Syslog.global.certificate.checkSSLCerts</code> in ESXi 7.0 Update 1 and later.
<code>Syslog.global.logFilters</code>	<code>esxcli system syslog config logfilter [add remove set] ...</code>	Specifies one or more log filtering specifications. Each log filter must be separated by a double vertical bar " ". The format of a log filter is: <code>numLogs ident logRegexp.numLogs</code> sets the maximum number of log entries for the specified log messages. After reaching this number, the specified log messages are filtered and ignored. <code>ident</code> specifies one or more system components to apply the filter to the log messages that these components generate. <code>logRegexp</code> specifies a case-sensitive phrase with Python regular expression syntax to filter the log messages by their content.

Option	ESXCLI command	Description
<code>Syslog.global.logFiltersEnable</code>		Enables the use of log filters.
<code>Syslog.global.logLevel</code>	<code>esxcli system syslog config set --log-level=<str></code>	Specifies the log filtering level. You must change this parameter only when troubleshooting an issue with the syslog daemon. You can use the values <code>debug</code> for the most detailed level, <code>info</code> for the default detail level, <code>warning</code> for only warnings or errors, or <code>error</code> , only for errors.
<code>Syslog.global.msgQueueDropMark</code>	<code>esxcli system syslog config --queue-drop-mark=<long></code>	Specifies the percent of the message queue capacity at which messages are dropped.
<code>Syslog.global.remoteHost.connectRetryDelay</code>	<code>esxcli system syslog config set --default-timeout=<long></code>	Specifies the delay before retrying to connect to a remote host after a connection attempt fails, in seconds.
<code>Syslog.global.remoteHost.maxMsgLen</code>	<code>esxcli system syslog config set --remote-host-max-msg-len=<long></code>	For the TCP and SSL protocols, this parameter specifies the maximum length of a syslog transmission before truncation occurs, in bytes. The default maximum length for remote host messages is 1 KiB. You can increase the maximum message length to up to 16 KiB. However, raising this value above 1 KiB does not ensure that long transmissions arrive untruncated to a syslog collector. For example, when the syslog infrastructure that issues a message is external to ESXi. This setting does not affect the UDP protocol. RFC 5426 specifies the UDP message lengths that can be safely accepted at 480 bytes for IPV4 and 1180 bytes for IPV6. Because of these restrictions, and because the networking infrastructure can arbitrary drop UDP packets, the use of UDP for transmitting critical syslog messages is not recommended.
<code>Syslog.global.vsanBacking</code>	<code>esxcli system syslog config set --vsan-backing=<bool></code>	Allows log files and the audit record storage directory to be placed on a vSAN cluster. However, enabling this parameter might cause the ESXi host to become unresponsive.

ESXi Log File Locations

ESXi records host activity in log files, using a syslog facility.

Table 248: ESXi Log File Locations

Component	Location	Purpose
Authentication	/var/log/auth.log	Contains all events related to authentication for the local system.
ESXi host agent log	/var/log/hostd.log	Contains information about the agent that manages and configures the ESXi host and its virtual machines.
Shell log	/var/log/shell.log	Contains a record of all commands typed into the ESXi Shell and shell events (for example, when the shell was enabled).
System messages	/var/log/syslog.log	Contains all general log messages and can be used for troubleshooting. This information was formerly located in the messages log file.
vCenter Server agent log	/var/log/vpxa.log	Contains information about the agent that communicates with vCenter Server (if the host is managed by vCenter Server).
Virtual machines	The same directory as the affected virtual machine's configuration files, named <code>vmware.log</code> and <code>vmware*.log</code> . For example, <code>/vmfs/volumes/datastore/virtual machine/vmware.log</code>	Contains virtual machine power events, system failure information, tools status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and so on.
VMkernel	/var/log/vmkernel.log	Records activities related to virtual machines and ESXi.
VMkernel summary	/var/log/vmksummary.log	Used to determine uptime and availability statistics for ESXi (comma separated).
VMkernel warnings	/var/log/vmkwarning.log	Records activities related to virtual machines and ESXi.
Quick Boot	/var/log/loadESX.log	Contains all events related to restarting an ESXi host through Quick Boot.
Trusted infrastructure agent	/var/run/log/kmxa.log	Records activities related to the Client Service on the ESXi Trusted Host.
Key Provider Service	/var/run/log/kmxd.log	Records activities related to the vSphere Trust Authority Key Provider Service.
Attestation Service	/var/run/log/attestd.log	Records activities related to the vSphere Trust Authority Attestation Service.
ESX Token Service	/var/run/log/esxtokend.log	Records activities related to the vSphere Trust Authority ESX Token Service.
ESX API Forwarder	/var/run/log/esxapiadapter.log	Records activities related to the vSphere Trust Authority API forwarder.

Securing Fault Tolerance Logging Traffic

VMware Fault Tolerance (FT) captures inputs and events that occur on a primary virtual machine and sends them to the secondary virtual machine, which is running on another host.

This logging traffic between the primary and secondary virtual machines is unencrypted and contains guest network and storage I/O data, as well as the memory contents of the guest operating system. This traffic might include sensitive data

such as passwords in plain text. To avoid such data being divulged, ensure that this network is secured, especially to avoid man-in-the-middle attacks. For example, use a private network for FT logging traffic. You can also encrypt the FT logging traffic.

Activate Fault Tolerance Encryption

You can encrypt Fault Tolerance log traffic.

FT encryption requires SMP-FT. Encryption on Legacy FT (Record-Replay FT) is not supported.

vSphere Fault Tolerance performs frequent checks between a primary VM and secondary VM so that the secondary VM can quickly resume from the last successful checkpoint. The checkpoint contains the VM state that has been modified since the previous checkpoint. You can encrypt Fault Tolerance log traffic.

When you turn on Fault Tolerance, FT encryption is set to **Opportunistic** by default, which means it activates encryption only if both the primary and secondary host are capable of encryption. Follow this procedure if you need to change the FT encryption mode manually.

Note: Fault Tolerance supports vSphere Virtual Machine Encryption with vSphere 7.0 Update 2 and later. In-guest and array-based encryption do not depend on or interfere with VM encryption. Having multiple encryption layers uses additional compute resources, which might impact virtual machine performance. The impact varies with hardware as well as the amount and type of I/O, but overall performance impact is negligible for most workloads. The effectiveness and compatibility of back-end storage features such as deduplication, compression, and replication might also be affected by VM encryption.

1. Select the VM and choose **Edit Settings**.
2. Under **VM Options** select the **Encrypted FT** drop-down menu.
3. Choose one of the following options:

Option	Description
Disabled	Do not turn on encrypted Fault Tolerance logging.
Opportunistic	Turn on encryption only if both sides are capable. A Fault Tolerance VM is allowed to move to an ESXi host which does not support encrypted Fault Tolerance logging.
Required	Choose hosts for Fault Tolerance primary and secondary that both support encrypted FT logging.

Note: While VM encryption is activated, FT encryption mode is set to **Required** by default and cannot be modified.

When FT encryption mode is set to **Required**:

- When you turn on FT, only FT encryption supported hosts are listed for the placement of FT secondary.
- FT failover can only happen on the FT encryption supported hosts.

4. Click **OK**.

Managing ESXi Audit Records

Audit records conform to RFC 5424 and contain information about events pertaining to items such as the time, status, description, and user information logged for events that have occurred from actions on ESXi hosts. Both local and remote audit record keeping are available. Audit record keeping is deactivated by default. You must manually activate both local and remote auditing modes.

Audit records, when enabled, are stored locally as a fixed-size buffer of recent additions. Once the audit records fill the buffer, new audit records overwrite the oldest records.

Audit records are stored in RFC 5424 format but are transmitted to remote hosts in compliance with the specified message formatting (RFC 3164 for ESXi 7.0 Update 3; and RFC 3164 or RFC 5424 for ESXi 8.0 and later). The audit records are part of the stream of syslog messages. You can configure the local storage and transmission of audit records

independently of each other. When you activate both local storage and transmission, the audit records are stored and transmitted simultaneously.

During a loss of connection between the ESXi host transmitting syslog data and a remote host, audit records are dropped if the available buffer space is exceeded. Upon reconnection, the system generates an audit message indicating potential message loss.

Configuring Audit Records

You use ESXCLI to configure the local audit record keeping. For more information, see *ESXCLI Concepts and Examples*.

Viewing Audit Records

You can view the audit records as follows.

- Local: Use the `ESXi/bin/viewAudit` application.
- Remote: Configure a remote audit server using ESXCLI. For more information, see [Enable the Transmission of Audit Records to a Remote Host with ESXCLI](#).

You can also use the `FetchAuditRecords` API (in the `DiagnosticsManager` managed object) to view audit records.

Securing the ESXi Configuration

In vSphere 7.0 Update 2 and later, the ESXi configuration is protected by encryption.

What Is a Secure ESXi Configuration

Many ESXi services store secrets in their configuration files. These configurations persist in an ESXi host's boot bank as an archived file. Before vSphere 7.0 Update 2, the archived ESXi configuration file is not encrypted. In vSphere 7.0 Update 2 and later, the archived configuration file is encrypted. As a result, attackers cannot read or alter this file directly, even if they have physical access to the ESXi host's storage.

In addition to preventing an attacker from accessing secrets, a secure ESXi configuration when used with a TPM can save virtual machine encryption keys across reboots. When the ESXi host is configured with a TPM, the TPM is used to "seal" the configuration to the host, providing a strong security guarantee. As a result, encrypted workloads can continue to function when a key server is unavailable or unreachable. See [vSphere Key Persistence on ESXi Hosts](#).

You do not need to activate the ESXi configuration encryption manually. When you install or upgrade to vSphere 7.0 Update 2 or later, the archived ESXi configuration file is encrypted.

For tasks associated with a secure ESXi configuration, see [Manage a Secure ESXi Configuration](#).

ESXi Configuration Files Before vSphere 7.0 Update 2

The configuration of an ESXi host consists of configuration files for each service that runs on the host. The configuration files typically reside in the `/etc/` directory, but they can also reside in other namespaces. The configuration files contain run-time information about the state of the services. Over time, the default values in the configuration files can change, for example, when you change settings on the ESXi host. A cron job backs up the ESXi configuration files periodically, or when ESXi shuts down gracefully, or on demand, and creates an archived configuration file in the boot bank. When ESXi reboots, it reads the archived configuration file and recreates the state that ESXi was in when the backup was taken. Before vSphere 7.0 Update 2, the archived configuration file is unencrypted. As a result, it is possible for an attacker who has access to the physical ESXi storage to read and alter this file while the system is offline.

How Is Secure ESXi Configuration Implemented

During the first boot after installing or upgrading the ESXi host to vSphere 7.0 Update 2 or later, the following occurs:

- If the ESXi host has a TPM, and it is activated in the firmware, the archived configuration file is encrypted by an encryption key stored in the TPM. From this point on, the configuration of the host is sealed by the TPM.
- If the ESXi host does not have a TPM, ESXi uses a Key Derivation Function (KDF) to generate a secure configuration encryption key for the archived configuration file. The inputs to the KDF are stored on disk in the `encryption.info` file.

Note: When an ESXi host has an activated TPM device, you gain additional protection.

When the ESXi host reboots after the first boot, the following occurs:

- If the ESXi host has a TPM, the host must obtain the encryption key from the TPM for that specific host. If the TPM measurements satisfy the sealing policy that was used when creating the encryption key, then the host obtains the encryption key from the TPM.
- If the ESXi host does not have a TPM, ESXi reads information from the `encryption.info` file to unlock the secure configuration.

Secure ESXi Configuration Requirements

- ESXi 7.0 Update 2 or later
- TPM 2.0 for configuration encryption and ability to use a sealing policy

Secure ESXi Configuration Recovery Key

A secure ESXi configuration includes a recovery key. If you must recover the ESXi secure configuration, you use a recovery key whose contents you enter as a command-line boot option. You can list the recovery key to create a recovery key backup. You can also rotate the recovery key as part of your security requirements.

Taking a backup of the recovery key is an important part of managing your secure ESXi configuration. vCenter Server generates an alarm to remind you to back up the recovery key.

Secure ESXi Configuration Recovery Key Alarm

Taking a backup of the recovery key is an important part of managing your secure ESXi configuration. Whenever an ESXi host in TPM mode is connected or reconnected to vCenter Server, vCenter Server generates an alarm to remind you to back up the recovery key. When you reset the alarm, it is not triggered again unless conditions change.

Best Practices for Secure ESXi Configuration

Follow these best practices for the secure ESXi recovery key:

- When you list a recovery key, it is temporarily displayed in an untrusted environment and is in memory. Remove traces of the key.
 - Rebooting the host removes the residual key in memory.
 - For enhanced protection, you can activate encryption mode on the host. See [Activate Host Encryption Mode Explicitly](#).
- When you perform a recovery:
 - To eliminate any traces of the recovery key in an untrusted environment, reboot the host.
 - For enhanced security, rotate the recovery key to use a new key after having recovered the key one time.

What Are TPM Sealing Policies

A TPM can use Platform Configuration Register (PCR) measurements to implement policies that restrict unauthorized access to sensitive data. When you install or upgrade an ESXi host with a TPM to vSphere 7.0 Update 2 and later, the TPM seals the sensitive information by using a policy that incorporates the secure boot setting. This policy checks that if secure boot was activated when data was first sealed with the TPM, then secure boot must still be activated when attempting to unseal the data on a subsequent boot.

Secure boot is part of the UEFI firmware standard. With UEFI Secure Boot activated, a host refuses to load any UEFI driver or app unless the operating system bootloader has a valid digital signature.

You can choose to deactivate or activate UEFI Secure Boot enforcement. See [Activate or Deactivate the Secure Boot Enforcement for a Secure ESXi Configuration](#).

Note: If you do not activate a TPM when you install or upgrade to vSphere 7.0 Update 2 or later, you can do so later with the following command.

```
esxcli system settings encryption set --mode=TPM
```

Once you have activated the TPM, you cannot undo the setting.

The `esxcli system settings encryption set` command fails on some TPMs even when the TPM is activated for the host.

- In vSphere 7.0 Update 2: TPMs from NationZ (NTZ), Infineon Technologies (IFX), and certain new models (like NPCT75x) from Nuvoton Technologies Corporation (NTC)
- In vSphere 7.0 Update 3: TPMs from NationZ (NTZ)

If an installation or upgrade of vSphere 7.0 Update 2 or later is unable to use the TPM during the first boot, the installation or upgrade continues, and the mode defaults to NONE (that is, `--mode=NONE`). The resulting behavior is as though the TPM is not activated.

The TPM can also enforce the setting for the `execInstalledOnly` boot option in the sealing policy. The `execInstalledOnly` enforcement is an advanced ESXi boot option that guarantees that the VMkernel executes only binaries that have been properly packaged and signed as part of a VIB. The `execInstalledOnly` boot option has a dependency on the secure boot option. The secure boot enforcement must be activated before you can enforce the `execInstalledOnly` boot option in the sealing policy. See [Activate or Deactivate the execInstalledOnly Enforcement for a Secure ESXi Configuration](#).

Manage a Secure ESXi Configuration

You can use ESXCLI commands to list the secure ESXi configuration recovery key, rotate the recovery key, and change the TPM policies (for example, enforcing UEFI Secure Boot).

List the Contents of the Secure ESXi Configuration Recovery Key

You can use ESXCLI to show the contents of the secure ESXi configuration recovery key.

- Have access to the ESXCLI command set. You can run ESXCLI commands remotely, or run them in the ESXi Shell.
- Required privilege for using ESXCLI standalone version or through PowerCLI: **Host > Config > Settings**

This task applies only to an ESXi host that has a TPM. In general, you list the contents of the secure ESXi configuration recovery key to create a backup, or as part of rotating recovery keys.

1. Run the following command on the ESXi host.

```
esxcli system settings encryption recovery list
```

2. Save the output in a secure, remote location as a backup, in case you must recover the secure configuration.

The recovery key ID and key are displayed.

List the Secure ESXi Configuration Recovery Key

```
[root@host1] esxcli system settings encryption recovery list
```

```
Recovery ID                                Key
-----
{2DDD5424-7F3F-406A-8DA8-D62630F6C8BC}
478269-039194-473926-430939-686855-231401-642208-184477-602511
```

-225586-551660-586542-338394-092578-687140-267425

Rotate the Secure ESXi Configuration Recovery Key

You can use ESXCLI to rotate the secure ESXi configuration recovery key.

- Have access to the ESXCLI command set. You can run ESXCLI commands remotely, or run them in the ESXi Shell.
- Required privilege for using ESXCLI standalone version or through PowerCLI: **Host > Config > Settings**

This task applies only to an ESXi host that has a TPM. You can rotate the ESXi secure configuration recovery key as part of your security best practices.

1. List the recovery key.

See [List the Contents of the Secure ESXi Configuration Recovery Key](#).

2. Run the following command.

```
esxcli system settings encryption recovery rotate [-k keyID] -u uuid
```

In this command, the optional *keyID* is the key ID in the VMkernel key cache and *uuid* is the Recovery ID (obtained from the `esxcli system settings encryption recovery list` command). If you do not supply the optional key ID, ESXi replaces the old recovery key with a new recovery key that is randomly generated.

The recovery key is now set to the contents of the key referenced by key ID, if provided. Otherwise, ESXi provides a new key ID.

Troubleshooting and Recovering the Secure ESXi Configuration

You can troubleshoot and recover from boot problems that you might encounter with a secure ESXi Configuration.

If you clear a TPM (that is, the seed values in the TPM are reset), if a TPM fails, or if you replace the motherboard or TPM device, or both, you must take steps to recover the ESXi secure configuration. You must have the recovery key to recover the configuration. Until you recover the configuration, the ESXi host cannot boot. See [Recover the Secure ESXi Configuration](#).

Although uncommon, it is possible that an ESXi host might fail to restore or decrypt the secure configuration, preventing the host from booting. Possible situations include:

- Change to secure boot setting (or other policy)
- Actual tampering
- The recovery key is unavailable

To troubleshoot these conditions, see the VMware knowledge base article at <https://kb.vmware.com/s/article/81446>.

Recover the Secure ESXi Configuration

If a TPM fails, or if you clear a TPM, you must recover the secure ESXi Configuration. Until you recover the configuration, the ESXi host cannot boot.

Get your recovery key. You should have previously listed and stored the recover key. See [List the Contents of the Secure ESXi Configuration Recovery Key](#).

Recovering the secure ESXi configuration refers to the following situations:

- You cleared the TPM (that is, the seeds in the TPM were reset).
- The TPM failed.
- You replaced the motherboard or the TPM device, or both.

To troubleshoot other secure ESXi configuration problems, see the VMware knowledge base article at <https://kb.vmware.com/s/article/81446>.

Perform the recovery manually. Do not perform the recovery as part of an installation or upgrade script.

1. Optional: If the TPM failed, move the disk (having the boot bank) to another host with a TPM.
2. Start the ESXi host.
3. When the ESXi installer window appears, press Shift+O to edit boot options.
4. To recover the configuration, at the command prompt, append the following boot option to any existing boot options.

```
encryptionRecoveryKey=recovery_key
```

The secure ESXi configuration is recovered and the ESXi host boots.

5. To persist the change, enter the following command:

```
/sbin/auto-backup.sh
```

When you enter the recovery key, it is temporarily displayed in an untrusted environment and is in memory. Though not necessary, as a best practice, you can remove residual traces of the key in memory by rebooting the host. Or, you can rotate the key. See [Rotate the Secure ESXi Configuration Recovery Key](#).

Activate or Deactivate the Secure Boot Enforcement for a Secure ESXi Configuration

You can choose to activate UEFI secure boot enforcement, or deactivate a previously activated UEFI secure boot enforcement. You must use ESXCLI to change the setting in the TPM on the ESXi host.

- Have access to the ESXCLI command set. You can run ESXCLI commands remotely, or run them in the ESXi Shell.
- Required privilege for using ESXCLI standalone version or through PowerCLI: **Host > Config > Settings**

This task applies only to ESXi hosts that have a TPM. UEFI Secure boot is a firmware setting for ensuring that the software launched by the firmware is trusted. To learn more, see [UEFI Secure Boot for ESXi Hosts](#). The enablement of UEFI Secure boot can be enforced upon every boot by using the TPM.

1. List the current settings on the ESXi host.

```
esxcli system settings encryption get
Mode: TPM
Require Executables Only From Installed VIBs: false
Require Secure Boot: true
```

If secure boot enforcement is activated, Require Secure Boot displays true. If secure boot enforcement is deactivated, Require Secure Boot displays false.

If Mode appears as NONE, you must activate the TPM in the firmware of the host and set the mode by running the following command:

```
esxcli system settings encryption set --mode=TPM
```

2. Activate or deactivate the secure boot enforcement.

Option	Description
Activate	1. Shut down the host gracefully. For example, right-click the ESXi host in the vSphere Client and select Power > Shut Down. 2. Activate secure boot in the firmware of the host. See your specific vendor hardware documentation. 3. Restart the host. 4. Run the following ESXCLI command. <pre>esxcli system settings encryption set --require-secure-boot=T</pre> 5. Verify the change. <pre>esxcli system settings encryption get</pre>

Option	Description
	Mode: TPM Require Executables Only From Installed VIBs: false Require Secure Boot: true Confirm that Required Secure Boot displays true. 6. To save the setting, run the following command. <pre>/bin/backup.sh 0</pre>
Deactivate	1. Run the following ESXCLI command. <pre>esxcli system settings encryption set --require-secure-boot=F</pre> 2. Verify the change. <pre>esxcli system settings encryption get</pre> Mode: TPM Require Executables Only From Installed VIBs: false Require Secure Boot: false Confirm that Require Secure Boot displays false. 3. To save the setting, run the following command. <pre>/bin/backup.sh 0</pre> You can choose to deactivate the secure boot in the firmware of the host, but at this point the dependency between the firmware setting and the TPM enforcement is no longer set.

The ESXi host runs with secure boot enforcement activated or deactivated, depending on your choice.

Note:

If you do not activate a TPM when you install or upgrade to vSphere 7.0 Update 2 or later, you can do so later with the following command.

```
esxcli system settings encryption set --mode=TPM
```

Once you have activated the TPM, you cannot undo the setting.

The `esxcli system settings encryption set` command fails on some TPMs even when the TPM is activated for the host.

- In vSphere 7.0 Update 2: TPMs from NationZ (NTZ), Infineon Technologies (IFX), and certain new models (like NPCT75x) from Nuvoton Technologies Corporation (NTC)
- In vSphere 7.0 Update 3: TPMs from NationZ (NTZ)

If an installation or upgrade of vSphere 7.0 Update 2 or later is unable to use the TPM during the first boot, the installation or upgrade continues, and the mode defaults to NONE (that is, `--mode=NONE`). The resulting behavior is as though the TPM is not activated.

Activate or Deactivate the `execInstalledOnly` Enforcement for a Secure ESXi Configuration

You can choose to activate `execInstalledOnly` enforcement, or deactivate a previously enabled `execInstalledOnly` enforcement. You must use ESXCLI to change the setting in the TPM on the ESXi host. UEFI secure boot enforcement must be activated before you can activate the `execInstalledOnly` enforcement.

- To activate the `execInstalledOnly` enforcement, you must first activate the UEFI secure boot enforcement. The `execInstalledOnly` enforcement is built on top of the UEFI secure boot enforcement. See [Activate or Deactivate the Secure Boot Enforcement for a Secure ESXi Configuration](#).

- Have access to the ESXCLI command set. You can run ESXCLI commands remotely, or run them in the ESXi Shell.
- Required privilege for using ESXCLI standalone version or through PowerCLI: **Host > Config > Settings**

This task applies only to ESXi hosts that have a TPM. The `execInstalledOnly` advanced ESXi boot option, when set to `TRUE`, guarantees that the VMkernel executes only those binaries that have been packaged and signed as part of a VIB. The enablement of this boot option can be enforced upon every boot by using the TPM.

1. List the current settings on the ESXi host.

```
esxcli system settings encryption get
Mode: TPM
Require Executables Only From Installed VIBs: false
Require Secure Boot: true
```

If `execInstalledOnly` enforcement is activated, `Require Executables Only From Installed VIBs` displays `true`. If `execInstalledOnly` enforcement is deactivated, `Require Executables Only From Installed VIBs` displays `false`. To activate the `execInstalledOnly` enforcement, the secure boot enforcement must be activated, and `Require Secure Boot` displays `true` in this case.

If `Mode` appears as `NONE`, you must enable the TPM in the firmware of the host and set the mode by running the following command:

```
esxcli system settings encryption set --mode=TPM
```

Also, if `Require Secure Boot` displays as `False`, see [Activate or Deactivate the Secure Boot Enforcement for a Secure ESXi Configuration](#) to activate the enforcement.

2. Activate or deactivate the `execInstalledOnly` enforcement.

Option	Description
Activate	1. Verify that the secure boot option is activated. <pre>esxcli system settings encryption get Mode: TPM Require Executables Only From Installed VIBs: false Require Secure Boot: true</pre> Confirm that <code>Require Secure Boot</code> displays <code>true</code>. If not, see Activate or Deactivate the Secure Boot Enforcement for a Secure ESXi Configuration. 2. To configure the runtime value of the <code>execInstalledOnly</code> boot option to <code>TRUE</code>, run the following ESXCLI command. <pre>esxcli system settings kernel set -s execInstalledOnly -v TRUE</pre> 3. Shut down the host gracefully. For example, right-click the ESXi host in the vSphere Client and select Power > Shut Down. 4. Restart the host. 5. To set the <code>execInstalledOnly</code> enforcement, run the following ESXCLI command. <pre>esxcli system settings encryption set --require-exec-installed-only=T</pre> 6. Verify the change. <pre>esxcli system settings encryption get Mode: TPM Require Executables Only From Installed VIBs: true</pre>

Option	Description
	Require Secure Boot: true Confirm that Require Executables Only From Installed VIBs displays true. 7. To save the setting, run the following command. <pre>/bin/backup.sh 0</pre>
Deactivate	1. Run the following ESXCLI command. <pre>esxcli system settings encryption set --require-exec-installed-only=F</pre> 2. Verify the change. <pre>esxcli system settings encryption get Mode: TPM Require Executables Only From Installed VIBs: false Require Secure Boot: true</pre> Confirm that Require Executables Only From Installed VIBs displays false. 3. To save the setting, run the following command. <pre>/bin/backup.sh 0</pre> The TPM no longer enforces the execInstalledOnly boot option.

The ESXi host runs with execInstalledOnly enforcement activated or deactivated, depending on your choice.

Deactivate the execInstalledOnly Internal Runtime Option

When you install or upgrade to ESXi 8.0 or later, the execInstalledOnly internal runtime option is activated on hosts by default. This option helps protect your hosts against ransomware attacks. If your ESXi 8.0 or later hosts still run non-VIB binaries from external sources, you can deactivate the execInstalledOnly internal runtime option.

To deactivate the execInstalledOnly internal runtime option, you must have root access to the ESXi host. You can use ESXCLI, PowerCLI, or the API. The task that follows uses ESXCLI.

CAUTION: Deactivating the execInstalledOnly internal runtime option leaves you more vulnerable to attacks.

The execInstalledOnly option helps protect your hosts against ransomware attacks by ensuring that the VMkernel executes only those binaries on a host that have been properly packaged and signed as part of a valid VIB.

The execInstalledOnly option is both a boot and an internal runtime option. The execInstalledOnly boot option, also called a kernel option, was introduced in ESXi 5.5. The execInstalledOnly boot option is deactivated by default. In vSphere 7.0 Update 2 and later, you can enforce the execInstalledOnly boot option upon every boot by using a TPM. For more information, see [Activate or Deactivate the execInstalledOnly Enforcement for a Secure ESXi Configuration](#).

The execInstalledOnly internal runtime option added in ESXi 8.0 is activated on hosts by default. The execInstalledOnly boot option continues to be deactivated by default, except a previously enabled execInstalledOnly boot option overwrites the internal runtime option if you set both.

Note: The execInstalledOnly option is independent of Secure Boot. Secure boot checks that all installed VIBs are signed. For more information, see [UEFI Secure Boot for ESXi Hosts](#).

When you deactivate the execInstalledOnly internal runtime option, vCenter Server warnings appear for the host.

1. Connect to the ESXi host by SSH.
2. To deactivate the execInstalledOnly internal runtime option, enter the following ESXCLI command.

```
esxcli system settings advanced set -o /User/execInstalledOnly -i 0
```

Securing vCenter Server Systems

Securing vCenter Server includes ensuring security of the host where vCenter Server is running, following best practices for assigning privileges and roles, and verifying the integrity of the clients that connect to vCenter Server.

Best Practices for vCenter Server Access Control

Strictly control access to different vCenter Server components to increase security for the system.

The following guidelines help ensure security of your environment.

Use Named Accounts to Access vCenter Server

- Grant the Administrator role only to those administrators who are required to have it. You can create custom roles or use the No cryptography administrator role for administrators with more limited privileges. Do not apply this role to any group whose membership is not strictly controlled.
- Make sure that applications use unique service accounts when connecting to a vCenter Server system.

Monitor Privileges of vCenter Server Administrator Users

Not all administrator users must have the Administrator role. Instead, create a custom role with the appropriate set of privileges and assign it to other administrators.

Users with the vCenter Server Administrator role have privileges on all objects in the hierarchy. For example, by default the Administrator role allows users to interact with files and programs inside the guest operating system of a virtual machine. Assigning that role to too many users can lessen virtual machine data confidentiality, availability, or integrity. Create a role that gives the administrators the privileges they need, but remove some of the virtual machine management privileges. See also [Privilege Recorder](#).

Minimize Access to the vCenter Server Appliance

Do not allow users to log directly in to the vCenter Server Appliance. Users who are logged in to the vCenter Server Appliance can cause harm, either intentionally or unintentionally, by altering settings and modifying processes. Those users also have potential access to vCenter Server credentials, such as the SSL certificate. Allow only users who have legitimate tasks to perform to log in to the system and ensure that login events are audited.

Grant Minimal Privileges to Database Users

The database user requires only certain privileges specific to database access.

Some privileges are required only for installation and upgrade. You can remove these privileges from the database administrator after vCenter Server is installed or upgraded.

Restrict Datastore Browser Access

Assign the **Datastore > Browse datastore** privilege only to users or groups who really need that privilege. Users with the privilege can view, upload, or download files on datastores associated with the vSphere deployment through the Web browser or the vSphere Client.

Restrict Users from Running Commands in a Virtual Machine

By default, a user with the Administrator role can interact with the files and the programs of a guest operating system within a virtual machine. To reduce the risk of breaching guest confidentiality, availability, or integrity, create a custom nonquest access role without the **Virtual machine > Guest operations** privilege. See [Restrict Users from Running Commands Within a Virtual Machine](#).

Consider Modifying the Password Policy for vpxuser

By default, vCenter Server changes the vpxuser password automatically every 30 days. Ensure that this setting meets company policy, or configure the vCenter Server password policy. See [Set the vCenter Server Password Policy](#).

Note: Make sure that password aging policy is not too short.

Check Privileges After Restarting vCenter Server

Check for privilege reassignment when you restart vCenter Server. If the user or group that has the Administrator role on the root folder cannot be validated during a restart, the role is removed from that user or group. In its place, vCenter Server grants the Administrator role to the vCenter Single Sign-On administrator, administrator@vsphere.local by default. This account can then act as the vCenter Server administrator.

Reestablish a named administrator account and assign the Administrator role to that account to avoid using the anonymous vCenter Single Sign-On administrator account (administrator@vsphere.local by default).

Use High Encryption Levels for Remote Desktop Protocol

On each Windows computer in the infrastructure, ensure that the Remote Desktop Protocol (RDP) Host Configuration settings are set to ensure the highest level of encryption appropriate for your environment.

Verify vSphere Client Certificates

Instruct users of the vSphere Client or other client applications to heed certificate verification warnings. Without certificate verification, the user might be the subject of a MiTM attack.

Set the vCenter Server Password Policy

By default, vCenter Server changes the vpxuser password automatically every 30 days. You can change that value from the vSphere Client.

1. Log in to the vCenter Server system by using the vSphere Client.
2. Select the vCenter Server system in the object hierarchy.
3. Click **Configure**.
4. Click **Advanced Settings** and click **Edit Settings**.
5. Click the **Filter** icon and enter `VimPasswordExpirationInDays`.
6. Set `VirtualCenter.VimPasswordExpirationInDays` to comply with your requirements.

Removing Expired or Revoked Certificates and Logs from Failed Installations

Leaving expired or revoked certificates or leaving vCenter Server installation logs for failed installation on your vCenter Server system can compromise your environment.

Removing expired or revoked certificates is required for the following reasons.

- If expired or revoked certificates are not removed from the vCenter Server system, the environment can be subject to a MiTM attack
- In certain cases, a log file that contains the database password in plain text is created on the system if vCenter Server installation fails. An attacker who breaks into the vCenter Server system, might gain access to this password and, at the same time, access to the vCenter Server database.

Limiting vCenter Server Network Connectivity

For improved security, avoid putting the vCenter Server system on any network other than a management network, and ensure that vSphere management traffic is on a restricted network. By limiting network connectivity, you limit certain types of attack.

vCenter Server requires access to a management network only. Avoid putting the vCenter Server system on other networks such as your production network or storage network, or on any network with access to the Internet. vCenter Server does not need access to the network where vMotion operates.

vCenter Server requires network connectivity to the following systems.

- All ESXi hosts.
- The vCenter Server database.
- Other vCenter Server systems (if the vCenter Server systems are part of a common vCenter Single Sign-On domain for purposes of replicating tags, permissions, and so on).
- Systems that are authorized to run management clients. For example, the vSphere Client, a Windows system where you use the PowerCLI, or any other SDK-based client.
- Infrastructure services such as DNS, Active Directory, and PTP or NTP.
- Other systems that run components that are essential to functionality of the vCenter Server system.

Use the firewall on the vCenter Server. Include IP-based access restrictions so that only necessary components can communicate with the vCenter Server system.

Evaluate the Use of Linux Clients with CLIs and SDKs

Communications between client components and a vCenter Server system or ESXi hosts are protected by SSL-based encryption by default. Linux versions of these components do not perform certificate validation. Consider restricting the use of these clients.

To improve security, you can replace the VMCA-signed certificates on the vCenter Server system and on the ESXi hosts with certificates that are signed by an enterprise or third-party CA. However, certain communications with Linux clients might still be vulnerable to machine-in-the-middle attacks. The following components are vulnerable when they run on the Linux operating system.

- ESXCLI commands
- vSphere SDK for Perl scripts
- Programs that are written using the vSphere Web Services SDK

You can relax the restriction against using Linux clients if you enforce proper controls.

- Restrict management network access to authorized systems only.
- Use firewalls to ensure that only authorized hosts are allowed to access vCenter Server.
- Use bastion hosts (jump-box systems) to ensure that the Linux clients are behind the "jump."

Examine vSphere Client Plug-Ins

vSphere Client extensions run at the same privilege level as the user who is logged in. A malicious extension can masquerade as a useful plug-in and perform harmful operations such as stealing credentials or changing the system configuration. To increase security, use an installation that includes only authorized extensions from trusted sources.

You must have privileges to access the vCenter Single Sign-On service. These privileges differ from vCenter Server privileges.

A vCenter Server installation includes an extensibility framework for the vSphere Client. You can use this framework to extend the client with menu selections or toolbar icons. The extensions can provide access to vCenter Server add-on components or external, Web-based functionality.

Using the extensibility framework results in a risk of introducing unintended capabilities. For example, if an administrator installs a plug-in in an instance of the vSphere Client, the plug-in can run arbitrary commands with the privilege level of that administrator.

To protect against a potential compromise of your vSphere Client, examine all installed plug-ins periodically and make sure that each plug-in comes from a trusted source.

1. Log in to the vSphere Client as administrator@vsphere.local or a user with vCenter Single Sign-On privileges.
2. From the Home page, select **Administration**, then select **Client Plug-Ins** under **Solutions**.
3. Examine the list of client plug-ins.

vCenter Server Security Best Practices

Follow all best practices for securing a vCenter Server system. Additional steps help you make your vCenter Server more secure.

Configure Precision Time Protocol or Network Time Protocol

Ensure that all systems use the same relative time source. This time source must be in sync with an agreed-upon time standard such as Coordinated Universal Time (UTC). Synchronized systems are essential for certificate validation. Precision Time Protocol (PTP) and Network Time Protocol (NTP) also make it easier to track an intruder in log files. Incorrect time settings make it difficult to inspect and correlate log files to detect attacks, and make auditing inaccurate. See [Synchronize the Time in vCenter Server with an NTP Server](#).

Restrict vCenter Server Network Access

Restrict access to components that are required to communicate with the vCenter Server. Blocking access from unnecessary systems reduces the potential for attacks on the operating system.

For the list of all supported ports and protocols in VMware products, including vSphere and , see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com/>. You can search ports by VMware product, create a customized list of ports, and print or save port lists.

Configure a Bastion Host

To help protect your assets, configure a bastion host (also called a jump box) to perform elevated administrative tasks. A bastion host is a special-purpose computer that hosts a minimal number of administrative applications. All other unnecessary services are removed. The host typically resides on the management network. A bastion host increases the protection of assets through restricting login to key individuals, requiring firewall rules to log in, and adding monitoring through auditing tools.

vCenter Password Requirements and Lockout Behavior

To manage your vSphere environment, you must be aware of the vCenter Single Sign-On password policy, of vCenter Server passwords, and of lockout behavior.

This section discusses vCenter Single Sign-On passwords. See [ESXi Passwords and Account Lockout](#) for a discussion of passwords of ESXi local users.

vCenter Single Sign-On Administrator Password Requirements

The password for the administrator of vCenter Single Sign-On, administrator@vsphere.local by default, is specified by the vCenter Single Sign-On password policy. By default, this password must meet the following requirements:

- At least eight characters
- At least one lowercase character
- At least one numeric character
- At least one special character

The password for this user cannot be more than 20 characters long. Non-ASCII characters are allowed. Administrators can change the default password policy. See the *vSphere Authentication* documentation.

vCenter Server Password Requirements

In vCenter Server, password requirements are dictated by vCenter Single Sign-On or by the configured identity source, which can be Active Directory or OpenLDAP.

vCenter Single Sign-On Lockout Behavior

Users are locked out after a preset number of consecutive failed attempts. By default, users are locked out after five consecutive failed attempts in three minutes and a locked account is unlocked automatically after five minutes. You can change these defaults using the vCenter Single Sign-On lockout policy. See the *vSphere Authentication* documentation.

The vCenter Single Sign-On domain administrator, administrator@vsphere.local by default, is not affected by the lockout policy. The user is affected by the password policy.

vCenter Server Password Changes

If you know your password, you can change the password by using the `dir-cli password change` command. If you forget your password, the vCenter Single Sign-On administrator can reset your password by using the `dir-cli password reset` command.

Search the VMware Knowledge Base for information on password expiration and related topics in different versions of vSphere.

Verify Thumbprints for Legacy ESXi Hosts

In vSphere 6.0 and later, hosts are assigned VMCA certificates by default. If you change the certificate mode to thumbprint, you can continue to use thumbprint mode for legacy hosts. You can verify the thumbprints in the vSphere Client.

Note: Certificates are preserved across upgrades by default.

1. Browse to the vCenter Server in the vSphere Client inventory.
2. Click **Configure**.
3. Under **Settings**, click **General**.
4. Click **Edit**.
5. Click **SSL settings**.
6. If any of your ESXi 5.5 or earlier hosts require manual validation, compare the thumbprints listed for the hosts to the thumbprints in the host console.

To obtain the host thumbprint, use the Direct Console User Interface (DCUI).

- a) Log in to the direct console and press F2 to access the System Customization menu.
- b) Select **View Support Information**.

The host thumbprint appears in the column on the right.

7. If the thumbprint matches, select the **Verify** check box next to the host.
Hosts that are not selected will be disconnected after you click **OK**.
8. Click **Save**.

Required Ports for vCenter Server

The vCenter Server system must be able to send data to every managed host and receive data from the vSphere Client. To enable migration and provisioning activities between managed hosts, the source and destination hosts must be able to receive data from each other through predetermined TCP and UDP ports.

vCenter Server is accessed through predetermined TCP and UDP ports. If you manage network components from outside a firewall, you might be required to reconfigure the firewall to allow access on the appropriate ports. For the list of all supported ports and protocols in vSphere, see the VMware Ports and Protocols Tool™ at <https://ports.broadcom.com>.

During installation, if a port is in use or is blocked using a denylist, the vCenter Server installer displays an error message. You must use another port number to proceed with the installation. There are internal ports that are used only for inter-process communication.

VMware uses designated ports for communication. Additionally, the managed hosts monitor designated ports for data from vCenter Server. If a built-in firewall exists between any of these elements, the installer opens the ports during the installation or upgrade process. For custom firewalls, you must manually open the required ports. If you have a firewall between two managed hosts and you want to perform source or target activities, such as migration or cloning, you must configure a means for the managed hosts to receive data.

To configure the vCenter Server system to use a different port to receive vSphere Client data, see the *vCenter Server and Host Management* documentation.

Securing Virtual Machines

The guest operating system that runs in the virtual machine is subject to the same security risks as a physical system. Secure virtual machines like physical machines, and follow best practices discussed in this document and in the *Security Configuration Guide* (formerly known as the *Hardening Guide*).

The *Security Configuration Guide* is available at <https://core.vmware.com/security>.

Activate or Deactivate UEFI Secure Boot for a Virtual Machine

UEFI Secure Boot is a security standard that helps ensure that your PC boots using only software that is trusted by the PC manufacturer. For certain virtual machine hardware versions and operating systems, you can activate secure boot just as you can for a physical machine.

You can activate secure boot only if all prerequisites are met. If prerequisites are not met, the check box is not visible in the vSphere Client.

- Verify that the virtual machine operating system and firmware support UEFI boot.
 - EFI firmware
 - Virtual hardware version 13 or later.
 - Operating system that supports UEFI secure boot.

Note: Some guest operating systems do not support changing from BIOS boot to UEFI boot without guest OS modifications. Consult your guest OS documentation before changing to UEFI boot. If you upgrade a virtual machine that already uses UEFI boot to an operating system that supports UEFI secure boot, you can activate Secure Boot for that virtual machine.

- Turn off the virtual machine. If the virtual machine is running, the check box is dimmed.

In an operating system that supports UEFI secure boot, each piece of boot software is signed, including the bootloader, the operating system kernel, and operating system drivers. The virtual machine's default configuration includes several code signing certificates.

- A Microsoft certificate that is used only for booting Windows.
- A Microsoft certificate that is used for third-party code that is signed by Microsoft, such as Linux bootloaders.
- A VMware certificate that is used only for booting ESXi inside a virtual machine.

The virtual machine's default configuration includes one certificate for authenticating requests to modify the secure boot configuration, including the secure boot revocation list, from inside the virtual machine, which is a Microsoft KEK (Key Exchange Key) certificate.

VMware Tools version 10.1 or later is required for virtual machines that use UEFI secure boot. You can upgrade those virtual machines to a later version of VMware Tools when it becomes available.

For Linux virtual machines, VMware Host-Guest Filesystem is not supported in secure boot mode. Remove VMware Host-Guest Filesystem from VMware Tools before you activate secure boot.

Note: If you turn on secure boot for a virtual machine, you can load only signed drivers into that virtual machine.

This task describes how to use the vSphere Client to activate and deactivate secure boot for a virtual machine. You can also write scripts to manage virtual machine settings. For example, you can automate changing the firmware from BIOS to EFI for virtual machines with the following PowerCLI code:

```
$vm = Get-VM TestVM

$spec = New-Object VMware.Vim.VirtualMachineConfigSpec
$spec.Firmware = [VMware.Vim.GuestOsDescriptorFirmwareType]::efi
$vm.ExtensionData.ReconfigVM($spec)
```

See *VMware PowerCLI User's Guide* for more information.

1. Browse to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click the **VM Options** tab, and expand **Boot Options**.
4. Under **Boot Options**, ensure that firmware is set to **EFI**.
5. Select your task.
 - Select the **Secure Boot** check box to activate secure boot.
 - Deselect the **Secure Boot** check box to deactivate secure boot.
6. Click **OK**.

When the virtual machine boots, only components with valid signatures are allowed. The boot process stops with an error if it encounters a component with a missing or invalid signature.

Virtual Machine Security Best Practices

Following virtual machine security best practices helps ensure the integrity of your vSphere deployment.

General Virtual Machine Protection

A virtual machine is, in most respects, the equivalent of a physical server. Employ the same security measures in virtual machines that you do for physical systems.

Follow these best practices to protect your virtual machines. For additional information, see the *vSphere Security Configuration Guide* at <https://core.vmware.com/security-configuration-guide>.

Patch Virtual Machines

Keep all security measures up-to-date, including applying appropriate patches. Keep track of updates for dormant virtual machines that are powered off, because it can be easy to overlook them. For example, ensure that anti-virus software, anti-spyware, intrusion detection, and other protections are enabled for virtual machines in your virtual infrastructure. Also, ensure that you have enough space for the virtual machine logs.

Scan Virtual Machines for Viruses

Because each virtual machine hosts a standard operating system, you must protect it from viruses by installing anti-virus software. Depending on how you use the virtual machine, you might also want to install a software firewall.

Stagger the schedule for virus scans, particularly in deployments with a large number of virtual machines. Performance of systems in your environment degrades significantly if you scan all virtual machines simultaneously. Because software firewalls and anti-virus software can be virtualization-intensive, balance the need for these two security measures against virtual machine performance, especially if you are confident that your virtual machines are in a fully trusted environment.

Deactivate Serial Ports on Virtual Machines

Serial ports are interfaces for connecting peripherals to the virtual machine. Administrators often use serial ports to provide a direct, low-level connection to the console of a server. A virtual serial port allows for the same access to a virtual machine. Because serial ports allow for low-level access, and do not have strong controls like logging or privileges, keep them deactivated on virtual machines.

Use Templates to Deploy Virtual Machines

When you manually install guest operating systems and applications on a virtual machine, you introduce a risk of misconfiguration. By using a template to capture a hardened base operating system image with no applications installed, you can ensure that all virtual machines are created with a known baseline level of security.

You can use templates that can contain a hardened, patched, and properly configured operating system to create other, application-specific templates, or you can use the application template to deploy virtual machines.

Provide templates for virtual machine creation that contain hardened, patched, and properly configured operating system deployments.

If possible, deploy applications in templates as well. Ensure that the applications do not depend on information specific to the virtual machine to be deployed.

For more information about templates, see the *vSphere Virtual Machine Administration* documentation.

Minimize Use of the Virtual Machine Console

The virtual machine console provides the same function for a virtual machine that a monitor provides on a physical server. Users with access to the virtual machine console have access to virtual machine power management and removable device connectivity controls. Console access might therefore allow a malicious attack on a virtual machine.

1. Use native remote management services, such as terminal services and SSH, to interact with virtual machines.

Grant access to the virtual machine console only when necessary.

2. Limit the connections to the virtual machine console.

For example, in a highly secure environment, limit the connection to one. In some environments, you can increase the limit if several concurrent connections are necessary to accomplish normal tasks.

- a) In the vSphere Client, power off the virtual machine.
- b) Right-click the virtual machine and select **Edit Settings**.
- c) Click the **VM Options** tab, and expand **VMware Remote Console Options**.
- d) Enter the maximum number of sessions, for example, 2.
- e) Click **OK**.

Prevent Virtual Machines from Taking Over Resources

When one virtual machine consumes so much of the host resources that other virtual machines on the host cannot perform their intended functions, a Denial of Service (DoS) might occur. To prevent a virtual machine from causing a DoS, use host resource management features such as setting Shares and using resource pools.

By default, all virtual machines on an ESXi host share resources equally. You can use Shares and resource pools to prevent a denial of service attack that causes one virtual machine to consume so much of the host's resources that other virtual machines on the same host cannot perform their intended functions.

Do not set limits or use resource pools until you fully understand the impact.

1. Provision each virtual machine with just enough resources (CPU and memory) to function properly.
2. Use Shares to guarantee resources to critical virtual machines.
3. Group virtual machines with similar requirements into resource pools.
4. In each resource pool, leave Shares set to the default to ensure that each virtual machine in the pool receives approximately the same resource priority.

With this setting, a single virtual machine cannot use more than other virtual machines in the resource pool.

See the *vSphere Resource Management* documentation for information about shares and limits.

Deactivate Unnecessary Functions Inside Virtual Machines

Any service that runs in a virtual machine provides the potential for attack. By deactivating system components that are not necessary to support the application or service that is running on the system, you reduce the attack potential.

Virtual machines do not usually require as many services or functions as physical servers. When you virtualize a system, evaluate whether a particular service or function is necessary.

Note: When possible, install guest operating systems using "minimal" or "core" installation modes to reduce the size, complexity, and attack surface of the guest operating system.

- Deactivate unused services in the operating system.
For example, if the system runs a file server, turn off any Web services.
- Disconnect unused physical devices, such as CD/DVD drives, floppy drives, and USB adapters.
- Deactivate unused functionality, such as unused display features, or VMware Shared Folders, which enables sharing of host files to the virtual machine (Host Guest File System).
- Turn off screen savers.
- Do not run the X Window system on top of Linux, BSD, or Solaris guest operating systems unless it is necessary.

Remove Unnecessary Hardware Devices from Virtual Machines

Any activated or connected device in a virtual machines represents a potential attack channel. Users and processes with privileges on a virtual machine can connect or disconnect hardware devices, such as network adapters and CD-ROM drives. Attackers can use this capability to breach virtual machine security. Removing unnecessary hardware devices can help prevent attacks.

An attacker with access to a virtual machine can connect a disconnected hardware device and access sensitive information on media that is left in a hardware device. The attacker can potentially disconnect a network adapter to isolate the virtual machine from its network, resulting in a denial of service.

- Do not connect unauthorized devices to the virtual machine.
- Remove unneeded or unused hardware devices.
- Deactivate unnecessary virtual devices from within a virtual machine.
- Ensure that only required devices are connected to a virtual machine. Virtual machines rarely use serial or parallel ports. As a rule, CD/DVD drives are connected only temporarily during software installation.

1. Browse to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and click **Edit Settings**.
3. In the **Virtual Hardware** tab, click the ellipsis icon and select **Remove device** to deactivate hardware devices that are not required.

Include checks for the following devices:

- Serial ports
- Parallel ports
- USB controllers
- CD-ROM drives

Note: You must use PowerCLI commands to manage floppy drive devices in vSphere 7.0 and later.

Deactivate Unused Display Features on Virtual Machines

Attackers can use an unused display feature as a vector for inserting malicious code into your environment. Deactivate features that are not in use in your environment.

Power off the virtual machine.

1. Browse to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and click **Edit Settings**.
3. Select **Advanced Parameters**.
4. If appropriate, add or edit the following parameters.

Name	Description
<code>svga.vgaonly</code>	If you set this parameter to TRUE, advanced graphics functions no longer work. Do not set this parameter to TRUE with modern-day guest operating systems as they do not operate correctly. When <code>svga.vgaonly</code> is set to TRUE, only character-cell console mode is available. If you use this setting, <code>mks.enable3d</code> has no effect. Note: Apply this setting only to virtual machines that do not need a virtualized video card.
<code>mks.enable3d</code>	Set this parameter to FALSE on virtual machines that do not require 3D functionality.

5. Click **OK**.

Deactivate Copy and Paste Operations Between Guest Operating System and Remote Console

Copy and paste operations between the guest operating system and remote console are deactivated by default. For a secure environment, retain the default setting. If you require copy and paste operations, you must activate them using the vSphere Client.

Turn off the virtual machine.

The default values for these options are set to ensure a secure environment. However, you must set them to true explicitly for audit tools to be able to check that the setting is correct.

1. Browse to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and click **Edit Settings**.
3. Select **Advanced Parameters**.
4. Add or edit the following parameters.

Name	Value
<code>isolation.tools.copy.disable</code>	true
<code>isolation.tools.paste.disable</code>	true
<code>isolation.tools.setGUIOptions.enable</code>	false

These options override any settings made in the guest operating system's VMware Tools control panel.

5. Click **OK**.
6. Optional: If you made changes to the configuration parameters, restart the virtual machine.

Limiting Exposure of Sensitive Data Copied to the Virtual Machine Console Clipboard

Copy and paste operations are deactivated by default for hosts to prevent exposing sensitive data that has been copied to the clipboard.

When copy and paste is activated on a virtual machine running VMware Tools, you can copy and paste between the guest operating system and remote console. When the console window gains focus, processes running in the virtual machine

and non-privileged users can access the virtual machine console clipboard. If a user copies sensitive information to the clipboard before using the console, the user might expose sensitive data to the virtual machine. To prevent this problem, copy and paste operations for the guest operating system are deactivated by default.

It is possible to activate copy and paste operations for virtual machines if necessary.

Restrict Users from Running Commands Within a Virtual Machine

By default, a user who has the vCenter Server Administrator role can interact with files and applications within the guest operating system of a virtual machine. To reduce the risk of breaching guest confidentiality, availability, or integrity, create a nonguest access role without the **Virtual machine > Guest operations** privilege. Assign that role to administrators who do not need virtual machine file access.

Verify that you have **Administrator** privileges on the vCenter Server system where you create the role.

For security, be as restrictive about allowing access to the virtual data center as you are to the physical data center. Apply a custom role that does not include the **Virtual machine > Guest operations** privilege to users who require administrator privileges, but who are not authorized to interact with guest operating system files and applications.

For example, a configuration might include a virtual machine on the infrastructure that has sensitive information on it.

If tasks such as vMotion migration require that data center administrators can access the virtual machine, deactivate some remote guest operating system operations to ensure that those administrators cannot access sensitive information.

1. Log in to the vSphere Client as a user who has **Administrator** privileges on the vCenter Server system where you want to create the role.
2. Select **Administration** and click **Roles**.
3. Click the Administrator role and click **Clone**.
4. Enter a role name and description and click **OK**.
For example, enter `Administrator No Guest Access`.
5. Select the cloned role and click **Edit**.
6. Under the **Virtual machine** privilege, deselect **Guests operations**.
7. Click **Save**.

Select the vCenter Server system or the host and assign a permission that pairs the user or group that should have the new privileges to the newly created role. Remove those users from the Administrator role.

Prevent a Virtual Machine User or Process from Disconnecting Devices

Users and processes without root or administrator privileges within virtual machines can connect or disconnect devices, such as network adapters and CD-ROM drives, and can modify device settings. To increase virtual machine security, remove these devices.

Turn off the virtual machine.

You can prevent virtual machine users in the guest operating system, and processes running in the guest operating system, from making any changes to the devices by changing the virtual machine advanced settings.

1. Browse to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and click **Edit Settings**.
3. Click the **Advanced Parameters** tab.
4. Verify the following parameter, or add it.

Name	Value
isolation.device.connectable.disable	true

This setting does not affect a vSphere administrator's ability to connect or disconnect the devices attached to the virtual machine.

5. Click **OK**.

Prevent Guest Operating System Processes from Sending Configuration Messages to the Host

To ensure that the guest operating system does not modify configuration settings, you can prevent these processes from writing any name-value pairs to the configuration file.

Turn off the virtual machine.

1. Browse to the virtual machine in the vSphere Client inventory.
2. Right-click the virtual machine and click **Edit Settings**.
3. Select **Advanced Parameters**.
4. Verify the following parameter, or add it.

Name	Value
isolation.tools.setinfo.disable	true

5. Click **OK**.

Avoid Using Independent Nonpersistent Disks with Virtual Machines

When you use independent nonpersistent disks with virtual machines, successful attackers can remove any evidence that the machine was compromised by shutting down or rebooting the system. Without a persistent record of activity on a virtual machine, administrators might be unaware of an attack. Therefore, avoid using independent nonpersistent disks.

1. Ensure that virtual machine activity is logged remotely on a separate server, such as a syslog server or equivalent Windows-based event collector.
2. If remote logging of events and activity is not configured for the guest, ensure that scsiX:Y.mode is one of the following settings:
 - Not present
 - Not set to independent nonpersistent

When nonpersistent mode is not enabled, you cannot roll a virtual machine back to a known state when you reboot the system.

Securing Virtual Machines with Intel Software Guard Extensions

vSphere enables you to configure Virtual Intel® Software Guard Extensions (vSGX) for virtual machines. Using vSGX enables you to provide additional security to your workloads.

Some modern Intel CPUs implement a security extension called Intel® Software Guard Extensions (Intel® SGX). Intel SGX is a processor-specific technology for application developers who seek to protect select code and data from disclosure or modification. Intel SGX allows user-level code to define private regions of memory, called enclaves. The enclave contents are protected such that code running outside the enclave cannot access the enclave contents.

vSGX enables virtual machines to use Intel SGX technology if available on the hardware. To use vSGX, the ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the BIOS of the ESXi host. You can use the vSphere Client to enable SGX for a virtual machine.

In vSphere 8.0 and later, you can use remote attestation for a vSGX-enabled virtual machine. Intel SGX remote attestation is a security mechanism that allows you to establish an authenticated and secure communication channel with a trusted remote entity. To use remote attestation for virtual machines using SGX enclaves, hosts with a single CPU socket do not require Intel registration. To enable the remote attestation on a virtual machine running on a host with multiple CPU sockets, you must first register the host with the Intel Registration Server. If an SGX-capable host with multiple CPU sockets is not registered with the Intel Registration Server, you can only power on vSGX-enabled virtual machines that do not require remote attestation.

See the *vCenter Server and Host Management* documentation for more information about registering a multi-socket ESXi host with the Intel Registration Server.

Getting Started with vSGX

Virtual machines can use Intel SGX technology, if available on the hardware.

vSphere Requirements for vSGX

To use vSGX, your vSphere environment must meet these requirements:

- Virtual machine requirements:
 - EFI firmware
 - Hardware version 17 or later
 - To enable remote attestation, hardware version 20 or later
- Component requirements:
 - vCenter Server 7.0 and later
 - ESXi 7.0 and later
 - The ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the BIOS of the ESXi host.
 - To enable the remote attestation for the host, register the host with the Intel Registration Server. This way, the virtual machine running on the host can use the remote attestation. For more information on how to register a multi-socket ESXi, see the *vCenter Server and Host Management* documentation.
- Guest OS support:
 - Linux
 - Windows Server 2016 (64 bit) and later
 - Windows 10 (64 bit) and later

Supported Intel Hardware for vSGX

For supported Intel hardware for vSGX, consult the VSphere Compatibility Guide at <https://www.vmware.com/resources/compatibility/search.php>.

You might need to turn off hyperthreading on certain CPUs to enable SGX on the ESXi host. For more information, see the VMware KB article at <https://kb.vmware.com/s/article/71367>.

Unsupported VMware Features on vSGX

The following features are not supported in a virtual machine when vSGX is enabled:

- vMotion/DRS migration
- Virtual machine suspend and resume
- Virtual machine snapshots (Virtual machine snapshots are supported if you do not snapshot the virtual machine's memory.)
- Fault tolerance
- Guest Integrity (GI, platform foundation for VMware AppDefense™ 1.0)

Note:

These VMware features are not supported due to how the Intel SGX architecture functions. They are not the result of a VMware shortcoming.

Enable vSGX on a Virtual Machine

You can enable vSGX on a virtual machine at the same time that you create a virtual machine.

See [vSphere Requirements for vSGX](#).

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, an ESXi host or a cluster.
3. Right-click the object, select **New Virtual Machine**, and follow the prompts to create a virtual machine.
4. On the **Customize hardware** page, click the **Virtual Hardware** tab and expand **Security Devices**.
5. To enable SGX, select the **Enable** check box.
6. In the **Enclave page cache size (MB)** text box, enter the size of the cache size in MB.

Note: The enclave page cache size must be multiple of 2 MB.

7. To prevent the virtual machine from powering on hosts that do not support SGX remote attestation, such as unregistered multi-socket SGX hosts, select the **Remote attestation** check box.
8. From the **Launch control configuration** drop-down menu, select the appropriate mode.

Option	Action
Unlocked	This option enables the launch enclave configuration of the guest operating system.
Locked	This option allows you to configure the launch enclave. 1. Select the Launch enclave public key hash option. 2. To use one of the public keys configured on the host, select Use from host and from the drop-down menu, select a public key hash. 3. To enter the public key manually, select Enter manually and enter a valid SHA256 hash (64) characters key.

9. Click **OK**.

Enable vSGX on an Existing Virtual Machine

You can enable vSGX on an existing virtual machine.

See [vSphere Requirements for vSGX](#).

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. In the **Virtual Hardware** tab, expand **Security Devices**.
4. To enable SGX, select the **Enable** check box.
5. In the **Enclave page cache size (MB)** text box, enter the size of the cache size in MB.

Note: The enclave page cache size must be multiple of 2 MB.

6. To prevent the virtual machine from powering on hosts that do not support SGX remote attestation, such as unregistered multi-socket SGX hosts, select the **Remote attestation** check box.
7. From the **Launch control configuration** drop-down menu, select the appropriate mode.

Option	Action
Unlocked	This option enables the launch enclave configuration of the guest operating system.
Locked	This option allows you to configure the launch enclave. 1. Select the Launch enclave public key hash option. 2. To use one of the public keys configured on the host, select Use from host and from the drop-down menu, select a public key hash. 3. To enter the public key manually, select Enter manually and enter a valid SHA256 hash (64) characters key.

8. Click **OK**.

Remove vSGX from a Virtual Machine

You can remove vSGX from a virtual machine.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. In the **Edit Settings** dialog box, under **Security devices**, deselect the **Enable** check box for SGX.
4. Click **OK**.

Verify that the vSGX entry no longer appears in the virtual machine **Summary** tab in the **VM Hardware** pane.

Securing Virtual Machines with AMD Secure Encrypted Virtualization-Encrypted State

Secure Encrypted Virtualization-Encrypted State (SEV-ES) is a hardware feature enabled in recent AMD CPUs that keeps the guest operating system's memory and register state encrypted, protecting it against access from the hypervisor.

You can add SEV-ES to your virtual machines as an extra security enhancement. SEV-ES prevents CPU registers from leaking information in registers to components like the hypervisor. SEV-ES can also detect malicious modifications to a CPU register state.

vSphere and AMD Secure Encrypted Virtualization-Encrypted State

In vSphere 7.0 Update 1 and later, you can activate Secure Encrypted Virtualization-Encrypted State (SEV-ES) on supported AMD CPUs and guest operating systems.

Currently, SEV-ES supports only AMD EPYC 7xx2 CPUs (code named "Rome") and later CPUs, and only versions of Linux kernels that include specific support for SEV-ES.

SEV-ES Components and Architecture

The SEV-ES architecture consists of the following components.

- AMD CPU, specifically, the Platform Security Processor (PSP) that manages encryption keys and handles encryption.
- Enlightened operating system, that is, an operating system that uses guest-initiated calls to the hypervisor.
- Virtual Machine Monitor (VMM) and Virtual Machine Executable (VMX), to initialize an encrypted virtual machine state during virtual machine power-on, and also to handle calls from the guest operating system.
- VMkernel driver, to communicate unencrypted data between the hypervisor and the guest operating system.

Implementing and Managing SEV-ES on ESXi

You must first activate SEV-ES in a system's BIOS configuration. See the documentation for your system for more information about accessing the BIOS configuration. After you have activated SEV-ES in the BIOS for your system, you can then add SEV-ES to a virtual machine.

You use either the vSphere Client (in vSphere 7.0 Update 2 and later) or PowerCLI commands to activate and deactivate SEV-ES on virtual machines. You can create new virtual machines with SEV-ES, or activate SEV-ES on existing virtual machines. Privileges to manage virtual machines activated with SEV-ES are the same as for managing regular virtual machines.

Unsupported VMware Features on SEV-ES

The following features are not supported when SEV-ES is activated.

- System Management Mode
- vMotion
- Powered-on snapshots (however, no-memory snapshots are supported)
- Hot add or remove of CPU or memory
- Suspend/resume
- VMware Fault Tolerance
- Clones and instant clones
- Guest Integrity
- UEFI Secure Boot

Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the vSphere Client

In vSphere 7.0 Update 2 and later, you can use the vSphere Client to add SEV-ES to a virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.
- SEV-ES must be enabled in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When enabling SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.
Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES enabled virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES enabled virtual machines per ESXi host.
- The ESXi host running in your environment must be at ESXi 7.0 Update 1 or later.

- The vCenter Server must be at vSphere 7.0 Update 2 or later.
- The guest operating system must support SEV-ES.
Currently, only Linux kernels with specific support for SEV-ES are supported.
- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option enabled, otherwise power-on fails.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, an ESXi host or a cluster.
3. Right-click the object, select **New Virtual Machine**, and follow the prompts to create a virtual machine.

Option	Action
Select a creation type	Create a virtual machine.
Select a name and folder	Specify a name and target location.
Select a compute resource	Specify an object for which you have privileges to create virtual machines.
Select storage	In the VM storage policy, select the storage policy. Select a compatible datastore.
Select compatibility	Ensure that ESXi 7.0 and later is selected.
Select a guest OS	Select Linux, and select a version of Linux with specific support for SEV-ES.
Customize hardware	Under VM Options > Boot Options > Firmware , ensure that EFI is selected. Under VM Options > Encryption , select the Enable check box for AMD SEV-ES.
Ready to complete	Review the information and click Finish .

The virtual machine is created with SEV-ES.

Add AMD Secure Encrypted Virtualization-Encrypted State to a Virtual Machine Using the Command Line

You can use the command line to add SEV-ES to a virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.
- SEV-ES must be enabled in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When enabling SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.

Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES enabled virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES enabled virtual machines per ESXi host.

- The ESXi host running in your environment must be at ESXi 7.0 Update 1 or later.
- The guest operating system must support SEV-ES.
Currently, only Linux kernels with specific support for SEV-ES are supported.
- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option enabled, otherwise power-on fails.
- PowerCLI 12.1.0 or later must be installed on a system with access to your environment.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator to the vCenter Server that manages the ESXi host on which you want to add a virtual machine with SEV-ES.

```
Connect-VIServer -server vCenter_Server_ip_address -User admin_user -Password 'password'
```

2. Create the virtual machine with the `New-VM` cmdlet, specifying `-SEVEnabled $true`.

For example, first assign the host information to a variable, then create the virtual machine.

```
$vmhost = Get-VMHost -Name 10.193.25.83
New-VM -Name MyVM1 $vmhost -NumCPU 2 -MemoryMB 4 -DiskMB 4 -SEVEnabled $true
```

If you must specify the virtual hardware version, run the `New-VM` cmdlet with the `-HardwareVersion vmx-18` parameter. For example:

```
New-VM -Name MyVM1 $vmhost -NumCPU 2 -MemoryMB 4 -DiskMB 4 -SEVEnabled $true -HardwareVersion vmx-18
```

The virtual machine is created with SEV-ES.

Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the vSphere Client

In vSphere 7.0 Update 2 and later, you can use the vSphere Client to add SEV-ES to an existing virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.
- SEV-ES must be activated in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When activating SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.

Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES activated virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES activated virtual machines per ESXi host.

- The ESXi host running in your environment must be at ESXi 7.0 Update 1 or later.
- The vCenter Server must be at vSphere 7.0 Update 2 or later.
- The guest operating system must support SEV-ES.
Currently, only Linux kernels with specific support for SEV-ES are supported.
- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option checked, otherwise power-on fails.
- Ensure that the virtual machine is powered off.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. Under **VM Options > Boot Options > Firmware**, ensure that EFI is selected.
4. In the **Edit Settings** dialog box, under **VM Options > Encryption**, select the **Enable** check box for AMD SEV-ES.
5. Click **OK**.

SEV-ES is added to the virtual machine.

Activate AMD Secure Encrypted Virtualization-Encrypted State on an Existing Virtual Machine Using the Command Line

You can use the command line to add SEV-ES to an existing virtual machine to provide enhanced security to the guest operating system.

- The system must be installed with an AMD EPYC 7xx2 (code named "Rome") or later CPU and supporting BIOS.

- SEV-ES must be activated in the BIOS.
- The number of SEV-ES virtual machines per ESXi host is controlled by the BIOS. When activating SEV-ES in the BIOS, enter a value for the **Minimum SEV non-ES ASID** setting equal to the number of SEV-ES virtual machines plus one. For example, if you have 12 virtual machines that you want to run concurrently, enter 13.
Note: vSphere 7.0 Update 1 and later supports 16 SEV-ES activated virtual machines per ESXi host. Using a higher setting in the BIOS does not prevent SEV-ES from working, however, the limit of 16 still applies. vSphere 7.0 Update 2 and later supports 480 SEV-ES activated virtual machines per ESXi host.
- The ESXi host running in your environment must be ESXi 7.0 Update 1 or later.
- The guest operating system must support SEV-ES.
Currently, only Linux kernels with specific support for SEV-ES are supported.
- The virtual machine must be at hardware version 18 or later.
- The virtual machine must have the **Reserve all guest memory** option checked, otherwise power-on fails.
- PowerCLI 12.1.0 or later must be installed on a system with access to your environment.
- Ensure that the virtual machine is powered off.

You can add SEV-ES to virtual machines running on ESXi 7.0 Update 1 or later.

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator to the vCenter Server that manages the ESXi host with the virtual machine to which you want to add SEV-ES.

For example:

```
Connect-VIServer -server vCenter_Server_ip_address -User admin_user -Password 'password'
```

2. Add SEV-ES to the virtual machine with the `Set-VM` cmdlet, specifying `-SEVEnabled $true`.

For example:

```
$vmhost = Get-VMHost -Name 10.193.25.83
Set-VM -Name MyVM2 $vmhost -SEVEnabled $true
```

If you must specify the virtual hardware version, run the `Set-VM` cmdlet with the `-HardwareVersion vmx-18` parameter. For example:

```
Set-VM -Name MyVM2 $vmhost -SEVEnabled $true -HardwareVersion vmx-18
```

SEV-ES is added to the virtual machine.

Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the vSphere Client

In vSphere 7.0 Update 2 and later, you can use the vSphere Client to deactivate SEV-ES on a virtual machine.

- Ensure that the virtual machine is powered off.
1. Connect to vCenter Server by using the vSphere Client.
 2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
 3. In the **Edit Settings** dialog box, under **VM Options > Encryption**, deselect the **Enable** check box for AMD SEV-ES.
 4. Click **OK**.

SEV-ES is deactivated on the virtual machine.

Deactivate AMD Secure Encrypted Virtualization-Encrypted State on a Virtual Machine Using the Command Line

You can use the command line to deactivate SEV-ES on a virtual machine.

- Ensure that the virtual machine is powered off.
 - PowerCLI 12.1.0 or later must be installed on a system that has access to your environment.
1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator to the vCenter Server that manages the ESXi host with the virtual machine from which you want to remove SEV-ES.

For example:

```
Connect-VIServer -server vCenter_Server_ip_address -User admin_user -Password 'password'
```

2. Deactivate SEV-ES on the virtual machine with the `Set-VM` cmdlet, specifying `-SEVEnabled $false`.

For example, first assign the host information to a variable, then deactivate SEV-ES for the virtual machine.

```
$vmhost = Get-VMHost -Name 10.193.25.83
Set-VM -Name MyVM2 $vmhost -SEVEnabled $false
```

SEV-ES is deactivated on the virtual machine.

Virtual Machine Encryption

With vSphere Virtual Machine Encryption, you can encrypt your sensitive workloads in an even more secure way. Access to encryption keys can be made conditional to the ESXi host being in a trusted state.

Before you can start with virtual machine encryption tasks, you must set up a key provider. The following key provider types are available.

Table 249: vSphere Key Providers

Key Provider	Description	For More Information
Standard key provider	Available in vSphere 6.5 and later, the standard key provider uses vCenter Server to request keys from an external key server. The key server generates and stores the keys, and passes them to vCenter Server for distribution.	See Configuring and Managing a Standard Key Provider .
Trusted key provider	Available in vSphere 7.0 and later, the vSphere Trust Authority trusted key provider makes access to the encryption keys conditional to the attestation state of a workload cluster. vSphere Trust Authority requires an external key server.	See vSphere Trust Authority .
VMware vSphere® Native Key Provider™	Available in vSphere 7.0 Update 2 and later, vSphere Native Key Provider is included in all vSphere editions and does not require an external key server.	See Configuring and Managing vSphere Native Key Provider .

Comparison of vSphere Key Providers

A high-level overview of the capabilities of the vSphere key providers requires your attention to help plan your encryption strategy.

In general, there is little difference in feature or product support between key provider daily operation. Though key providers look and behave similarly, you might have requirements and regulations to consider when choosing a key provider, as shown in the following table.

Table 250: Key Provider Considerations

Key Provider	External Key Server Required?	Quick Setup?	Works Only with vSphere?	Encryption Keys Stored Permanently on Host?	Rekey While Cloning?
Standard key provider	Yes	No	No	No	Yes
Trusted key provider	Yes	No	No	No	Yes
vSphere Native Key Provider	No	Yes	Yes	Yes	Yes

Note: At host startup, vSphere Native Key Provider always writes the encryption key to the ESXi hosts in the cluster. If you are concerned about the physical security of the cluster, consider using either a standard key provider or trusted key provider, both of which require that the key server be available for encrypted virtual machines to function.

Key Provider Encryption Features

The following encryption features are supported by each key provider type.

- Rekey using the same key provider or to another key provider
- Rotate keys
- Virtual Trusted Platform Module (vTPM)
- Disk encryption
- vSphere Virtual Machine Encryption
- Co-existence with other key providers
- Upgrade to a different key provider

Key Provider Support for vSphere Features

The following describes key provider support for some important vSphere features.

- Encrypted vSphere vMotion: Supported by all key provider types. The same key provider must be available on the destination host. See [What Is Encrypted vSphere vMotion](#).
- vCenter Server File-based Backup and Restore: Standard key provider and vSphere Native Key Provider support vCenter Server file-based backup and restore. Because most vSphere Trust Authority configuration information is stored on the ESXi hosts, the vCenter Server file-based backup mechanism does not back up this information. To ensure the configuration information for your vSphere Trust Authority deployment is saved, see [Backing Up the vSphere Trust Authority Configuration](#).

Key Provider Support for VMware Products

The following table compares key provider support for some VMware products.

Table 251: Comparison of Support for VMware Products

Key Provider	vSAN Data-At-Rest Encryption	Site Recovery Manager	vSphere Replication
Standard key provider	Yes	Yes	Yes
Trusted key provider	No	Yes	No

Key Provider	vSAN Data-At-Rest Encryption	Site Recovery Manager	vSphere Replication
		If the same vSphere Trust Authority services configuration is available on the recovery side, then SRM with array-based replication is supported.	
vSphere Native Key Provider	Yes	Yes	Yes

Note:

Standard key provider, trusted key provider, and vSphere Native Key Provider support vSphere Virtual Machine Encryption on top of vSAN.

Required Hardware for Key Providers

The following table compares some minimum key provider hardware requirements.

Table 252: Comparison of Required Hardware for Key Providers

Key Provider	TPM on ESXi Host
Standard key provider	Not required
Trusted key provider	Required on Trusted Hosts (hosts in the Trusted Cluster). Note: Currently, the ESXi hosts in the Trust Authority Cluster do not require a TPM. However, as a matter of best practice, consider installing new ESXi hosts with TPMs.
vSphere Native Key Provider	Not required vSphere Native Key Provider availability can optionally be restricted to hosts with a TPM.

Key Provider Naming

vSphere uses a key provider name to look up a key identifier. If two key providers have the same name, vSphere assumes that they are equivalent and have access to the same keys. Each logical key provider, regardless of its type (Standard, Trusted, and Native Key Provider), must have a unique name across all vCenter Server systems.

In a few instances, you configure the same key provider across multiple vCenter Server systems, such as:

- Migrating encrypted virtual machines between vCenter Server systems
- Setting up a vCenter Server as a disaster recovery site

How vSphere Virtual Machine Encryption Protects Your Environment

Regardless of which key provider you use, with vSphere Virtual Machine Encryption you can create encrypted virtual machines and encrypt existing virtual machines. Because all virtual machine files with sensitive information are encrypted, the virtual machine is protected. Only administrators with encryption privileges can perform encryption and decryption tasks.

Important: ESXi Shell users also have cryptographic operation privileges. For more information, see [Prerequisites and Required Privileges for Virtual Machine Encryption Tasks](#).

What Storage Does vSphere Virtual Machine Encryption Support

vSphere Virtual Machine Encryption works with any supported storage type (NFS, iSCSI, Fibre Channel, direct-attached storage, and so on), including VMware vSAN. For more information about using encryption on a vSAN cluster, see the *Administering VMware vSAN* documentation.

vSphere Virtual Machine Encryption and vSAN use the same encryption libraries but they have different profiles. Virtual Machine Encryption is a per-VM encryption and vSAN is a datastore level encryption.

vSphere Encryption Keys and Key Providers

vSphere uses two levels of encryption in the form of a Key Encryption Key (KEK) and a Data Encryption Key (DEK). Briefly, an ESXi host generates a DEK to encrypt virtual machines and disks. The KEK is provided by a key server, and encrypts (or "wraps") the DEK. The KEK encrypts the DEK using the AES256 algorithm and the DEK encrypts the VMDK using the XTS-AES-256 (512-bit key size) algorithm. Depending on the type of key provider, different methods are used to create and manage the DEK and KEK.

Standard key provider operates as follows.

1. The ESXi host generates and uses internal keys to encrypt virtual machines and disks. These keys are used as DEKs.
2. vCenter Server requests keys from the key server (KMS). These keys are used as the KEKs. vCenter Server stores only the ID of each KEK, but not the key itself.
3. ESXi uses the KEK to encrypt the internal keys, and stores the encrypted internal key on disk. ESXi does not store the KEK on disk. If a host reboots, vCenter Server requests the KEK with the corresponding ID from the key server and makes it available to ESXi. ESXi can then decrypt the internal keys as needed.

vSphere Trust Authority trusted key provider operates as follows.

1. The vCenter Server of the Trusted Cluster checks if the default trusted key provider is accessible to the ESXi host where the encrypted virtual machine is to be created.
2. The vCenter Server of the Trusted Cluster adds the trusted key provider to the virtual machine ConfigSpec.
3. The virtual machine creation request is sent to the ESXi host.
4. If an attestation token is not already available to the ESXi host, it requests one from the Attestation Service.
5. The Key Provider Service validates the attestation token and creates a KEK to be sent to the ESXi host. The KEK is wrapped (encrypted) with the primary key that is configured on the key provider. Both the KEK ciphertext and KEK plaintext are returned to the Trusted Host.
6. The ESXi host generates a DEK to encrypt the virtual machine disks.
7. The KEK is used to wrap the DEK generated by the ESXi host, and the ciphertext from the key provider is stored alongside the encrypted data.
8. The virtual machine is encrypted and written to storage.

Note:

The ESXi hosts in vSphere clusters hold the KEK for encrypted virtual machines in host memory to enable availability features such as High Availability, vMotion, DRS, and so on. When a virtual machine is deleted or unregistered, the ESXi hosts in the cluster delete the KEK from their memory. Thus, the ESXi hosts can no longer use the KEK. This behavior is the same for standard key providers and trusted key providers.

vSphere Native Key Provider operates as follows.

1. When you create the key provider, vCenter Server generates a primary key and pushes it to the ESXi hosts in the cluster. (No external key server is involved.)
2. The ESXi hosts generate a DEK on demand.
3. When you perform an encryption activity, data is encrypted with the DEK.
Encrypted DEKs are stored alongside the encrypted data.
4. When you decrypt data, the primary key is used to decrypt the DEK, and then the data.

What Components Does vSphere Virtual Machine Encryption Encrypt

vSphere Virtual Machine Encryption supports encryption of virtual machine files, virtual disk files, and core dump files.

Virtual machine files

Most virtual machine files, in particular, guest data that is not stored in the VMDK file, are encrypted. This set of files includes but is not limited to the NVRAM, VSWP, and VMSN files. The key from the key provider unlocks an encrypted bundle in the VMX file that contains internal keys and other secrets. The key retrieval works as follows, depending on the key provider:

- Standard key provider: vCenter Server manages the keys from the key server and the ESXi hosts cannot directly access the key provider. The hosts wait for vCenter Server to push the keys.
- Trusted key provider and vSphere Native Key Provider: The ESXi hosts directly access the key providers, and so fetch the requested keys either from the vSphere Trust Authority service directly or the vSphere Native Key Provider.

When you use the vSphere Client to create an encrypted virtual machine, you can encrypt and decrypt virtual disks separate from virtual machine files. All virtual disks are encrypted by default. For other encryption tasks, such as encrypting an existing virtual machine, you can encrypt and decrypt virtual disks separate from virtual machine files.

Note: You cannot associate an encrypted virtual disk with a virtual machine that is not encrypted.

Virtual disk files

Data in an encrypted virtual disk (VMDK) file is never written in cleartext to storage or physical disk, and is never transmitted over the network in cleartext. The VMDK descriptor file is mostly cleartext, but contains a key ID for the KEK and the internal key (DEK) in the encrypted bundle.

You can use the vSphere Client or the vSphere API to perform a shallow reencrypt operation with a new KEK, or use the vSphere API to perform a deep reencrypt operation with a new internal key.

Core dumps

Core dumps on an ESXi host that has encryption mode enabled are always encrypted. See [vSphere Virtual Machine Encryption and Core Dumps](#). Core dumps on the vCenter Server system are not encrypted. Protect access to the vCenter Server system.

Virtual machine swap file

The virtual machine swap file is encrypted whenever you add a vTPM to a virtual machine. Environments low on RAM can experience encryption-related paging that could impact performance.

vTPMs

When you configure a vTPM, the virtual machine files are encrypted but not the disks. You can choose to add encryption explicitly for the virtual machine and its disks. For more information, see [Securing Virtual Machines with Virtual Trusted Platform Module](#).

Note: For information on some limitations concerning devices and features that vSphere Virtual Machine Encryption can interoperate with, see [Virtual Machine Encryption Interoperability](#).

What Components Does vSphere Virtual Machine Encryption Not Encrypt

Some of the files that are associated with a virtual machine are not encrypted or partially encrypted.

Log files	Log files are not encrypted because they do not contain sensitive data.
Virtual machine configuration files	Most of the virtual machine configuration information, stored in the VMX and VMSS files, is not encrypted.
Virtual disk descriptor file	To support disk management without a key, most of the virtual disk descriptor file is not encrypted.

What Privileges Are Required to Perform Cryptographic Operations

Only users that are assigned the **Cryptographic Operations** privileges can perform cryptographic operations. The privilege set is fine grained. The default Administrator system role includes all **Cryptographic Operations** privileges. The No Cryptography Administrator role supports all Administrator privileges except for the **Cryptographic Operations** privileges.

In addition to using the **Cryptographer.*** privileges, vSphere Native Key Provider can use the **Cryptographer.ReadKeyServersInfo** privilege, which is specific to vSphere Native Key Providers.

See [Cryptographic Operations Privileges](#) for more information.

You can create additional custom roles, for example, to allow a group of users to encrypt virtual machines but to prevent them from decrypting virtual machines.

How Do You Perform Cryptographic Operations

The vSphere Client supports many of the cryptographic operations. For other tasks, you can use PowerCLI or the vSphere API.

Table 253: Interfaces for Performing Cryptographic Operations

Interface	Operations	Information
vSphere Client	Create encrypted virtual machine Encrypt and decrypt virtual machines Perform a shallow reencrypt of a virtual machine (use a different KEK)	This book
PowerCLI	Create encrypted virtual machine Encrypt and decrypt virtual machines Configure vSphere Trust Authority	<i>VMware PowerCLI Cmdlets Reference</i>
vSphere Web Services SDK	Create encrypted virtual machine Encrypt and decrypt virtual machines Perform a deep reencrypt of a virtual machine (use a different DEK) Perform a shallow reencrypt of a virtual machine (use a different KEK)	<i>vSphere Web Services SDK Programming Guide</i> <i>vSphere Web Services API Reference</i>
crypto-util	Decrypt encrypted core dumps Check whether files are encrypted Perform other management tasks directly on the ESXi host	Command-line help vSphere Virtual Machine Encryption and Core Dumps

How Do You Recrypt (Rekey) an Encrypted Virtual Machine

You can recrypt (also called rekey) a virtual machine with new keys, for example, in case a key expires or becomes compromised. The following rekeying options are available.

- A shallow recrypt, which replaces only the Key Encryption Key (KEK)
- A deep recrypt, which replaces both the Disk Encryption Key (DEK) and the KEK

A deep recrypt requires that the virtual machine is powered off and contains no snapshots. You can perform a shallow recrypt operation while the virtual machine is powered on, and if the virtual machine has snapshots present. Shallow recrypt of an encrypted virtual machine with snapshots is permitted only on a single snapshot branch (disk chain). Multiple snapshot branches are not supported. Also, shallow recrypt is not supported on a linked clone of a virtual machine or disk. If the shallow recrypt fails before updating all links in the chain with the new KEK, you can still access the encrypted virtual machine if you have the old and new KEKs. However, it is best to reissue the shallow recrypt operation before performing any snapshot operations.

You can perform a rekey of a virtual machine by using the vSphere Client, the CLI, or the API. See [Rekey an Encrypted Virtual Machine Using the vSphere Client](#), [Rekey an Encrypted Virtual Machine Using the CLI](#), and [vSphere Web Services SDK Programming Guide](#).

vSphere Virtual Machine Encryption Components

Depending on which key provider you use, an external key server, the vCenter Server system, and your ESXi hosts are potentially contributing to the encryption solution.

The following components comprise vSphere Virtual Machine Encryption:

- An external key server, also called a KMS (not required for vSphere Native Key Provider)
- vCenter Server
- ESXi hosts

What Is the Role of a Key Server in vSphere Virtual Machine Encryption

The key server is a Key Management Interoperability Protocol (KMIP) management server that is associated with a key provider. A standard key provider and a trusted key provider require a key server. vSphere Native Key Provider does not require a key server. The following table describes the differences in key provider and key server interaction.

Table 254: Key Providers and Key Server Interaction

Key Provider	Interaction with Key Server
Standard key provider	A standard key provider uses vCenter Server to request keys from a key server. The key server generates and stores the keys, and passes them to vCenter Server for distribution to the ESXi hosts.
Trusted key provider	A trusted key provider uses a Key Provider Service that enables the trusted ESXi hosts to fetch the keys directly. See What Is the vSphere Trust Authority Key Provider Service .
vSphere Native Key Provider	vSphere Native Key Provider does not require a key server. vCenter Server generates a primary key and pushes it to the ESXi hosts. The ESXi hosts then generate data encryption keys (even when not connected to vCenter Server). See vSphere Native Key Provider Overview .

You can use the vSphere Client or the vSphere API to add key provider instances to the vCenter Server system. If you use multiple key provider instances, all instances must be from the same vendor and must replicate keys.

If your environment uses different key server vendors in different environments, you can add a key provider for each key server and specify a default key provider. The first key provider that you add becomes the default key provider. You can explicitly specify the default later.

As a KMIP client, vCenter Server uses the Key Management Interoperability Protocol (KMIP) to make it easy to use the key server of your choice.

What Is the Role of vCenter Server in vSphere Virtual Machine Encryption

The following table describes the role of vCenter Server in the encryption process.

Table 255: Key Providers and vCenter Server

Key Provider	Role of vCenter Server	How Are Privileges Checked
Standard key provider	Only vCenter Server has the credentials for logging in to the key server. Your ESXi hosts do not have those credentials. vCenter Server obtains keys from the key server and pushes them to the ESXi hosts. vCenter Server does not store the key server keys, but keeps a list of key IDs.	vCenter Server checks the privileges of users who perform cryptographic operations.
Trusted key provider	vSphere Trust Authority removes the need for vCenter Server to request keys from the key server, and makes access to the encryption keys conditional to the attestation state of a workload cluster. You must use separate vCenter Server systems for the Trusted Cluster and Trust Authority Cluster.	vCenter Server checks the privileges of users who perform cryptographic operations. Only users who are members of the TrustedAdmins SSO group can perform administrative operations.
vSphere Native Key Provider	The vCenter Server generates the keys.	vCenter Server checks the privileges of users who perform cryptographic operations.

You can use the vSphere Client to assign cryptographic operation privileges or to assign the **No cryptography administrator** custom role to groups of users. See [Prerequisites and Required Privileges for Virtual Machine Encryption Tasks](#).

vCenter Server adds cryptography events to the list of events that you can view and export from the vSphere Client Event Console. Each event includes the user, time, key ID, and cryptographic operation.

The keys that come from the key server are used as key encryption keys (KEKs).

What Is the Role of ESXi Hosts in vSphere Virtual Machine Encryption

ESXi hosts are responsible for several aspects of the encryption workflow.

Table 256: Key Providers and ESXi Hosts

Key Provider	ESXi Host Aspects
Standard key provider	vCenter Server pushes keys to an ESXi host when the host needs a key. The host must have encryption mode enabled.

Key Provider	ESXi Host Aspects
	- Ensuring that guest data for encrypted virtual machines is encrypted when stored on disk. - Ensuring that guest data for encrypted virtual machines is not sent over the network without encryption.
Trusted key provider	The ESXi hosts run vSphere Trust Authority services, depending on if they are Trusted Hosts or Trust Authority Hosts. Trusted ESXi hosts run workload virtual machines that can be encrypted using key providers published by the Trust Authority Hosts. See vSphere Trust Authority Trusted Infrastructure .
vSphere Native Key Provider	The ESXi hosts fetch keys directly from the vSphere Native Key Provider.

The keys that the ESXi host generates are called internal keys in this document. These keys typically act as data encryption keys (DEKs).

Encryption Process Flow

After you set up a key provider, users with the required privileges can create encrypted virtual machines and disks. Those users can also encrypt existing virtual machines and decrypt encrypted virtual machines, and add Virtual Trusted Platform Modules (vTPMs) to virtual machines.

Depending on the key provider type, the process flow can involve a key server, the vCenter Server, and the ESXi host.

Standard Key Provider Encryption Process Flow

During the encryption process, different vSphere components interact as follows.

1. When the user performs an encryption task, for example, creating an encrypted virtual machine, vCenter Server requests a new key from the default key server. This key is used as the KEK.
2. vCenter Server stores the key ID and passes the key to the ESXi host. If the ESXi host is part of a cluster, vCenter Server sends the KEK to each host in the cluster.
The key itself is not stored on the vCenter Server system. Only the key ID is known.
3. The ESXi host generates internal keys (DEKs) for the virtual machine and its disks. It keeps the internal keys in memory only, and uses the KEKs to encrypt internal keys.
Unencrypted internal keys are never stored on disk. Only encrypted data is stored. Because the KEKs come from the key server, the host continues to use the same KEKs.
4. The ESXi host encrypts the virtual machine with the encrypted internal key.
Any hosts that have the KEK and that can access the encrypted key file can perform operations on the encrypted virtual machine or disk.

Trusted Key Provider Encryption Process Flow

The vSphere Trust Authority encryption process flow includes the vSphere Trust Authority services, the trusted key providers, the vCenter Server, and the ESXi hosts.

Encrypting a virtual machine with a trusted key provider looks the same as the virtual machine encryption user experience when using a standard key provider. Virtual machine encryption under vSphere Trust Authority continues to rely on either virtual machine encryption storage policies, or the presence of a vTPM device, to decide when to encrypt a virtual machine. You still use a default configured key provider (called a KMS cluster in vSphere 6.5 and 6.7) when encrypting a virtual machine from the vSphere Client. And, you can still use the APIs in a similar way to specify the key provider manually. The existing Cryptographic privileges added for vSphere 6.5 are still relevant in vSphere 7.0 and later for vSphere Trust Authority.

The encryption process for the trusted key provider has some important differences from the standard key provider:

- Trust Authority administrators do not specify information directly when setting up a key server for a vCenter Server instance, and they do not establish the key server trust. Instead, vSphere Trust Authority publishes trusted key providers that the Trusted Hosts can use.
- vCenter Server no longer pushes keys to ESXi hosts and instead it can treat each trusted key provider as a single top-level key.
- Only Trusted Hosts can request encryption operations from Trust Authority Hosts.

vSphere Native Key Provider Encryption Process Flow

vSphere Native Key Provider is included in vSphere 7.0 Update 2 and later. When you configure a vSphere Native Key Provider, vCenter Server pushes a primary key to all ESXi hosts in the cluster. Likewise, if you update or delete a vSphere Native Key Provider, the change is pushed to the hosts in the cluster. The encryption process flow is similar to how a trusted key provider works. The difference is that vSphere Native Key Provider generates the keys and wraps them with the primary key, then hands them back to perform encryption.

Custom Attributes for Key Servers

The Key Management Interoperability Protocol (KMIP) supports adding custom attributes intended for vendor-specific purposes. Custom attributes enable you to more specifically identify keys stored in your key server. vCenter Server adds the following custom attributes for virtual machine keys and host keys.

Table 257: Virtual Machine Encryption Custom Attributes

Custom Attribute	Value
x-Vendor	VMware, Inc.
x-Product	VMware vSphere
x-Product_Version	vCenter Server version
x-Component	Virtual Machine
x-Name	Virtual machine name (gathered from ConfigInfo or ConfigSpec)
x-Identifier	Virtual machine's instanceUuid (gathered from ConfigInfo or ConfigSpec)

Table 258: Host Encryption Custom Attributes

Custom Attribute	Value
x-Vendor	VMware, Inc.
x-Product	VMware vSphere
x-Product_Version	vCenter Server version
x-Component	ESXi Server
x-Name	Host name
x-Identifier	Host's hardware Uuid

vCenter Server adds the `x-Vendor`, `x-Product`, and `x-Product_Version` attributes when the key server creates a key. When the key is used to encrypt a virtual machine or host, vCenter Server sets the `x-Component`, `x-Identifier`, and `x-Name` attributes. You might be able to view these custom attributes in your key server user interface. Check with your key server vendor.

Both the host key and virtual machine key have the six custom attributes. `x-Vendor`, `x-Product`, and `x-Product_Version` might be the same for both keys. These attributes are set when the key is generated. Depending on if the key is for a virtual machine or a host, it might have `x-Component`, `x-Identifier`, and `x-Name` attributes appended.

Encryption Key Errors

When an error occurs sending keys from the key server to an ESXi host, vCenter Server generates a message in the event log for the following events:

- Adding keys to the ESXi host failed due to host connection or host support issues.
- Getting keys from the key server failed due to key missing in the key server.
- Getting keys from the key server failed due to the key server connection.

Decrypting Encrypted Virtual Machines

If you later want to decrypt an encrypted virtual machine, you change its storage policy. You can change the storage policy for the virtual machine and all disks. If you want to decrypt individual components, decrypt selected disks first, then decrypt the virtual machine by changing the storage policy for VM Home. Both keys are required for decryption of each component. See [Decrypt an Encrypted Virtual Machine or Virtual Disk](#).

Virtual Disk Encryption

When you create an encrypted virtual machine from the vSphere Client, you can decide which disks to exclude from encryption. You can later add disks and set their encryption policies. You cannot add an encrypted disk to a virtual machine that is not encrypted, and you cannot encrypt a disk if the virtual machine is not encrypted.

Encryption for a virtual machine and its disks is controlled through storage policies. The storage policy for VM Home governs the virtual machine itself, and each virtual disk has an associated storage policy.

- Setting the storage policy of VM Home to an encryption policy encrypts only the virtual machine itself.
- Setting the storage policy of VM Home and all the disks to an encryption policy encrypts all components.

Consider the following use cases.

Table 259: Virtual Disk Encryption Use Cases

Use Case	Details
Create an encrypted virtual machine.	If you add disks while creating an encrypted virtual machine, the disks are encrypted by default. You can change the policy not to encrypt one or more of the disks. After virtual machine creation, you can explicitly change the storage policy for each disk. See Change the Encryption Policy for Virtual Disks .
Encrypt a virtual machine.	To encrypt an existing virtual machine, you change its storage policy. You can change the storage policy for the virtual machine and all virtual disks. To encrypt just the virtual machine, you can specify an encryption policy for VM Home and select a different storage policy, such as Datastore Default, for each virtual disk. See Create an Encrypted Virtual Machine .
Add an existing unencrypted disk to an encrypted virtual machine (encryption storage policy).	Fails with an error. You have to add the disk with the default storage policy, but can later change the storage policy. See Change the Encryption Policy for Virtual Disks .

Use Case	Details
Add an existing unencrypted disk to an encrypted virtual machine with a storage policy that does not include encryption, for example Datastore Default.	The disk uses the default storage policy. You can explicitly change the storage policy after adding the disk if you want an encrypted disk. See Change the Encryption Policy for Virtual Disks .
Add an encrypted disk to an encrypted virtual machine. VM Home storage policy is Encryption.	When you add the disk, it remains encrypted. The vSphere Client displays the size and other attributes, including encryption status.
Add an existing encrypted disk to an unencrypted virtual machine.	This use case is not supported. However, if you use the vSphere Client to encrypt the VM Home files, you can then reconfigure the unencrypted virtual machine with the encrypted disk.
Register an encrypted virtual machine.	If you remove an encrypted virtual machine from vCenter Server but do not delete it from disk, you can return it to the vCenter Server inventory by registering the VM's virtual machine configuration (.vmx) file. To register the encrypted VM, the user must have the Cryptographic operations > Register VM privilege. If the VM was encrypted using a standard key provider, when the encrypted VM is registered, vCenter Server pushes the required keys to the ESXi host. If the user registering the VM does not have the Cryptographic operations > Register VM privilege, vCenter Server locks the VM upon registration, and the VM is not usable until it is unlocked. If the VM was encrypted using a trusted key provider or vSphere Native Key Provider, when the encrypted VM is registered, vCenter Server no longer pushes keys to the ESXi host. Instead, the keys are fetched from the host when the VM is registered. If the user registering the VM does not have the Cryptographic operations > Register VM privilege, vCenter Server does not permit the operation.

Virtual Machine Encryption Errors

If vCenter Server detects a critical error with virtual machine encryption, it creates an event. You can view these events to help troubleshoot and resolve encryption errors.

vCenter Server creates events for the following virtual machine encryption critical errors.

- Failure to generate a KEK.
- Insufficient disk space on the datastore to create an encrypted virtual machine.
- Insufficient user privilege to initiate encryption operation.
- The specified key is missing on the key provider and so the ESXi host key is renewed with a new key.
- An error occurred on the key provider with the specified key and so the ESXi host key is renewed with a new key.

Prerequisites and Required Privileges for Virtual Machine Encryption Tasks

Virtual machine encryption tasks are possible only in environments that include vCenter Server. Also, the ESXi host must have encryption mode activated for most encryption tasks. The user who performs the task must have the appropriate privileges. A set of **Cryptographic Operations** privileges allows fine-grained control. If virtual machine encryption tasks require a change to the host encryption mode, additional privileges are required.

Note: vSphere Trust Authority has additional prerequisites and required privileges. See [Prerequisites and Required Privileges for vSphere Trust Authority](#).

Using Cryptography Privileges and Roles

By default, the user with the vCenter Server Administrator role has all privileges, including cryptographic operations privileges. The **No cryptography administrator** role does not have the following privileges that are required for cryptographic operations.

Important: ESXi Shell users also have cryptographic operation privileges.

- **AddCryptographic Operations** privileges.
- **Global > Diagnostics**
- **Host > Inventory > Add host to cluster**
- **Host > Inventory > Add standalone host**
- **Host > Local operations > Manage user groups**

You can assign the **No cryptography administrator** role to vCenter Server administrators that do not need **Cryptographic Operations** privileges.

To impose more limits on what users can do, you can clone the **No cryptography administrator** role and create a custom role with only some of the **Cryptographic Operations** privileges. For example, you can create a role that allows users to encrypt but not to decrypt virtual machines. See [using-roles-to-assign-privileges.dita](#).

What Is Host Encryption Mode

Host encryption mode determines if an ESXi host is ready to accept cryptographic material for encrypting virtual machines and virtual disks. Before any cryptographic operations can occur on a host, encryption mode must be activated. Host encryption mode is often set automatically when it is required, but you can set it explicitly. You can check and explicitly set the current host encryption mode from the vSphere Client or by using the vSphere API.

When host encryption mode is activated, vCenter Server installs a host key on the host, which ensures that the host is cryptographically "safe." With the host key in place, other cryptographic operations can proceed, including vCenter Server obtaining keys from the key provider and pushing them to the ESXi hosts.

In "safe" mode, user worlds (that is, hostd) and encrypted virtual machines have their core dumps encrypted. Unencrypted virtual machines do not have their core dumps encrypted.

For more information about encrypted core dumps and how they are used by VMware Technical Support, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2147388>.

For instructions, see [Activate Host Encryption Mode Explicitly](#).

After host encryption mode is set, it cannot be deactivated easily. See [Deactivate Host Encryption Mode Using the API](#).

Automatic changes occur when encryption operations attempt to set host encryption mode. For example, assume that you add an encrypted virtual machine to a standalone host. Host encryption mode is not set. If you have the required privileges on the host, encryption mode is automatically set.

Assume that a cluster has three ESXi hosts, host A, B, and C. You create an encrypted virtual machine on host A. What happens depends on several factors.

- If hosts A, B, and C already have host encryption mode set, you need only **Cryptographic operations > Encrypt new** privileges to create the virtual machine.

- If hosts A and B are set for host encryption, and C is not, the system proceeds as follows.
 - Assume that you have both the **Cryptographic operations > Encrypt new** and the **Cryptographic operations > Register host** privileges on each host. In this case, the encryption process sets host encryption mode on host C, and pushes the key to each host in the cluster.
For this case, you can also explicitly set host encryption mode on host C.
 - Assume that you have only **Cryptographic operations > Encrypt new** privileges on the virtual machine or virtual machine folder. In that case, virtual machine creation succeeds and the key becomes available on host A and host B. Host C remains deactivated for encryption and does not have the virtual machine key.
- If none of the hosts has host encryption mode set, and you have **Cryptographic operations > Register host** privileges on host A, then the virtual machine creation process sets host encryption mode on that host. Otherwise, for hosts B and C, an error results.
- You can also use the vSphere API to set the encryption mode of a cluster to "force enable." Force enable causes all hosts in the cluster to be cryptographically "safe," that is, vCenter Server has installed a host key on the host. See *vSphere Web Services SDK Programming Guide*.

Disk Space Requirements When Encrypting Virtual Machines

When you encrypt an existing virtual machine, you need at least twice the space that the virtual machine is currently using.

What Is Encrypted vSphere vMotion

Encrypted vSphere vMotion secures confidentiality, integrity, and authenticity of data that is transferred with vSphere vMotion. vSphere supports encrypted vMotion of unencrypted and encrypted virtual machines across vCenter Server instances.

vSphere vMotion always uses encryption when migrating encrypted virtual machines. For virtual machines that are not encrypted, you can select one of the encrypted vSphere vMotion options.

What Is Encrypted in Encrypted vSphere vMotion

For encrypted disks, the data is transmitted encrypted in all cases. For unencrypted disks, the following applies:

- If disk data is transferred within a host, that is without changing the host, you change only the datastore, the transfer is unencrypted.
- If disk data is transferred between hosts and encrypted vMotion is used, the transfer is encrypted. If encrypted vMotion is not used the transfer is unencrypted.

For virtual machines that are encrypted, migration with vSphere vMotion always uses encrypted vSphere vMotion. You cannot turn off encrypted vSphere vMotion for encrypted virtual machines.

Encrypted vSphere vMotion States for Unencrypted Virtual Machines

For virtual machines that are not encrypted, you can set encrypted vSphere vMotion to one of the following states. The default is Opportunistic.

Disabled

Do not use encrypted vSphere vMotion.

Opportunistic

Use encrypted vSphere vMotion if the source and the destination hosts support it. Only ESXi hosts of version 6.5 and later support encrypted vSphere vMotion.

Required

Allow only encrypted vSphere vMotion. If the source or the destination host does not support encrypted vSphere vMotion, migration with vSphere vMotion is not allowed.

When you encrypt a virtual machine, the virtual machine keeps a record of the current encrypted vSphere vMotion setting. If you later deactivate encryption for the virtual machine, the encrypted vMotion setting remains at Required until you change the setting explicitly. You can change the settings using **Edit Settings**.

Note: Currently, you must use the vSphere APIs to migrate or clone encrypted virtual machines across vCenter Server instances. See *vSphere Web Services SDK Programming Guide* and *vSphere Web Services API Reference*.

Migrating or Cloning Encrypted Virtual Machines Across vCenter Server Instances

vSphere vMotion supports migrating and cloning encrypted virtual machines across vCenter Server instances.

When migrating or cloning encrypted virtual machines across vCenter Server instances, the source and destination vCenter Server instances must be configured to share the key provider that was used to encrypt the virtual machine. In addition, the key provider name must be the same on both the source and destination vCenter Server instances and have the following characteristics:

- Standard key provider: The same key server (or key servers) must be in the key provider.
- Trusted key provider: The same vSphere Trust Authority service must be configured on the destination host.
- vSphere Native Key Provider: Must have the same KDK.

Note: You cannot clone or migrate an encrypted virtual machine using vSphere Native Key Provider to a standalone host, no matter that the source host resides in a cluster.

The destination vCenter Server ensures the destination ESXi host has encryption mode set, ensuring the host is cryptographically "safe."

The following privileges are required when using vSphere vMotion to migrate or clone an encrypted virtual machine across vCenter Server instances.

- Migrating: **Cryptographic operations > Migrate** on the virtual machine
- Cloning: **Cryptographic operations > Clone** on the virtual machine

Also, the destination vCenter Server must have the **Cryptographic operations > EncryptNew** privilege. If the destination ESXi host is not in "safe" mode, the **Cryptographic operations > RegisterHost** privilege must also be on the destination vCenter Server.

Certain tasks are not allowed when migrating virtual machines (non-encrypted or encrypted), either on the same vCenter Server or across vCenter Server instances.

- You cannot change the VM Storage Policy.
- You cannot perform a key change.

Note: You can change the VM Storage Policy while cloning virtual machines.

Minimum Requirements for Migrating or Cloning Encrypted Virtual Machines Across vCenter Server Instances

The minimum version requirements for migrating or cloning standard key provider encrypted virtual machines across vCenter Server instances using vSphere vMotion are:

- Both the source and destination vCenter Server instances must be on version 7.0 or later.
- Both the source and destination ESXi hosts must be on version 6.7 or later.

The minimum version requirements for migrating or cloning trusted key provider encrypted virtual machines across vCenter Server instances using vSphere vMotion are:

- The vSphere Trust Authority service must be configured for the destination host and the destination host must be attested.

- Encryption cannot change on migration. For example, an unencrypted disk cannot be encrypted while the virtual machine is migrated to the new storage.
- You can migrate a standard encrypted virtual machine onto a Trusted Host. The key provider name must be the same on both the source and destination vCenter Server instances.
- You cannot migrate a vSphere Trust Authority encrypted virtual machine onto a non-Trusted Host.

Trusted Key Provider vMotion and Cross-vCenter Server vMotion

Trusted key provider fully supports vMotion across ESXi hosts.

Cross-vCenter Server vMotion is supported, but with the following restrictions.

1. The required trusted service must be configured on the destination host and the destination host must be attested.
2. Encryption cannot change on migration. For example, a disk cannot be encrypted while the virtual machine is migrated to the new storage.

When performing cross-vCenter Server vMotion, vCenter Server checks that the trusted key provider is available on the destination host, and if the host has access to it.

vSphere Native Key Provider vMotion and Cross-vCenter Server vMotion

vSphere Native Key Provider supports vMotion and Encrypted vMotion across ESXi hosts. Cross-vCenter Server vMotion is supported if vSphere Native Key Provider is configured on the destination host.

Virtual Machine Encryption Best Practices

Follow virtual machine encryption best practices to avoid problems later, for example, when you generate a `vm-support` bundle.

Best Practices for Virtual Machine Encryption to Get Started With

To avoid problems when using virtual machine encryption, follow these general best practices.

- Do not encrypt any vCenter Server appliance virtual machines.
- If your ESXi host fails, retrieve the support bundle as soon as possible. The host key must be available for generating a support bundle that uses a password, or for decrypting a core dump. If the host is rebooted, it is possible that the host key changes. If that happens, you can no longer generate a support bundle with a password or decrypt core dumps in the support bundle with the host key.
- Manage key provider names carefully. If the key provider name changes for a key server that is already in use, a VM that is encrypted with keys from that key server enters a locked state during power-on or register. In that case, remove the key server from the vCenter Server and add it with the key provider name that you used initially.
- Do not edit VMX files and VMDK descriptor files. These files contain the encryption bundle. It is possible that your changes make the virtual machine unrecoverable, and that the recovery problem cannot be fixed.
- The vSphere Virtual Machine Encryption process encrypts data on the host before writing the data to storage. The effectiveness of back-end storage features such as deduplication, compression, replication, and so on, might be affected when encrypting virtual machines in this manner.
- If you use multiple layers of encryption, for example, vSphere Virtual Machine Encryption and in-guest encryption (BitLocker, dm-crypt, and so on), overall virtual machine performance might be affected, because the encryption processes use additional CPU and memory resources.
- Ensure that replicated copies of virtual machines encrypted with vSphere Virtual Machine Encryption have access to the encryption keys at the recovery site. For standard key providers, this is handled as part of the design of the Key Management System, outside of vSphere. For vSphere Native Key Provider, ensure that a backup copy of the Native Key Provider key exists and is protected against loss. For more information, see [Back up a vSphere Native Key Provider](#).

- Encryption is CPU intensive. AES-NI significantly improves encryption performance. Enable AES-NI in your BIOS.

Best Practices for Encrypted Core Dumps

Follow these best practices to avoid having problems when you want to examine a core dump to diagnose a problem.

- Establish a policy regarding core dumps. Core dumps are encrypted because they can contain sensitive information such as keys. If you decrypt a core dump, consider it sensitive information. ESXi core dumps might contain keys for the ESXi host and for the virtual machines on it. Consider changing the host key and recrypting encrypted virtual machines after you decrypt a core dump. You can perform both tasks by using the vSphere API. See [vSphere Virtual Machine Encryption and Core Dumps](#) for details.
- Always use a password when you collect a `vm-support` bundle. You can specify the password when you generate the support bundle from the vSphere Client or using the `vm-support` command.
The password recrypts core dumps that use internal keys to use keys that are based on the password. You can later use the password to decrypt any encrypted core dumps that might be included in the support bundle. Unencrypted core dumps and logs are not affected by using the password option.
- The password that you specify during `vm-support` bundle creation is not persisted in vSphere components. You are responsible for keeping track of passwords for support bundles.
- Before you change the host key, generate a `vm-support` bundle with a password. You can later use the password to access any core dumps that might have been encrypted with the old host key.

Best Practices for Key Lifecycle Management

Implement best practices that guarantee key server availability and monitor keys on the key server.

- You are responsible for having policies in place that guarantee key server availability.
If the key server is not available, virtual machine operations that require that vCenter Server request the key from the key server are not possible. That means running virtual machines continue to run, and you can power on, power off, and reconfigure those virtual machines. However, you cannot relocate the virtual machine to a host that does not have the key information.
Most key server solutions include high availability features. You can use the vSphere Client or the API to specify a key provider and the associated key servers.
- Note:** Starting in version 7.0 Update 2, encrypted virtual machines and virtual TPMs can continue to function even when the key server is temporarily offline or unavailable. The ESXi hosts can persist the encryption keys to continue encryption and vTPM operations. See [vSphere Key Persistence on ESXi Hosts](#).

- You are responsible for keeping track of keys and for performing remediation if keys for existing virtual machines are not in the Active state.

The KMIP standard defines the following states for keys.

- Pre-Active
- Active
- Deactivated
- Compromised
- Destroyed
- Destroyed Compromised

vSphere Virtual Machine Encryption uses only Active keys for encryption. If a key is Pre-Active, vSphere Virtual Machine Encryption activates it. If the key state is Deactivated, Compromised, Destroyed, Destroyed Compromised, you cannot encrypt a virtual machine or disk with that key.

For keys that are in other states, virtual machines using those keys continue to work. Whether a clone or migration operation succeeds depends on whether the key is already on the host.

- If the key is on the destination host, the operation succeeds even if the key is not Active on the key server.
- If the required virtual machine and virtual disk keys are not on the destination host, vCenter Server has to fetch the keys from the key server. If the key state is Deactivated, Compromised, Destroyed, or Destroyed Compromised, vCenter Server displays an error and the operation does not succeed.

A clone or migration operation succeeds if the key is already on the host. The operation fails if vCenter Server has to pull the keys from the key server.

If a key is not Active, perform a rekey operation using the API. See the *vSphere Web Services SDK Programming Guide*.

- Develop key rotation policies so that keys are retired and rolled over after a specific time.
 - Trusted key provider: Change the primary key of a trusted key provider.
 - vSphere Native Key Provider: Change the `key_id` of a vSphere Native Key Provider.

Best Practices for Backup and Restore

Set up policies on backup and restore operations.

- Not all backup architectures are supported. See [Virtual Machine Encryption Interoperability](#).
- Set up policies for restore operations. Because backup is always in cleartext, plan to encrypt virtual machines right after restore is finished. You can specify that the virtual machine is encrypted as part of the restore operation. If possible, encrypt virtual machine as part of the restore process to avoid exposing sensitive information. To change the encryption policy for any disks that are associated with the virtual machine, change the storage policy for the disk.
- Because the VM home files are encrypted, ensure that the encryption keys are available at the time of a restore.

Best Practices for Encryption Performance

- Encryption performance depends on the CPU and storage speed.
- Encrypting existing virtual machines is more time consuming than encrypting a virtual machine during creation. Encrypt a virtual machine when you create it if possible.

Best Practices for the Sample Storage Policy

Do not modify the bundled VM Encryption sample storage policy. Instead, clone the policy and edit the clone.

Note: No automated way of returning VM Encryption Policy to its original settings exists.

See the *vSphere Storage* documentation for details customizing storage policies.

Best Practices for Removing Encryption Keys

To ensure that encryption keys are removed from a cluster, after deleting, unregistering, or moving the encrypted virtual machine to another vCenter Server, reboot the ESXi hosts in the cluster.

Virtual Machine Encryption Caveats

Review Virtual Machine Encryption caveats to avoid problems later.

To understand which devices and features cannot be used with Virtual Machine Encryption, see [Virtual Machine Encryption Interoperability](#).

Encrypted Virtual Machine Limitations

Consider the following caveats when you plan your virtual machine encryption strategy.

- When you clone an encrypted virtual machine or perform a Storage vMotion operation, you can attempt to change the disk format. Such conversions do not always succeed. For example, if you clone a virtual machine and attempt to change the disk format from lazy-zeroed thick format to thin format, the virtual machine disk keeps the lazy-zeroed thick format.
- When you detach a disk from a virtual machine, the storage policy information for the virtual disk is not retained.
 - If the virtual disk is encrypted, you must explicitly set the storage policy to VM Encryption Policy or to a storage policy that includes encryption.
 - If the virtual disk is not encrypted, you can change the storage policy when you add the disk to a virtual machine. See [Virtual Disk Encryption](#) for details.
- Decrypt core dumps before moving a virtual machine to a different cluster.
 The vCenter Server does not store the key server keys but only tracks the key IDs. As a result, vCenter Server does not store the ESXi host key persistently. However, in vSphere 7.0 Update 2 and later, encrypted devices can function even when access to a key server is disrupted. See [vSphere Key Persistence on ESXi Hosts](#).
 Under certain circumstances, for example, when you move the ESXi host to a different cluster and reboot the host, vCenter Server assigns a new host key to the host. You cannot decrypt any existing core dumps with the new host key.
- OVF Export is not supported for an encrypted virtual machine.
- Using the VMware Host Client to register an encrypted virtual machine is not supported.

Virtual Machine Locked State

If the virtual machine key or one or more of the virtual disk keys are missing, the virtual machine enters a locked state. In a locked state, you cannot perform virtual machine operations.

- When you encrypt both a virtual machine and its disks from the vSphere Client, the same key is used for both.
- When you perform the encryption using the API, you can use different encryption keys for the virtual machine and for disks. In that case, if you attempt to power on a virtual machine, and one of the disk keys is missing, the power on operation fails. If you remove the virtual disk, you can power on the virtual machine.

See [Resolve Missing Encryption Key Issues](#) for troubleshooting suggestions.

Virtual Machine Encryption Interoperability

vSphere Virtual Machine Encryption has some limitations regarding devices and features that it can interoperate with.

The following limitations and remarks refer to using vSphere Virtual Machine Encryption. For similar information about using vSAN encryption, see the *Administering VMware vSAN* documentation.

Limitations on Certain Encryption Tasks

Some restrictions apply when performing certain tasks on an encrypted virtual machine.

- For most virtual machine encrypted operations, you must power off the virtual machine. You can clone an encrypted virtual machine and you can perform a shallow recrypt while the virtual machine is powered on.
Note: Virtual machines configured with IDE controllers must be powered off to perform a shallow rekey operation.
- You cannot perform a deep recrypt on a virtual machine with snapshots. You can perform a shallow recrypt on a virtual machine with snapshots.

Virtual Trusted Platform Module Devices and vSphere Virtual Machine Encryption

A virtual Trusted Platform Module (vTPM) is a software-based representation of a physical Trusted Platform Module 2.0 chip. You can add a vTPM to either a new or an existing virtual machine. To add a vTPM to a virtual machine, you must configure a key provider in your vSphere environment. When you configure a vTPM, the virtual machine “home” files are encrypted (memory swap, NVRAM files, and so on). The disk files, or VMDK files, are not automatically encrypted. You can choose to add encryption explicitly for the virtual machine disks.

CAUTION: Cloning a virtual machine duplicates the entire virtual machine, including the virtual devices such as a vTPM. Information stored in the vTPM, including properties of the vTPM that software can use to determine a system’s identity, is also duplicated.

In vSphere 8.0 and later, when cloning a virtual machine that includes a vTPM, you can choose to start with a new, blank vTPM, which gets its own secrets and identity.

vSphere Virtual Machine Encryption and Suspended State and Snapshots

You can resume from a suspended state of an encrypted virtual machine, or revert to a memory snapshot of an encrypted machine. You can migrate an encrypted virtual machine with memory snapshot and suspended state between ESXi hosts.

vSphere Virtual Machine Encryption and IPv6

You can use vSphere Virtual Machine Encryption with pure IPv6 mode or in mixed mode. You can configure the key server with IPv6 addresses. You can configure both the vCenter Server and the key server with only IPv6 addresses.

Limitations on Cloning in vSphere Virtual Machine Encryption

For all key provider types, cloning is supported conditionally. You can change encryption keys on clone. Certain cloning features do not work with vSphere Virtual Machine Encryption.

- Full clones are supported. The clone inherits the parent encryption state including keys. You can encrypt the full clone, re-encrypt the full clone to use new keys, or decrypt the full clone.
Linked clones are supported and the clone inherits the parent encryption state including keys. You cannot decrypt the linked clone or re-encrypt a linked clone with different keys.
Note: Verify that other applications support linked clones. For example, VMware Horizon® 7 supports both full clones and instant clones, but not linked clones.
- Instant clone is supported by all key provider types, but you cannot change encryption keys on clone.
- You can create a linked clone virtual machine from an encrypted virtual machine. The linked clone virtual machine contains the same keys. You can rekey the encrypted virtual machine “home” files of a linked clone, but you cannot rekey the disks.

Limitations with vSphere Native Key Provider

Certain operations are not supported with vSphere Native Key Provider.

- You cannot use vSphere Native Key Provider to encrypt virtual machines on a standalone host. The host must reside in a cluster to use vSphere Native Key Provider.
- You cannot move a host that contains virtual machines encrypted using vSphere Native Key Provider to a different cluster unless the target cluster contains the same vSphere Native Key Provider. (The encrypted virtual machines on the moved host become locked when the encryption keys are not present and the target cluster does not have the same vSphere Native Key Provider.)
- You cannot register a virtual machine encrypted by vSphere Native Key Provider to a legacy host because of the lack of support for vSphere Native Key Provider.
- You cannot register a virtual machine encrypted by vSphere Native Key Provider to a standalone host because of the requirements for the host to reside in a cluster.

Unsupported Disk Configurations with vSphere Virtual Machine Encryption

Certain types of virtual machine disk configurations are not supported with vSphere Virtual Machine Encryption.

- RDM (Raw Device Mapping). However, vSphere Virtual Volumes (vVols) are supported.
- Multi-writer or shared disks (MSCS, WSFC, or Oracle RAC). Encrypted virtual machine “home” files are supported for multi-writer disks. Encrypted virtual disks are not supported for multi-writer disks. If you attempt to select Multi-writer in the **Edit Settings** page of the virtual machine with encrypted virtual disks, the **OK** button is deactivated.

Miscellaneous Limitations in vSphere Virtual Machine Encryption

Other features that do not work with vSphere Virtual Machine Encryption include the following.

- vSphere ESXi Dump Collector
- Content Library
 - Content libraries support two types of templates, the OVF Template type and the VM Template type. You cannot export an encrypted virtual machine to the OVF Template type. The OVF Tool does not support encrypted virtual machines. You can create encrypted VM templates using the VM Template type. In vSphere 8.0 and later, the `ovftool` command includes an option to add a vTPM placeholder to the OVF descriptor file. When deploying a virtual machine from such a template, vCenter Server creates a vTPM with unique secrets on the destination virtual machine. See the *vSphere Virtual Machine Administration* documentation.
- Software for backing up encrypted virtual disks must use the VMware vSphere Storage API - Data Protection (VADP) to either back up the disks in hot-add mode or NBD mode with SSL enabled. However, not all backup solutions that use VADP for virtual disk backup are supported. Check with your backup vendor for details.
 - VADP SAN transport mode solutions are not supported for backing up encrypted virtual disks.
 - VADP Hot-Add solutions are supported for encrypted virtual disks. The backup software must support encryption of the proxy VM that is used as part of the hot-add backup workflow. The vendor must have the privilege **Cryptographic Operations > Encrypt Virtual Machine**.
 - Backup solutions using the NBD-SSL transport modes are supported for backing up encrypted virtual disks. The vendor application must have the privilege **Cryptographic Operations > Direct Access**.
- You cannot send output from an encrypted virtual machine to a serial port or parallel port. Even if the configuration appears to succeed, output is sent to a file.
- vSphere Virtual Machine Encryption is not supported in VMware Cloud on AWS. See the *Managing the VMware Cloud on AWS Data Center* documentation.

vSphere Key Persistence on ESXi Hosts

In vSphere 7.0 Update 2 and later, encrypted virtual machines and virtual TPMs can continue to optionally function even when the key server is temporarily offline or unavailable. The ESXi hosts can persist the encryption keys to continue encryption and vTPM operations.

Before vSphere 7.0 Update 2, encrypted virtual machines and vTPMs require that the key server always is available to function. In vSphere 7.0 Update 2 and later, encrypted devices can function even when access to a key server is disrupted.

In vSphere 7.0 Update 3 and later, encrypted vSAN clusters can also function even when access to a key provider is disrupted.

Note: Key persistence is not necessary when using vSphere Native Key Provider. vSphere Native Key Provider is designed out-of-the-box to run without requiring access to a key server. See the following section, "Key Persistence and vSphere Native Key Provider."

How Does Key Persistence on ESXi Hosts Work

When using a standard key provider, the ESXi host relies on vCenter Server to manage the encryption keys. When using a trusted key provider, the ESXi host relies directly on the Trust Authority Hosts for keys, and vCenter Server is not involved. vSphere Native Key Provider handles keys differently. See the next section for more information.

Regardless of the type of key provider, the ESXi host obtains the keys initially and retains them in its key cache. If the ESXi host reboots, it loses its key cache. The ESXi host then requests the keys again, either from the key server (standard key provider), or the Trust Authority Hosts (trusted key provider). When the ESXi host tries to obtain keys and the key server is offline or unreachable, vTPMs and workload encryption cannot function. For edge-style deployments, in which a key server is typically not deployed on site, a loss of connectivity to a key server can cause unnecessary downtime for encrypted workloads.

In vSphere 7.0 Update 2 and later, encrypted workloads can continue to function even when the key server is offline or unreachable. If the ESXi host has a TPM, the encryption keys are persisted in the TPM across reboots. So, even if an ESXi host reboots, the host does not need to request encryption keys. Also, encryption and decryption operations can continue when the key server is unavailable, because the keys have persisted in the TPM. In essence, depending on the key provider, when either the key server or Trust Authority Hosts are unavailable, you can keep running encrypted workloads "key server free." Also, vTPMs can likewise continue to function even when the key server is unreachable.

Key Persistence and vSphere Native Key Provider

When using a vSphere Native Key Provider, vSphere generates encryption keys and no key server is required. The ESXi hosts get a Key Derivation Key (KDK), which is used to derive other keys. After receiving the KDK and generating other keys, the ESXi hosts do not need access to vCenter Server to do encryption operations. In essence, a vSphere Native Key Provider always runs "key server free."

The KDK persists on an ESXi host by default even after reboot, and even when vCenter Server is not available after the host reboots.

You can activate key persistence with vSphere Native Key Provider, but it is normally unnecessary. The ESXi hosts have complete access to vSphere Native Key Provider and so additional persistence of keys is redundant. One use case for activating key persistence with vSphere Native Key Provider is when you also have a standard key provider (external KMIP server) configured.

How Do You Set Up Key Persistence

To activate or deactivate key persistence, see [Activate and Deactivate Key Persistence on an ESXi Host](#).

Configuring and Managing a Standard Key Provider

Using a standard key provider in your vSphere environment requires some preparation. After your environment is set up, you can create encrypted virtual machines and virtual disks and encrypt existing virtual machines and disks.

After your environment is set up for a standard key provider, you can use the vSphere Client to create encrypted virtual machines and virtual disks and encrypt existing virtual machines and disks. See [Using Encryption in Your vSphere Environment](#).

You can perform additional tasks by using the API and by using the `crypto-util` CLI. See the *vSphere Web Services SDK Programming Guide* for the API documentation and the `crypto-util` command-line help for details about that tool.

What Is a Standard Key Provider

You can use a standard key provider to perform virtual machine encryption tasks.

In vSphere, a standard key provider gets encryption keys directly from a key server, and the vCenter Server distributes the keys to the required ESXi hosts in a data center.

You can add separate standard key providers for different users and set the default standard key provider.

Standard Key Provider Requirements

- vSphere 6.5 or later
- An external key server (KMS)

The key server must support the Key Management Interoperability Protocol (KMIP) 1.1 standard. See the *vSphere Compatibility Matrices* for details.

You can find information about VMware certified key server (KMS) vendors in the [VMware Compatibility Guide](#) under Platform and Compute. If you select Compatibility Guides, you can open the Key Management Server (KMS) compatibility documentation. This documentation is updated frequently.

Standard Key Provider Privileges

Standard key providers use the **Cryptographer.*** privileges. See [Cryptographic Operations Privileges](#).

Setting Up the Standard Key Provider

Before you can start with virtual machine encryption tasks, you must set up the standard key provider.

Setting up a standard key provider includes adding the key provider and establishing trust with the key server. When you add a key provider, you are prompted to make it the default. You can explicitly change the default key provider. vCenter Server provisions keys from the default key provider.

Note: What was previously called a Key Management Server cluster in vSphere 6.5 and 6.7 is now called a key provider.

Add a Standard Key Provider Using the vSphere Client

You can add a standard key provider to your vCenter Server system from the vSphere Client or by using the public API.

- Verify that the key server (KMS) is in the *VMware Compatibility Guide for Key Management Servers (KMS)* and is KMIP 1.1 compliant, and that it can be a symmetric key foundry and server.
- Verify that you have the required privileges: **Cryptographic operations > Manage key servers**.
- Ensure that the key server is highly available. Loss of connection to the key server, such as during a power outage or a disaster recovery event, renders encrypted virtual machines inaccessible.

Note: In vSphere 7.0 Update 2 and later, encrypted virtual machines and virtual TPMs can continue to function even when the key server is temporarily offline or unavailable. See [vSphere Key Persistence on ESXi Hosts](#).

- Consider your infrastructure's dependencies on the key server carefully. Some KMS solutions are delivered as virtual appliances, making it possible to create a dependency loop or other availability problem with poor placement of the KMS appliance.

The vSphere Client enables you to add a standard key provider to your vCenter Server system, and establish trust between the key server and vCenter Server.

- You can add multiple key servers from the same vendor.
 - If your environment supports solutions from different vendors, you can add multiple key providers.
 - If your environment includes multiple key providers, and you delete the default key provider, you must set another default explicitly.
 - You can configure the key server with IPv6 addresses.
 - Both the vCenter Server system and the key server can be configured with only IPv6 addresses.
- Log in to the vCenter Server system with the vSphere Client.
 - Browse the inventory list and select the vCenter Server instance.
 - Click **Configure**, and under **Security** click **Key Providers**.
 - Click **Add Standard Key Provider** and enter the key provider information.

Option	Value
Name	Name for the key provider. Each logical key provider, regardless of its type (Standard, Trusted, and Native Key Provider), must have a unique name across all vCenter Server systems. For more information, see Key Provider Naming .
KMS	Alias for the key server (KMS).
Address	IP address or FQDN of the key server.
Port	Port on which vCenter Server connects to the key server.
Proxy server	Optional proxy server address for connecting to the key server.
Proxy port	Optional proxy port for connecting to the key server.
Username	Some key server vendors allow users to isolate encryption keys that are used by different users or groups by specifying a user name and password. Specify a user name only if your key server supports this functionality, and if you intend to use it.
Password	Some key server vendors allow users to isolate encryption keys that are used by different users or groups by specifying a user name and password. Specify a password only if your key server supports this functionality, and if you intend to use it.

You can click **Add KMS** to add more key servers.

- Click **Add Key Provider**.

- Click **Trust**.

vCenter Server adds the key provider and displays the status as Connected.

See [Establish a Standard Key Provider Trusted Connection by Exchanging Certificates](#).

Establish a Standard Key Provider Trusted Connection by Exchanging Certificates

After you add the standard key provider to the vCenter Server system, you can establish a trusted connection. The exact process depends on the certificates that the key provider accepts, and on your company policy.

Add the standard key provider.

1. Navigate to the .
2. Click **Configure** and select **Key Providers** under **Security**.
3. Select the key provider.
The KMS for the key provider is displayed.
4. Select the KMS.
5. From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
6. Select the option appropriate for your server and follow the steps.

Option	See
vCenter Server Root CA certificate	Use the Root CA Certificate Option to Establish a Standard Key Provider Trusted Connection.
vCenter Server Certificate	Use the Certificate Option to Establish a Standard Key Provider Trusted Connection.
Upload certificate and private key	Use the Upload Certificate and Private Key Option to Establish a Standard Key Provider Trusted Connection.
New Certificate Signing Request	Use the New Certificate Signing Request Option to Establish a Standard Key Provider Trusted Connection.

Use the Root CA Certificate Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors require that you upload your root CA certificate to the KMS. All certificates that are signed by your root CA are then trusted by this KMS.

The root CA certificate that vSphere Virtual Machine Encryption uses is a self-signed certificate that is stored in a separate store in the VMware Endpoint Certificate Store (VECS) on the vCenter Server system.

Note: Generate a root CA certificate only if you want to replace existing certificates. If you do, other certificates that are signed by that root CA become invalid. You can generate a new root CA certificate as part of this workflow.

1. Navigate to the .
2. Click **Configure** and select **Key Providers** under **Security**.
3. Select the key provider with which you want to establish a trusted connection.
The key server (KMS) for the key provider is displayed.
4. From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
5. Select **vCenter Root CA Certificate** and click **Next**.
The Download Root CA Certificate dialog box is populated with the root certificate that vCenter Server uses for encryption. This certificate is stored in VECS.
6. Copy the certificate to the clipboard or download the certificate as a file.
7. Follow the instructions from your KMS vendor to upload the certificate to their system.

Note: Some KMS vendors require that the KMS vendor restarts the KMS to pick up the root certificate that you upload.

Finalize the certificate exchange. See [Finish the Trust Setup for a Standard Key Provider](#).

Use the Certificate Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors require that you upload the vCenter Server certificate to the KMS. After the upload, the KMS accepts traffic that comes from a system with that certificate.

vCenter Server generates a certificate to protect connections with the KMS. The certificate is stored in a separate key store in the VMware Endpoint Certificate Store (VECS) on the vCenter Server system.

1. Navigate to the .
 2. Click **Configure** and select **Key Providers** under **Security**.
 3. Select the key provider with which you want to establish a trusted connection.
The key server (KMS) for the key provider is displayed.
 4. From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
 5. Select **vCenter Certificate** and click **Next**.
The Download Certificate dialog box is populated with the root certificate that vCenter Server uses for encryption. This certificate is stored in VECS.
- Note:** Do not generate a new certificate unless you want to replace existing certificates.
6. Copy the certificate to the clipboard or download it as a file.
 7. Follow the instructions from your KMS vendor to upload the certificate to the KMS.

Finalize the trust relationship. See [Finish the Trust Setup for a Standard Key Provider](#).

Use the Upload Certificate and Private Key Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors require that you upload the KMS server certificate and private key to the vCenter Server system.

- Request a certificate and private key from the KMS vendor. The files are X509 files in PEM format.

Some KMS vendors generate a certificate and private key for the connection and make them available to you. After you upload the files, the KMS trusts your vCenter Server instance.

1. Navigate to the .
2. Click **Configure** and select **Key Providers** under **Security**.
3. Select the key provider with which you want to establish a trusted connection.
The key server (KMS) for the key provider is displayed.
4. From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
5. Select **KMS certificate and private key** and click **Next**.
6. Paste the certificate that you received from the KMS vendor into the top text box or click **Upload a File** to upload the certificate file.
7. Paste the key file into the bottom text box or click **Upload a File** to upload the key file.
8. Click **Establish Trust**.

Finalize the trust relationship. See [Finish the Trust Setup for a Standard Key Provider](#).

Use the New Certificate Signing Request Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors require that vCenter Server generate a Certificate Signing Request (CSR) and send that CSR to the KMS. The KMS signs the CSR and returns the signed certificate. You can upload the signed certificate to vCenter Server.

Using the **New Certificate Signing Request** option is a two-step process. First you generate the CSR and send it to the KMS vendor. Then you upload the signed certificate that you receive from the KMS vendor to vCenter Server.

1. Navigate to the .
2. Click **Configure** and select **Key Providers** under **Security**.
3. Select the key provider with which you want to establish a trusted connection.
The key server (KMS) for the key provider is displayed.
4. From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
5. Select **New Certificate Signing Request (CSR)** and click **Next**.
6. In the dialog box, copy the full certificate in the text box to the clipboard or download it as a file.
Use the **Generate new CSR** button in the dialog box only if you explicitly want to generate a CSR.
7. Follow the instructions from your KMS vendor to submit the CSR.
8. When you receive the signed certificate from the KMS vendor, click **Key Providers** again, select the key provider, and from the **Establish Trust** drop-down menu, select **Upload Signed CSR Certificate**.
9. Paste the signed certificate into the bottom text box or click **Upload File** and upload the file, and click **Upload**.

Finalize the trust relationship. See [Finish the Trust Setup for a Standard Key Provider](#).

Finish the Trust Setup for a Standard Key Provider

Unless the **Add Standard Key Provider** dialog box prompted you to trust the KMS, you must explicitly establish trust after certificate exchange is complete.

You can complete the trust setup, that is, make vCenter Server trust the KMS, either by trusting the KMS or by uploading a KMS certificate. You have two options:

- Trust the certificate explicitly by using the **Upload KMS certificate** option.
- Upload a KMS leaf certificate or the KMS CA certificate to vCenter Server by using the **Make vCenter Trust KMS** option.

Note: If you upload the root CA certificate or the intermediate CA certificate, vCenter Server trusts all certificates that are signed by that CA. For strong security, upload a leaf certificate or an intermediate CA certificate that the KMS vendor controls.

1. Navigate to the .
2. Click **Configure** and select **Key Providers** under **Security**.
3. Select the key provider with which you want to establish a trusted connection.
The key server (KMS) for the key provider is displayed.

4. Select the KMS.
5. Select one of the following options from the **Establish Trust** drop-down menu.

Option	Action
Make vCenter Trust KMS	In the dialog box that appears, click Trust .
Upload KMS certificate	1. In the dialog box that appears, either paste in the certificate, or click Upload a file and browse to the certificate file. 2. Click Upload.

Set Up Separate Key Providers for Different Users

You can set up your environment with different key providers for different users of the same KMS instance. Having multiple key providers is helpful, for example, if you want to grant different departments in your company access to different sets of encryption keys.

Set up the connection with the KMS.

You can use multiple key providers for the same KMS to separate keys. Having separate sets of keys is essential for use cases like different BUs or different customers.

Note: Not all KMS vendors support multiple users.

1. Create two users with corresponding user names and passwords, for example C1 and C2, on the KMS.
2. Log in to vCenter Server and create the first key provider.
3. When prompted for a user name and password, give information that is unique to the first user.
4. Create a second key provider and add the same KMS, but use the second user name and password (C2).

The two key providers have independent connections to the KMS and use a different set of keys.

Delete a Standard Key Provider

You can use the vSphere Client to delete a standard key provider from vCenter Server.

Required privilege: **Cryptographic operations > Manage key servers**

Before you delete a standard key provider, rekey any encrypted virtual machines and datastores that were encrypted using that key provider to another key provider. See [Rekey an Encrypted Virtual Machine Using the vSphere Client](#).

In addition, maintain a backup of the standard key provider in case you must rekey an encrypted virtual machine after deleting the key provider.

After you delete a standard key provider, virtual machines that have vTPMs or that are encrypted continue to run. If you reboot the ESXi host, its encrypted virtual machines enter a locked state. After you unregister these virtual machines,

they enter a locked state when you try to re-register them. The only way to unlock the virtual machines is to restore the previous standard key provider.

1. Log in to the vCenter Server system with the vSphere Client.
2. Browse the inventory list and select the vCenter Server instance.
3. Click **Configure**, and under **Security** click **Key Providers**.
4. Select the standard key provider you want to delete.
5. Click **Delete**.
6. Read the warning message and slide the slider all the way to the right.
7. Click **Delete**.

The standard key provider is removed from the vCenter Server.

Configuring and Managing vSphere Native Key Provider

Using a VMware vSphere®Native Key Provider™ in your vSphere environment requires some preparation. After you configure vSphere Native Key Provider, you can create virtual Trusted Platform Modules (vTPMs) on your virtual machines.

After you set up your environment for vSphere Native Key Provider, you can use the vSphere Client and API to create vTPMs. If you purchase the VMware vSphere®Enterprise Plus Edition™, you can also encrypt virtual machines and virtual disks, and encrypt existing virtual machines and disks.

vSphere Native Key Provider Overview

In vSphere 7.0 Update 2 and later, you can use the built-in vSphere Native Key Provider to enable encryption technologies, such as virtual TPMs (vTPM).

vSphere Native Key Provider is included in all vSphere editions and does not require an external key server (also called a Key Management Server (KMS) in the industry). You can also use vSphere Native Key Provider for vSphere Virtual Machine Encryption, but you must purchase the VMware vSphere®Enterprise Plus Edition™.

What Is vSphere Native Key Provider

With a standard key provider or trusted key provider, you must configure an external key server. In a standard key provider setup, vCenter Server fetches the keys from the external key server and distributes them to the ESXi hosts. In a trusted key provider (vSphere Trust Authority) setup, the trusted ESXi hosts fetch the keys directly.

With vSphere Native Key Provider, you no longer need an external key server. vCenter Server generates a primary key, called the Key Derivation Key (KDK), and pushes it to all ESXi hosts in the cluster. The ESXi hosts then generate data encryption keys (even when not connected to vCenter Server) to enable security functionality such as vTPMs. vTPM functionality is included in all vSphere editions. To use vSphere Native Key Provider for vSphere Virtual Machine Encryption, you must have purchased the vSphere Enterprise Plus Edition. vSphere Native Key Provider can coexist with an existing key server infrastructure.

vSphere Native Key Provider:

- Enables the use of vTPMs, vSphere Virtual Machine Encryption, and vSAN Data at Rest Encryption, when you do not require or want an external key server.
- Works only with VMware infrastructure products.
- Does not provide external interoperability, KMIP support, hardware security modules, or other features that a traditional, third-party, external key server can offer for interoperability or regulatory compliance. If your organization requires this functionality for non-VMware products and components, install a traditional, third-party key server.

- Helps address the needs of organizations that either cannot use, or do not want not to use, an external key server.
- Improves data sanitization and system reuse practices by enabling earlier use of encryption technologies on media that is difficult to sanitize, such as flash and SSD.
- Provides a transition path between key providers. vSphere Native Key Provider is compatible with the VMware standard key provider and the vSphere Trust Authority trusted key provider.
- Works with multiple vCenter Server systems using an Enhanced Linked Mode configuration or a vCenter Server High Availability configuration.
- Can be used to enable vTPM in all editions of vSphere, and encrypt virtual machines with the purchase of the vSphere Enterprise Plus Edition that includes vSphere Virtual Machine Encryption. vSphere Virtual Machine Encryption works with vSphere Native Key Provider as it does with VMware standard and trusted key providers.
- Can be used to enable vSAN Data at Rest Encryption with the use of an appropriate vSAN license.
- Can use a Trusted Platform Module (TPM) 2.0 to increase security when one is installed in an ESXi host. You can also configure vSphere Native Key Provider to be available only to hosts where a TPM 2.0 is installed. If you use a TPM, it must be TPM 2.0. vSphere Native Key Provider does not support TPM 1.2.

Note: An ESXi host does not require a TPM 2.0 to use a vSphere Native Key Provider. However, a TPM 2.0 does provide enhanced security.

As with all security solutions, consider the system design, implementation considerations, and tradeoffs of using Native Key Provider. For example, ESXi key persistence avoids the dependency on a key server always being available. However, because key persistence stores the Native Key Provider cryptographic information on the clustered hosts, you still are at risk if malicious actors steal the ESXi hosts themselves. Because environments differ, assess and implement your security controls in accordance with your organization's regulatory and security needs, operational requirements, and tolerance for risk.

For more overview information about vSphere Native Key Provider, see <https://core.vmware.com/native-key-provider>.

vSphere Native Key Provider Requirements

To use vSphere Native Key Provider, you must:

- Ensure that both the vCenter Server system and ESXi hosts are running vSphere 7.0 Update 2 or later.
- Configure the ESXi hosts in a cluster.
- While not required, as a best practice, use ESXi hosts that are as identical as possible, including TPMs. Cluster management and feature enablement is much easier when cluster hosts are identical.
- Configure the vCenter Server file-based backup and restore, and store the backups securely as they contain the Key Derivation Key. See the topic on vCenter Server backup and restore in the *vCenter Server Installation and Setup* documentation.

To perform vSphere Virtual Machine Encryption or vSAN encryption using vSphere Native Key Provider, you must purchase the edition of those products containing the appropriate license.

vSphere Native Key Provider and Enhanced Linked Mode

You can configure a single vSphere Native Key Provider that is shareable across vCenter Server systems configured in an Enhanced Linked Mode configuration. The high-level steps in this scenario are:

1. Creating the vSphere Native Key Provider on one of the vCenter Server systems
2. Backing up the Native Key Provider on the vCenter Server on which it was created
3. Exporting the Native Key Provider
4. Restoring the Native Key Provider to the other vCenter Server systems in the Enhanced Link Mode configuration (see [Restore a vSphere Native Key Provider Using the vSphere Client](#))

vSphere Native Key Provider Privileges

As with standard and trusted key providers, vSphere Native Key Provider uses the **Cryptographer.*** privileges. In addition, vSphere Native Key Provider uses the **Cryptographer.ReadKeyServersInfo** privilege, which is specific to vSphere Native Key Providers, to list vSphere Native Key Providers. See [Cryptographic Operations Privileges](#).

vSphere Native Key Provider Alarms

You must back up a vSphere Native Key Provider. When a vSphere Native Key Provider is not backed up, vCenter Server generates an alarm. When you back up the vSphere Native Key Provider for which an alarm was generated, vCenter Server resets the alarm. By default, vCenter Server checks for backed-up vSphere Native Key Providers once a day. You can change the checking interval by modifying the `vpzd.KMS.backupCheckInterval` option.

vSphere Native Key Provider Periodic Remediation Check

vCenter Server checks periodically that the vSphere Native Key Provider configuration on vCenter Server and ESXi hosts matches. When a host state changes, for example, when you add a host to the cluster, the key provider configuration on the cluster drifts away from the configuration on the host. If the configuration (keyID) differs on the host, vCenter Server updates the configuration of the host automatically. No manual intervention is required.

By default, vCenter Server checks the configuration every five minutes. You can modify the interval by using the `vpzd.KMS.remediationInterval` option.

Using vSphere Native Key Provider with a Disaster Recovery Site

You can use vSphere Native Key Provider with a backup disaster recovery site. Importing the vSphere Native Key Provider backup from the primary vCenter Server to the vCenter Server backup at the disaster recovery site enables that cluster to decrypt and run your encrypted virtual machines.

Always test your DR solution. Never assume that your solution works without trying a recovery. Ensure that a copy of the vSphere Native Key Provider backup is also available to your DR site.

Unsupported Features in vSphere Native Key Provider

vSphere Native Key Provider does not support First Class Disk (FCD) encryption in vSphere 8.0 through vSphere 8.0 Update 2. Starting in vSphere 8.0 Update 3, vSphere Native Key Provider does support FCD encryption.

Migrating Virtual Machines Using vSphere Native Key Provider Across Non-Linked vCenter Server Systems

The high-level steps to migrate a virtual machine, either encrypted or enabled with a vTPM through vSphere Native Key Provider, from one non-linked vCenter Server system to another, include:

1. Restoring the vSphere Native Key Provider to the to-be-migrated-to vCenter Server system.
2. Migrating the virtual machine by using vMotion.

vSphere Native Key Provider Process Flows

Understanding vSphere Native Key Provider process flows is essential to learning how to configure and manage your vSphere Native Key Provider.

You can use the built-in vSphere Native Key Provider to power encryption-based virtual TPMs (vTPM). vSphere Native Key Provider is included in all vSphere editions and does not require an external key server (KMS). To use vSphere Native Key Provider for vSphere Virtual Machine Encryption, you must purchase the vSphere Enterprise+ edition.

Configuring vSphere Native Key Provider

Configuring vSphere Native Key Provider involves these basic operations:

1. A user with the appropriate administrative privileges uses the vSphere Client to create a vSphere Native Key Provider on a vCenter Server.
2. The vCenter Server then configures the vSphere Native Key Provider for all clusters of ESXi hosts.
In this step, vCenter Server pushes a primary key to all ESXi hosts in the cluster. Likewise, if you update or delete a vSphere Native Key Provider, the change is pushed to the hosts in the cluster.
3. Users with the appropriate cryptographic privileges create vTPMs and encrypted virtual machines (provided you have purchased the vSphere Enterprise+ edition).
See [Using Encryption in Your vSphere Environment](#) and [Securing Virtual Machines with Virtual Trusted Platform Module](#).

vSphere Native Key Provider Encryption Process Flow

To understand how different components interact to perform an encryption task using vSphere Native Key Provider, see [vSphere Native Key Provider Encryption Process Flow](#).

Configure a vSphere Native Key Provider

A key provider is required for performing encryption tasks. You can use the vSphere Client to configure a vSphere Native Key Provider on vCenter Server.

Required privilege: **Cryptographic operations > Manage key servers**

vSphere 7.0 Update 2 and later includes a key provider called vSphere Native Key Provider. vSphere Native Key Provider enables encryption-related functionality without requiring an external key server (KMS). Initially, vCenter Server is not configured with a vSphere Native Key Provider. You must manually configure a vSphere Native Key Provider.

An ESXi host does not require a TPM 2.0 to use a vSphere Native Key Provider. However, a TPM 2.0 does provide enhanced security.

Note: When you configure vSphere Native Key Provider, the key providers are available on all clusters for the vCenter Server on which you configure them. As a result, all hosts attached to the vCenter Server have access to all the vSphere Native Key Providers that you configure.

1. Log in to the vCenter Server system with the vSphere Client.
2. Browse the inventory list and select the vCenter Server instance.
3. Click **Configure**, and under **Security** click **Key Providers**.
4. Click **Add** then click **Add Native Key Provider**.
5. Enter a name for the vSphere Native Key Provider.

Each logical key provider, regardless of its type (Standard, Trusted, and Native Key Provider), must have a unique name across all vCenter Server systems.

For more information, see [Key Provider Naming](#).

6. If you want this vSphere Native Key Provider to be used only by hosts with a TPM 2.0, select the **Use key provider only with TPM protected ESXi hosts** check box.
If enabled, the vSphere Native Key Provider is available only on hosts with a TPM 2.0.

7. Click **Add Key Provider**.

Note: It takes about five minutes for all the clustered ESXi hosts in a data center to get the key provider, and for the vCenter Server to update its cache. Because of the way the information is propagated, you might have to wait for a few minutes to use the key provider for key operations on some of the hosts.

The vSphere Native Key Provider is added and appears in the **Key Provider** pane. At this point, the vSphere Native Key Provider is not backed up. You must back up the vSphere Native Key Provider before you can use it.

See [Back up a vSphere Native Key Provider](#).

Back up a vSphere Native Key Provider

In case you must restore the key provider configuration, backing up a vSphere Native Key Provider is required as part of a disaster recovery scenario. You can use the vSphere Client, PowerCLI, or API to back up the vSphere Native Key Provider.

Required privilege: **Cryptographic operations > Manage key servers**

Note: In an Enhanced Link Mode configuration, you must perform the backup on the vCenter Server that the key provider belongs to.

vSphere Native Key Provider is backed up as part of the vCenter Server file-based backup. However, you must back up the vSphere Native Key Provider at least once before you can use it. When you create a vSphere Native Key Provider, it is not backed up.

A backup is necessary in case you must restore the configuration. To restore a vSphere Native Key Provider, see [Restore a vSphere Native Key Provider Using the vSphere Client](#).

Keep the backup file in a secure location. You can password-protect the backup when you create it. The backup file is in PKCS#12 format.

vCenter Server creates an alarm if a vSphere Native Key Provider has not been backed up. You can acknowledge the alarm, but it reappears every 24 hours until you have backed up the vSphere Native Key Provider.

1. Log in to the vCenter Server system with the vSphere Client.
2. Browse the inventory list and select the vCenter Server instance.
3. Click **Configure**, and under **Security** click **Key Providers**.
4. Select the vSphere Native Key Provider you want to back up.
A status of "Not backed up" appears for key providers that you have not backed up.
5. Click **Back Up**.
6. To password-protect the backup, check the **Protect Native Key Provider data with password** box.
 - a) Enter a password and save it in a secure location.
 - b) Check the **I have saved the password in a secure place** box, indicating that you have saved the password to a secure place.
7. Click **Back Up Key Provider**.
The backup file is in PKCS#12 format.
8. Save the backup file in a secure location.

The status of the vSphere Native Key Provider changes from Not Backed Up, to Warning, to Active. Warning indicates that the vCenter Server is still pushing the information to all the ESXi hosts in the data center. Active means that the information has been pushed to all the hosts.

To add vTPMs to virtual machines, see [Securing Virtual Machines with Virtual Trusted Platform Module](#). To encrypt virtual machines, see [Using Encryption in Your vSphere Environment](#).

Recovering a vSphere Native Key Provider

You can recover the vSphere Native Key Provider either through the vSphere Client or from the vCenter Server Appliance Backup.

When necessary, you can recover a vSphere Native Key Provider in the following ways.

1. If you do not need to rebuild your vCenter Server Appliance, use the vSphere Client to restore the key provider. See [Restore a vSphere Native Key Provider Using the vSphere Client](#).
2. If you must rebuild your vCenter Server Appliance, you must restore the key provider from your vCenter Server Appliance Backup. When you perform a vCenter Server Appliance Backup, it saves the Native Key Provider. See <https://blogs.vmware.com/vsphere/2018/05/vcenter-server-appliance-6-7-file-based-backup-and-restore-walkthroughs.html> for information about restoring the vCenter Server Appliance from backup.

Restore a vSphere Native Key Provider Using the vSphere Client

You can use the vSphere Client to restore the vSphere Native Key Provider.

- Required privilege: **Cryptographic operations > Manage key servers**
- The key provider backup file.
- The password for the key provider file, if you entered one when you backed up the key provider.

You can restore a vSphere Native Key Provider in case it was accidentally deleted or if you must perform a disaster recovery.

When you restore a vSphere Native Key Provider, you do not need to back up the key provider again. The initial backup suffices. Continue to maintain the backup file in a secure location.

Note: You can also use this task to configure vSphere Native Key Provider for vCenter Server systems in an Enhanced Linked Mode configuration. After you create the vSphere Native Key Provider on one vCenter Server system in the Enhanced Linked Mode configuration, use the **Restore** function to import the encrypted key file to the other ELM-connected vCenter Server systems.

1. Log in to the vCenter Server system with the vSphere Client.
2. Browse the inventory list and select the vCenter Server instance.
3. Click **Configure**, and under **Security** click **Key Providers**.
4. Select the vSphere Native Key Provider and click **Restore**.
5. Browse to the file location and select the backup encrypted key file.
The file was saved in PKCS#12 format.
6. Optional: If the file is password protected, enter the password.
7. Click **Next**.
8. Optional: If you decided to use this key provider only with TPM-protected ESXi hosts, select the check box.
9. Click **Finish**.

The vSphere Native Key Provider is imported to the vCenter Server. To use the vSphere Native Key Provider for encryption tasks, ensure that you first select it in the **Key Provider** pane and click **Set as Default**.

Update a vSphere Native Key Provider

As part of your regular key rotation plans, you can use PowerCLI to update a vSphere Native Key Provider.

- Required privilege: **Cryptographic operations > Manage key servers**
- PowerCLI 12.3.0

If you have a policy for key rotation, you can update the vSphere Native Key Provider and rekey the virtual machines that you encrypted with that key provider. You must use PowerCLI to update the vSphere Native Key Provider. You can also rekey the encrypted virtual machines without updating the key provider. In this case, only the virtual machine keys are changed. To rekey a virtual machine, see [Rekey an Encrypted Virtual Machine Using the vSphere Client](#).

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator user to the vCenter Server where you configured the vSphere Native Key Provider that you want to update.

```
Connect-VIServer -server VC_ip_address -User admin_user -Password 'password'
```

2. To get your vSphere Native Key Provider names, run the `Get-KeyProvider` cmdlet with the optional `Type` parameter.

```
Get-KeyProvider -Type NativeKeyProvider
```

3. To update the key provider, run the `Set-KeyProvider` cmdlet, specifying your key provider name and GUID. You can generate a GUID to use by running the `New-Guid` cmdlet.

```
Set-KeyProvider -KeyProvider KeyProvider_name -KeyId Guid
```

A warning appears about backing up the configuration.

4. To back up the key provider, run the `Export-KeyProvider` cmdlet.

```
Export-KeyProvider -KeyProvider KeyProvider_name -FilePath path_file_name
```

You can also back up the key provider using the vSphere Client. See [Back up a vSphere Native Key Provider](#).

When a key provider is updated, its status changes to Not Backed Up. After you back up the key provider, its status changes to Active.

Delete a vSphere Native Key Provider

You can use the vSphere Client to delete a vSphere Native Key Provider from vCenter Server.

Required privilege: **Cryptographic operations > Manage key servers**

Before you delete a vSphere Native Key Provider, rekey any encrypted virtual machines and datastores that were encrypted using that key provider to another key provider. See [Rekey an Encrypted Virtual Machine Using the vSphere Client](#).

In addition, maintain a backup of the vSphere Native Key Provider in case you must rekey an encrypted virtual machine after deleting the key provider.

After you delete a vSphere Native Key Provider, virtual machines that have vTPMs or that are encrypted continue to run. If you reboot the ESXi host, its encrypted virtual machines enter a locked state. After you unregister these virtual machines,

they enter a locked state when you try to re-register them. The only way to unlock the virtual machines is to restore the previous vSphere Native Key Provider.

1. Log in to the vCenter Server system with the vSphere Client.
2. Browse the inventory list and select the vCenter Server instance.
3. Click **Configure**, and under **Security** click **Key Providers**.
4. Select the key provider you want to delete.
5. Click **Delete**.
6. Read the warning message and slide the slider all the way to the right.
7. Click **Delete**.

The vSphere Native Key Provider is removed from the vCenter Server.

vSphere Trust Authority

With vSphere 7.0 and later, you can take advantage of VMware®vSphere Trust Authority™. vSphere Trust Authority is a foundational technology that enhances workload security. vSphere Trust Authority establishes a greater level of trust in your organization by associating an ESXi host's hardware root of trust to the workload itself.

vSphere Trust Authority Concepts and Features

vSphere Trust Authority secures your SDDC against malicious attacks by extending the trustworthiness of a trusted computing base to your organization's entire computing infrastructure. vSphere Trust Authority uses remote attestation and controlled access to advanced cryptographic capabilities.

vSphere Trust Authority is a set of services that satisfies high security requirements. With vSphere Trust Authority, you can set up and maintain a secure infrastructure. You can ensure that sensitive workloads run only on ESXi hosts proven to have booted authentic software.

How vSphere Trust Authority Protects Your Environment

You configure vSphere Trust Authority services to attest your ESXi hosts, which then become capable of performing trusted cryptographic operations.

vSphere Trust Authority uses remote attestation for ESXi hosts to prove the authenticity of their booted software. Attestation verifies that the ESXi hosts are running authentic VMware software, or VMware-signed partner software. Attestation relies on measurements that are rooted in a Trusted Platform Module (TPM) 2.0 chip installed in the ESXi host. In vSphere Trust Authority, an ESXi can access encryption keys and perform cryptographic operations only after it has been attested.

vSphere Trust Authority Glossary

vSphere Trust Authority introduces specific terms and definitions that are important to understand.

Table 260: vSphere Trust Authority Glossary

Term	Definition
VMware vSphere®Trust Authority™	Specifies a set of services that enables a Trusted Infrastructure. It is responsible for ensuring that ESXi hosts are running trusted software and for releasing encryption keys only to trusted ESXi hosts.

Term	Definition
vSphere Trust Authority Components	The vSphere Trust Authority components are: • Attestation Service • Key Provider Service
Attestation Service	Attests the state of a remote ESXi host. Uses TPM 2.0 to establish a hardware root of trust, and verifies software measurements against a list of administrator approved ESXi versions.
Key Provider Service	Encapsulates one or more key servers and exposes Trusted Key Providers that can be specified when encrypting virtual machines. Currently, key servers are limited to the KMIP protocol.
Trusted Infrastructure	A Trusted Infrastructure consists of: • A Trust Authority vCenter Server • A Workload vCenter Server • At least one vSphere Trust Authority Cluster (configured as part of the Trust Authority vCenter Server) • At least one Trusted Cluster (configured as part of the Workload vCenter Server) • Encrypted workload virtual machines running in the Trusted Cluster • At least one KMIP-compliant key management server Note: You must use separate vCenter Server systems for the Trust Authority Cluster and the Trusted Cluster.
Trust Authority Cluster	Consists of a vCenter Server cluster of ESXi hosts that run the vSphere Trust Authority components (the Attestation Service and the Key Provider Service).
Trust Authority Host	An ESXi host running vSphere Trust Authority components (the Attestation Service and the Key Provider Service).
Trusted Cluster	Consists of a vCenter Server cluster of Trusted ESXi hosts that are remotely attested by the Trust Authority Cluster. Although not strictly required, a configured Key Provider Service greatly increases the value provided by a Trusted Cluster.
Trusted Host	An ESXi host whose software has been validated by the Trust Authority Cluster Attestation Service. This host runs workload virtual machines that can be encrypted using Key Providers published by the Trust Authority Cluster Key Provider Service.
vSphere Encryption for Virtual Machines	With vSphere Virtual Machine Encryption, you can create encrypted virtual machines and encrypt existing virtual machines. vSphere Virtual Machine Encryption was introduced in vSphere 6.5. See vSphere Encryption Keys and Key Providers for differences in how key providers handle encryption keys.
Trusted Key Provider	A Key Provider that encapsulates a single encryption key on a key server. Access to the encryption key requires the Attestation Service to acknowledge that the ESXi software has been verified on the Trusted Host.
Standard Key Provider	A Key Provider that gets encryption keys directly from a key server, and distributes keys to the required hosts in a data center. Previously referred to in vSphere as KMS Cluster.
Key Server	A KMIP key management server (KMS) that is associated with a Key Provider.

Term	Definition
Workload vCenter Server	The vCenter Server that manages and is used to configure one or more Trusted Clusters.

vSphere Trust Authority Basics

With vSphere Trust Authority, you can:

- Provide ESXi hosts with a hardware root of trust and remote attestation capabilities
- Restrict encryption key management by releasing keys only to attested ESXi hosts
- Create a more secure administrative environment for managing trust
- Centralize management of multiple key servers
- Continue to perform cryptographic operations on virtual machines but with an enhanced level of encryption key management

In vSphere 6.5 and 6.7, virtual machine encryption depends on vCenter Server to obtain encryption keys from a key server and push them to ESXi hosts as needed. vCenter Server authenticates with the key server by using client and server certificates, which are stored in VMware Endpoint Certificate Store (VECS). Encryption keys that are sent from the key server pass through vCenter Server memory to the required ESXi hosts (with data encryption provided by TLS over the wire). In addition, vSphere depends on privilege checks in vCenter Server to validate user permissions and enforce key server access restrictions. Although this architecture is secure, it does not address the potential for a compromised vCenter Server, a rogue vCenter Server administrator, or a management or configuration error that might result in leaked or stolen secrets.


In vSphere 7.0 and later, vSphere Trust Authority addresses these problems. You can create a trusted computing base, which consists of a secure, manageable set of ESXi hosts. vSphere Trust Authority implements a remote attestation service for the ESXi hosts you want to trust. Furthermore, vSphere Trust Authority improves upon TPM 2.0 attestation support (added to vSphere beginning in the 6.7 release), to implement access restrictions on encryption keys and so better protect virtual machine workload secrets. In addition, vSphere Trust Authority allows only authorized Trust Authority administrators to configure vSphere Trust Authority services, and configure Trust Authority hosts. The Trust Authority administrator can be the same user as the vSphere administrator user, or a separate user.

In the end, vSphere Trust Authority enables you to run your workloads in a more safe and secure environment by:

- Detecting tampering
- Disallowing unauthorized changes
- Preventing malware and modifications
- Limiting sensitive workloads to run only on a verified, safe hardware and software stack

vSphere Trust Authority Architecture

The following figure shows a simplified view of the vSphere Trust Authority architecture.

Figure 118: vSphere Trust Authority Architecture

In this figure:


1. **vCenter Server systems**
Separate vCenter Server systems manage the Trust Authority Cluster and Trusted Clusters.
2. **Trust Authority Cluster**
Consists of the ESXi hosts that run the vSphere Trust Authority components.
3. **Key Servers**
Store encryption keys that are used by the Key Provider Service when encryption operations are performed. The key servers are external to vSphere Trust Authority.
4. **Trusted Clusters**
Consist of the ESXi Trusted Hosts, which have been remotely attested with a TPM, and that run encrypted workloads.
5. **Trust Authority Administrator**
Administrator who is a member of the vCenter Server TrustedAdmins group, and configures the Trusted Infrastructure. vSphere Trust Authority enables flexibility in how you designate Trust Authority administrators. The Trust Authority administrators in the figure can be separate users. It is also possible for the Trust Authority administrators to be the same user, using credentials that are linked across the vCenter Server systems. In this case, it is the same user and the same TrustedAdmins group.
6. **VM Administrator**
Administrator who has been granted privileges to manage the encrypted workload virtual machines on the Trusted Hosts.

vSphere Trust Authority Trusted Infrastructure

vSphere Trust Authority services, at least one external KMIP-compliant key server, the vCenter Server systems, and your ESXi hosts contribute to the Trusted Infrastructure.

What Is a Trusted Infrastructure

A Trusted Infrastructure consists of at least one vSphere Trust Authority Cluster, at least one Trusted Cluster, and at least one external KMIP-compliant key server. Each cluster contains ESXi hosts that run specific vSphere Trust Authority services, as shown in the following figure.

Figure 119: vSphere Trust Authority Services

Configuring the Trust Authority Cluster enables two services:

- Attestation Service
- Key Provider Service

When you configure vSphere Trust Authority, the ESXi hosts in the Trusted Cluster communicate with the Attestation Service. The Key Provider Service interposes between the Trusted Hosts and one or more trusted key providers.

Note: Currently, the ESXi hosts in the Trust Authority Cluster do not require a TPM. However, as a matter of best practice, consider installing new ESXi hosts with TPMs.

What Is the vSphere Trust Authority Attestation Service

The Attestation Service generates a signed document that contains assertions describing the binary and configuration state of the remote ESXi hosts in the Trusted Cluster. The Attestation Service attests the state of the ESXi hosts using a Trusted Platform Module (TPM) 2.0 chip as its basis for software measurement and reporting. The TPM on the remote ESXi host measures the software stack and sends the configuration data to the Attestation Service. The Attestation Service verifies that the software measurement signature can be attributed to a previously configured trusted TPM endorsement key (EK). The Attestation Service also ensures that the software measurement matches one of a set of previously blessed ESXi images. The Attestation Service signs a JSON Web Token (JWT) that it issues to the ESXi host, providing the assertions about the identity, validity, and configuration of the ESXi host.

What Is the vSphere Trust Authority Key Provider Service

The Key Provider Service removes the need for the vCenter Server and the ESXi hosts from requiring direct key server credentials. In vSphere Trust Authority, for an ESXi host to have access to an encryption key, it must authenticate with the Key Provider Service.

For the Key Provider Service to connect to a key server, the Trust Authority administrator must configure a trust setup. For most KMIP-compliant servers, configuring a trust setup involves configuring client and server certificates.

To ensure that the keys are released only to ESXi Trusted Hosts, the Key Provider Service acts as a gatekeeper to the key servers. The Key Provider Service hides the key server specifics from the rest of the data center software stack by using the concept of a trusted key provider. Each trusted key provider has a single configured primary encryption key, and

references one or more key servers. The Key Provider Service can have several configured trusted key providers. For example, you might want to have a separate trusted key provider for each department in an organization. Each trusted key provider uses a different primary key, but can reference the same backing key server.

After you create a trusted key provider, the Key Provider Service can accept requests from the ESXi Trusted Hosts to run cryptographic operations against that trusted key provider.

When an ESXi Trusted Host requests operations against a trusted key provider, the Key Provider Service makes sure that the ESXi host that is trying to obtain the encryption key is attested. After passing all the checks, the ESXi Trusted Host receives encryption keys from the Key Provider Service.

What Ports Are Used by vSphere Trust Authority

The vSphere Trust Authority services listen for connections behind the ESXi host's reverse proxy. All communication occurs over HTTPS on port 443.

What Are vSphere Trust Authority Trusted Hosts

The ESXi Trusted Hosts are configured to use trusted key providers to perform cryptographic operations. The ESXi Trusted Hosts perform key operations by communicating to the Key Provider Service and the Attestation Service. For authentication and authorization, the ESXi Trusted Hosts use a token obtained from the Attestation Service. To get a valid token, the ESXi Trusted Host must successfully attest to the Attestation Service. The token contains certain claims that are used to decide whether the ESXi Trusted Host is authorized to access a trusted key provider.

vSphere Trust Authority and Key Server Requirement

vSphere Trust Authority requires the use of at least one key server. In previous vSphere releases, a key server was called a Key Management Server or KMS. Currently, vSphere Virtual Machine Encryption supports KMIP 1.1 compliant key servers.

How Does vSphere Trust Authority Store Configuration and State Information

vCenter Server is mostly a pass-through service for vSphere Trust Authority configuration and state information. Most vSphere Trust Authority configuration and state information is stored on the ESXi hosts in the ConfigStore database. Some state information is stored in the vCenter Server database as well.

Note: Because most vSphere Trust Authority configuration information is stored on the ESXi hosts, the vCenter Server file-based backup mechanism does not back up this information. To ensure the configuration information for your vSphere Trust Authority deployment is saved, see [Backing Up the vSphere Trust Authority Configuration](#).

How Does vSphere Trust Authority Integrate with vCenter Server

You configure separate vCenter Server instances to manage the Trust Authority Cluster and Trusted Cluster. See [Configuring vSphere Trust Authority](#).

On a Trusted Cluster, the vCenter Server manages the Trust Authority API calls and passes them through to the ESXi hosts. The vCenter Server replicates the API calls across all ESXi hosts in the Trusted Cluster.

After you configure vSphere Trust Authority initially, you can add or remove ESXi hosts to or from a Trust Authority Cluster or a Trusted Cluster. See [Adding and Removing vSphere Trust Authority Hosts](#).

vSphere Trust Authority Process Flows

Understanding vSphere Trust Authority process flows is essential to learning how to configure and manage your Trusted Infrastructure.

How Do You Configure vSphere Trust Authority

vSphere Trust Authority is not activated by default. You must manually configure vSphere Trust Authority in your environment. See [Configuring vSphere Trust Authority](#).

When you configure vSphere Trust Authority, you must specify which versions of ESXi software the Attestation Service accepts, and also which Trusted Platform Modules (TPMs) are trustworthy.

TPM and Attestation in vSphere Trust Authority

This guide uses the following definitions when discussing TPMs and attestation.

Table 261: TPM and Attestation Glossary

Term	Definition
Endorsement Key (EK)	A TPM is manufactured with an RSA public/private key pair built into the hardware, called the endorsement key (EK). The EK is unique to a particular TPM.
EK Public Key	The public portion of the EK key pair.
EK Private Key	The private portion of the EK key pair.
EK Certificate	The EK public key wrapped with a signature. The EK certificate is created by the TPM manufacturer who uses its Certificate Authority private key to sign the EK public key. Not all TPMs contain an EK certificate. In this case, the EK public key is not signed.
TPM Attestation	The ability of the Attestation Service to verify the software that is running on a remote host. TPM attestation is accomplished through cryptographic measurements made by the TPM while the remote host starts, and is relayed to the Attestation Service on request. The Attestation Service establishes trust in the TPM through either the EK public key or the EK certificate.

Configuring TPM Trust on the Trusted Hosts

An ESXi Trusted Host must contain a TPM. A TPM is manufactured with a public/private key pair built into the hardware, called the endorsement key (EK). Although TPM 2.0 permits many key/certificate pairs, the most common is an RSA-2048 key pair. When a TPM EK public key is signed by a CA, the result is the EK certificate. The TPM manufacturer typically pre-generates at least one EK, signs the public key with a Certificate Authority, and embeds the signed certificate in the TPM's non-volatile memory.

You can configure the Attestation Service to trust TPMs as follows:

- Trust all CA certificates with which the manufacturer signed the TPM (the EK public key). The default setting for the Attestation Service is to trust CA certificates. In this approach, the same CA certificate covers many ESXi hosts, and so reduces your administrative overhead.
- Trust the ESXi host's TPM CA certificate and EK public key. The latter can be either the EK certificate or the EK public key. Although this approach provides more security, it requires you to configure information about each Trusted Host.
- Some TPMs do not contain an EK certificate. In this case, you trust the EK public key.

Deciding to trust all TPM CA certificates is operationally convenient. You configure new certificates only when you add a new class of hardware to your data center. By trusting individual EK certificates, you can limit access to specific ESXi hosts.

You can also decide not to trust TPM CA certificates. Though an uncommon situation, you can use this configuration when an EK is not signed by a CA. Currently, this functionality is not fully implemented.

Note: Some TPMs do not include EK certificates. If you want to trust individual ESXi hosts, the TPM must include an EK certificate.

How Does vSphere Trust Authority Attest TPMs

To begin the attestation process, the ESXi Trusted Host in the Trusted Cluster sends the pre-configured EK public key and EK certificate to the Attestation Service on the Trust Authority Cluster. When the Attestation Service receives the request, it looks up the EK in its configuration, which can be the EK public key or the EK certificate, or both, depending on the configuration. If no case is valid, the Attestation Service rejects the attestation request.

The EK is not directly used for signing, so an Attestation Key (AK or AIK) is negotiated. The negotiation protocol ensures that a newly created AK is bound to the previously verified EK, preventing a man-in-the-middle situation, or an impersonator. After an AK is negotiated, it is reused on future attestation requests rather than generating a new one each time.


The ESXi Trusted Host reads the Quote and PCR values from the TPM. The Quote is signed by the AK. The ESXi Trusted Host also reads the TCG Event Log, which includes all the events that resulted in the current PCR state. This TPM information is sent to the Attestation Service for validation. The Attestation Service verifies the PCR values using the event log.

How Do Key Providers Work with Key Servers

The Key Provider Service uses the concept of a trusted key provider to hide the key server specifics from the rest of the data center software. Each trusted key provider has a single configured primary encryption key, and references one or more key servers. The primary encryption key is present in the key servers. As part of configuring vSphere Trust Authority, you must provision the primary key as a separate activity and activate it. The Key Provider Service can have several configured trusted key providers. Each trusted key provider uses a different primary key, but can reference the same backing key server.

When a new trusted key provider is added, the Trust Authority administrator must specify the key server and an existing key identifier on that key server.

The following figure shows the relationship between the Key Provider Service and key servers.

Figure 120: Key Provider and Key Server

After you configure a trusted key provider for a Trusted Cluster, the Key Provider Service can accept requests to run cryptographic operations against that trusted key provider. For example, in this figure, three trusted key providers are configured, two for KMS-1, and one for KMS-2. The Trusted Host requests an encryption operation against key-provider-2. The Trusted Host requests an encryption key to be generated and returned, and uses this encryption key to perform encryption operations.

The Key Provider Service uses the primary key referenced by key-provider-2 to encrypt the specified plaintext data and return the corresponding ciphertext. Later, the Trusted Host can provide the same ciphertext to a decrypt operation and get back the original plaintext.

vSphere Trust Authority Authentication and Authorization

vSphere Trust Authority administrative operations require a user that is a member of the TrustedAdmins group. Having only Trust Authority administrator privileges is not sufficient to perform all administrative operations that involve the ESXi hosts. For more information, see [Prerequisites and Required Privileges for vSphere Trust Authority](#).

Adding a Trusted Host to a Trusted Cluster

The steps to add ESXi hosts initially to the Trusted Cluster are described in [Configuring vSphere Trust Authority](#).

Later, if you want to add ESXi hosts to the Trusted Cluster, the workflow is different. See [Adding and Removing vSphere Trust Authority Hosts](#).

When you initially add ESXi hosts to the Trusted Cluster, you must gather the following information:

- TPM certificate for each type of hardware in the cluster
- ESXi image for each version of ESXi in the cluster
- vCenter Server principal information

If you later add ESXi hosts to a Trusted Cluster, you might need to collect some additional information. That is, if the new ESXi hosts differ in hardware or ESXi version from the original hosts, you must collect the new ESXi host information and

import it to the Trust Authority Cluster. You only must collect the vCenter Server principal information one time per vCenter Server system.

vSphere Trust Authority Topology

vSphere Trust Authority requires separate vCenter Server systems for the Trust Authority Cluster and the Trusted Cluster.

The Trust Authority Cluster is configured and managed on an independent, isolated vCenter Server. The vCenter Server of the Trust Authority Cluster cannot also be the vCenter Server of the Trusted Cluster. The Trusted Cluster must have its own, separate vCenter Server. A single vCenter Server can manage multiple Trusted Clusters. Multiple vCenter Server systems for Trusted Clusters can participate in enhanced linked mode. The vCenter Server for the Trust Authority Cluster cannot participate in enhanced linked mode with other Trust Authority Cluster vCenter Server systems or Trusted Cluster vCenter Server systems.

The Trust Authority administrator manages the Trust Authority Cluster and its associated vCenter Server independently from other vCenter Server instances, because this approach provides the best security isolation.

The Trust Authority administrator documents or publishes the hostnames and SSL certificates that Trusted Cluster administrators use to configure their clusters. The Trust Authority administrator also provisions trusted key providers for the organization and its departments, or even individual administrators.

You cannot deploy vSphere Trust Authority services directly on the Trusted Cluster managed by the Workload vCenter Server, because the workload administrator has high privilege access to the ESXi hosts. This type of deployment does not achieve the necessary separation of roles that is required to meet the security objectives of vSphere Trust Authority.

Prerequisites and Required Privileges for vSphere Trust Authority

You must consider hardware and software requirements when configuring vSphere Trust Authority. You must set cryptographic privileges and roles to use encryption. The user who performs vSphere Trust Authority tasks must have the appropriate privileges.

Requirements for vSphere Trust Authority

To use vSphere Trust Authority, your vSphere environment must meet these requirements:

- ESXi Trusted Host hardware requirements:
 - TPM 2.0
 - Secure boot must be enabled
 - EFI firmware
- Component requirements:
 - vCenter Server 7.0 or later
 - A dedicated vCenter Server system for the vSphere Trust Authority Cluster and ESXi hosts
 - A separate vCenter Server system for the Trusted Cluster and ESXi Trusted Hosts
 - A key server (called a Key Management Server, or KMS, in prior vSphere releases)
- Virtual machine requirements:
 - EFI firmware
 - Secure Boot Enabled

Note: Before you can begin configuring vSphere Trust Authority, ensure that you have set up your vCenter Server systems for the Trust Authority Cluster and Trusted Cluster, and added ESXi hosts to each cluster.

vSphere Trust Authority and Cryptography Privileges

vSphere Trust Authority does not introduce any new cryptography privileges. The same cryptography privileges described in [Using Cryptography Privileges and Roles](#) apply to vSphere Trust Authority.

vSphere Trust Authority and Host Encryption Mode

vSphere Trust Authority does not introduce any new requirements for enabling host encryption mode on the ESXi Trusted Hosts. See [Prerequisites and Required Privileges for Virtual Machine Encryption Tasks](#) for more information about host encryption mode.

Using the vSphere Trust Authority Roles and the TrustedAdmins Group

vSphere Trust Authority operations require a user that is a member of the TrustedAdmins group. This user is called the Trust Authority administrator. vSphere administrators must either add themselves to the TrustedAdmins group or add other users to the group to gain the Trusted Infrastructure administrator role. The Trusted Infrastructure administrator role is necessary for vCenter Server authorization. The TrustedAdmins group is necessary for authentication on the ESXi hosts that are part of the Trusted Infrastructure. Users with the **Cryptographic Operations > Register host** privilege on ESXi hosts can manage the Trusted Cluster. The vCenter Server permissions are not propagated to the Trust Authority hosts, only to the Trusted Hosts. Only members of the TrustedAdmins group are granted privileges on the Trust Authority hosts. Group membership is verified on the ESXi host itself.

Note: vSphere administrators and members of the Administrators group are assigned the Trusted Infrastructure administrator role, but this role by itself does not permit a user to perform vSphere Trust Authority operations. Membership in the TrustedAdmins group is also required.

After vSphere Trust Authority is enabled, Trust Authority administrators can assign trusted key providers to Trusted Hosts. Those Trusted Hosts can then use the trusted key providers to perform cryptographic tasks.

In addition to the Trusted Infrastructure administrator role, vSphere Trust Authority provides the No Trusted Infrastructure administrator role, which contains all privileges in vCenter Server except the ones that call the vSphere Trust Authority APIs.

vSphere Trust Authority groups, roles, and users function as follows:

- On first boot, vSphere grants the TrustedAdmins group the Trusted Infrastructure administrator role, which has global permissions.
- The Trusted Infrastructure administrator role is a system role that has the required privileges to call the vSphere Trust Authority APIs (`TrustedAdmin.*`), and the system privileges **System.Read**, **System.View**, and **System.Anonymous** to view inventory objects.
- The No Trusted Infrastructure administrator role is a system role that contains all privileges in vCenter Server except the ones to call the vSphere Trust Authority APIs. Adding new privileges to vCenter Server also adds them to the No Trusted Infrastructure administrator role. (The No Trusted Infrastructure administrator role is similar to the No cryptography administrator role.)
- The vSphere Trust Authority privileges (`TrustedAdmin.*` APIs) are not included in the No cryptography administrator role, preventing users with this role from setting up a Trusted Infrastructure or performing cryptographic operations.

The use cases for these users, groups, and roles, are shown in the following table.

Table 262: vSphere Trust Authority Users, Groups, and Roles

User, Group, or Role	Can Call vSphere Trust Authority vCenter Server API (Includes Calls to vSphere Trust Authority ESXi API)	Can Call vSphere Trust Authority vCenter Server API (Does Not Include Calls to vSphere Trust Authority ESXi API)	Can Perform Host Operations in Cluster Not Related to vSphere Trust Authority	Comment
User in both Administrators@system.domain group and TrustedAdmins@system.domain group	Yes	Yes	Yes	NA
User in TrustedAdmins@system.domain group only	Yes	Yes	No	Such a user cannot perform regular cluster management operations.
User in Administrators@system.domain group only	Yes	No	Yes	NA
User with Trusted Infrastructure administrator role but not in TrustedAdmins@system.domain group	Yes	No	No	The ESXi host checks the group membership of the user to grant permissions.
User with No Trusted Infrastructure administrator role only	No	No	Yes	Such a user is similar to an administrator who cannot perform vSphere Trust Authority operations.

vSphere Trust Authority Best Practices, Caveats, and Interoperability

The vSphere Trust Authority architecture results in some additional recommendations. As you are planning your vSphere Trust Authority strategy, consider interoperability limitations.

Trusted Infrastructure Interoperability

For ESXi versions, the Attestation Service is backward and forward compatible. For example, you can have a cluster of ESXi hosts running ESXi 7.0 in the vSphere Trust Authority Cluster, and upgrade or patch ESXi hosts in the Trusted Cluster to a newer ESXi version. Similarly, you can upgrade or patch the ESXi hosts in the Trust Authority Cluster while keeping the ESXi hosts in the Trusted Cluster at the current version.

You cannot have a cluster function as both a Trust Authority Cluster and a Trusted Cluster. This configuration is not supported.

Trusted Cluster Configuration Limitation

You can configure only one Trust Authority Cluster per workload vCenter Server. A Trusted Cluster cannot be configured to reference multiple Trust Authority Clusters.

vSphere Features Supported in vSphere Trust Authority

vSphere Trust Authority supports the following:

- vCenter High Availability (vCenter HA)
- VMware vSphere High Availability
- DRS
- DPM
- SRM, with the following understanding:
 - SRM with array-based replication is supported, if the same vSphere Trust Authority services configuration is available on the recovery side.
 - SPPG
- VADP
 - Support is the same as with standard encryption. Hot-add and NFC modes are supported, but not SAN mode. Backups are decrypted. VADP partners have the option of recovering the backed-up virtual machine with the same encryption key as the original virtual machine.
- vSAN
 - Virtual machine encryption is fully supported on top of vSAN.
- OVF
 - Encrypted virtual machines cannot be exported to OVF. However, virtual machines can be encrypted while being imported from an OVF.
- VVol

vSphere Features Not Supported in vSphere Trust Authority

Currently, vSphere Trust Authority does not support the following:

- vSAN data-at-rest encryption
- First Class Disk (FCD) encryption
- vSphere Replication
- vSphere Host Profiles

vSphere Trust Authority Life Cycle

The vSphere Trust Authority services are packaged and installed as part of the base ESXi image.

Starting and Stopping vSphere Trust Authority Services

In the vSphere Client, you can start, stop, and restart vSphere Trust Authority services that are running on an ESXi host. You can restart services upon a configuration change or if you suspect functional or performance problems. To restart the service on an ESXi Trusted Host, you must log in to the host itself to restart the service. See [Start, Stop, and Restart vSphere Trust Authority Services](#).

Upgrading and Patching vSphere Trust Authority

Each time you upgrade or patch an ESXi Trusted Host, you must update the vSphere Trust Authority Cluster with the new ESXi version information. One way to do so is to upgrade or patch a test ESXi host, export the ESXi base image information, import the image file to the Trust Authority Cluster, then upgrade or patch the ESXi Trusted Hosts.

Best Practices for Upgrading vSphere Trust Authority

Best practice for upgrading a vSphere Trust Authority infrastructure is to upgrade the Trust Authority vCenter Server and Trust Authority Hosts first. In this way, you get the most benefit from the latest vSphere Trust Authority features. However, you can perform separate, standalone upgrades of vCenter Server and ESXi hosts to fit specific business reasons.

In general, follow this order for upgrading your vSphere Trust Authority infrastructure:

1. Upgrade the Trust Authority Cluster vCenter Server.
2. Upgrade the Trust Authority Hosts.
3. Upgrade the Trusted Cluster vCenter Server.
4. Upgrade the Trusted Hosts.

To ensure a smooth process, upgrade your Trust Authority Hosts and Trusted Hosts gradually, one-by-one.

Upgrading vSphere Trust Authority with Quick Booted ESXi Trusted Hosts

Quick Boot is a setting that you can use with clusters that you manage with vSphere Lifecycle Manager images and vSphere Lifecycle Manager baselines. Using Quick Boot optimizes the ESXi host patching and upgrade operations.

When you upgrade an ESXi host using the Quick Boot optimization, host attestation continues to report the previously booted ESXi version in the root of trust measurement.

Thus, when you upgrade an ESXi trusted host that is enabled for Quick Boot, and that is part of a vSphere Trust Authority deployment, pay attention to the following:

1. Do not remove the ESXi base image version that you initially trusted from the Attestation Service until all the ESXi hosts have completed a full reboot after upgrade. (If you need to reboot the host, disable Quick Boot.)
2. If you have used Quick Boot for multiple upgrades, and want to remove an intermediate ESXi version that is no longer trustworthy, use the `base-images` API to confirm the ESXi version that you last attested.
3. When you export the ESXi base image of an ESXi host enabled for Quick Boot, a message appears that the host was upgraded by Quick Boot. The resulting file contains the latest metadata of the ESXi base image.

If you upgrade a regular cluster's hosts using Quick Boot, then later add that cluster to vSphere Trust Authority, the hosts do not attest until you reboot them. The attestation failure occurs because the exported ESXi base image file of the hosts contains only the latest metadata, whereas the host attestation is based on the metadata from the last full boot. Thus, if the cluster is not part of vSphere Trust Authority and the ESXi base image metadata is not imported to vSphere Trust Authority for the full boot, attestation fails.

To get the base image, you can use the following PowerCLI commands.

```
$vTA = Get-TrustAuthorityCluster -name trustedCluster
$bm = Get-TrustAuthorityVMHostBaseImage $vTA
$bm | select *
```

Troubleshooting vSphere Trust Authority Upgrade Problems

If you encounter an unsuccessful upgrade of a Trust Authority Host, follow these steps.

1. Remove the Trust Authority Host from the Trusted Cluster.
2. Revert to the previous version of ESXi.
3. Re-add the Trust Authority Host to the cluster as described in the VMware knowledge base article at <https://kb.vmware.com/s/article/77234>.
4. Verify that the Trust Authority Host's configuration is consistent with the other Trust Authority Hosts in the Trust Authority Cluster. See [Check Trusted Cluster Health](#).

When you upgrade to a new version of ESXi on a Trusted Host, attestation fails until you update the Trust Authority Cluster with the new ESXi base image information. This behavior is to be expected. You can no longer encrypt virtual

machines or use existing virtual machines that were encrypted before upgrade until you fix the problem. Attestation error messages appear in the vSphere Client **Recent Tasks** pane and the `attestd.log`, `kmxa.log`, and `vpzd.log` files.

To correct the problem, follow these steps.

1. Run the `Export-VMHostImageDb` cmdlet to re-export the ESXi base images. See Step 5 in [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).
2. Run the `New-TrustAuthorityVMHostBaseImage` cmdlet to reimport the new base image to the vCenter Server of the Trust Authority Cluster. See Step 8 in [Import the Trusted Host Information to the Trust Authority Cluster](#).
3. If you no longer must attest the older versions of ESXi (all the Trusted Hosts have been upgraded), run the `Remove-TrustAuthorityVMHostBaseImage` cmdlet to remove the versions. For example:

```
$vTA = Get-TrustAuthorityCluster 'vTA Cluster'
$baseImages = Get-TrustAuthorityVMHostBaseImage -TrustAuthorityCluster $vTA
Remove-TrustAuthorityVMHostBaseImage -VMHostBaseImage $baseImages
```

Backing Up the vSphere Trust Authority Configuration

Because most vSphere Trust Authority configuration information is stored on the ESXi hosts, the vCenter Server Backup does not back up this vSphere Trust Authority information. See [Backing Up the vSphere Trust Authority Configuration](#).

Configuring vSphere Trust Authority

vSphere Trust Authority is not enabled by default. You must configure your environment for vSphere Trust Authority before you can start using it.

You enable vSphere Trust Authority services on a dedicated vCenter Server cluster, known as the vSphere Trust Authority Cluster. The Trust Authority Cluster acts as a centralized, secure management platform. You then enable a workload vCenter Server cluster as the Trusted Cluster. The Trusted Cluster contains the ESXi Trusted Hosts.


The Trust Authority Cluster attests the ESXi hosts in the Trusted Cluster remotely. The Trust Authority Cluster releases encryption keys only to attested ESXi hosts in the Trusted Cluster to encrypt virtual machines and virtual disks using trusted key providers.

Before you begin configuring vSphere Trust Authority, see [Prerequisites and Required Privileges for vSphere Trust Authority](#) for information on the required setup of vCenter Server systems and ESXi hosts.

You manage different aspects of vSphere Trust Authority in the following ways.

- Configure the vSphere Trust Authority services and trusted connections using PowerCLI cmdlets or the vSphere APIs. See *VMware PowerCLI Cmdlets Reference* and *vSphere Automation SDKs Programming Guide*.
- Manage the configuration of trusted key providers using the PowerCLI cmdlets or from the vSphere Client.
- Perform encryption workflows, as in prior vSphere releases, using the vSphere Client and APIs.

Figure 121: vSphere Trust Authority Workflow



To configure and manage vSphere Trust Authority, you use VMware PowerCLI, though some functionality is available in the vSphere Client.

When you configure vSphere Trust Authority, you must complete setup tasks on both the Trust Authority Cluster and the Trusted Cluster. Some of these tasks are order-specific. Use the task sequence outlined in this guide.

Note: When adding more ESXi hosts to the Trusted Cluster after completing the initial vSphere Trust Authority setup, you might need to export and import the Trusted Host information again. That is, if the new ESXi hosts differ from the original hosts, you must collect the new ESXi host information and import it to the Trust Authority Cluster. See [Adding and Removing vSphere Trust Authority Hosts](#).

Set up Your Workstation to Configure vSphere Trust Authority

To configure a vSphere Trust Authority deployment, you must first prepare a workstation with the necessary software and setup.

Perform the following steps on a workstation that has access to your vSphere Trust Authority environment.

1. Install PowerCLI 12.1.0 or later. See *PowerCLI User's Guide*.
2. Verify that Microsoft .NET Framework 4.8 or later is installed.
3. Create a local folder in which to save the Trust Authority information that you export as files.

Continue with [Enable the Trust Authority Administrator](#).

Enable the Trust Authority Administrator

To enable vSphere Trust Authority, you must add a user to the vSphere TrustedAdmins group. This user becomes the Trust Authority administrator. You use the Trust Authority administrator for most vSphere Trust Authority configuration tasks.

Either create or a user, or identify an existing user, to be the Trust Authority administrator.

Use a separate user from the vCenter Server administrator as your Trust Authority administrator. Having a separate user enhances the security of your environment. You must enable a Trust Authority administrator for both the Trust Authority Cluster and the Trusted Cluster.

1. Connect to the vCenter Server of the Trust Authority Cluster by using the vSphere Client.
2. Log in as an administrator.
3. From the **Home** menu, select **Administration**.
4. Under **Single Sign On**, click **Users and Groups**.
5. Click **Groups** and click the **TrustedAdmins** group.
If the TrustedAdmins group does not appear initially, use the **Filter** icon to filter for it, or navigate through the groups by clicking the right arrow at the bottom of the pane.
6. In the **Group Members** area, click **Add Members**.
Make sure that the local identity source is selected (vsphere.local is the default, but you might have selected a different domain during installation), and search for the member (user) you want to add to the group as the Trust Authority administrator.
7. Select the member.
8. Click **Save**.
9. Repeat steps 1 through 8 for the vCenter Server of the Trusted Cluster.

Continue with [Enable the Trust Authority State](#).

Enable the Trust Authority State

Making a vCenter Server cluster into a vSphere Trust Authority Cluster (also called enabling the Trust Authority State) starts the required Trust Authority services on the ESXi hosts in the cluster.

- [Enable the Trust Authority Administrator](#).

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as the Trust Authority administrator user to the vCenter Server of the Trust Authority Cluster.

```
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

2. To check the current state of the cluster, run the `Get-TrustAuthorityCluster` cmdlet.
For example, this command shows the cluster, vTA Cluster, and that its state is disabled.

```
Get-TrustAuthorityCluster
```

Name	State	Id
----	-----	--
vTA Cluster	Disabled	TrustAuthorityCluster-domain-c8

The output shows either Disabled or Enabled in the State column for each cluster found. Disabled means that the Trust Authority services are not running.

3. To enable the Trust Authority Cluster, run the `Set-TrustAuthorityCluster` cmdlet.
For example, this command enables the cluster vTA Cluster.

```
Set-TrustAuthorityCluster -TrustAuthorityCluster 'vTA Cluster' -State Enabled
```

The system responds with a confirmation prompt.

```
Confirmation
Setting TrustAuthorityCluster 'vTA Cluster' with new State 'Enabled'. Do you want to proceed?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):
```

4. At the confirmation prompt, press Enter. (The default is Y.)

The output shows the state of the cluster. For example, the following shows that cluster vTA Cluster has been enabled:

Name	State	Id
----	-----	--
vTA Cluster	Enabled	TrustAuthorityCluster-domain-c8

Two services start on the ESXi hosts in the Trust Authority Cluster: the Attestation Service and the Key Provider Service.

Enable the Trusted State on the Trust Authority Cluster

This example shows how to use PowerCLI to enable services on the Trust Authority Cluster. The following table shows the example components and values that are used.

Table 263: Example vSphere Trust Authority Setup

Component	Value
vCenter Server for Trust Authority Cluster	192.168.210.22
Trust Authority Cluster name	vTA Cluster
Trust Authority administrator	trustedadmin@vsphere.local

```
PS C:\Users\Administrator.CORP> Disconnect-VIServer -server * -Confirm:$false
PS C:\Users\Administrator.CORP> Connect-VIServer -server 192.168.210.22 -User trustedadmin@vsphere.local -Password 'VMware1!'
```

Name	Port	User
----	----	----
192.168.210.22	443	VSPHERE.LOCAL\trustedadmin

```
PS C:\Users\Administrator.CORP> Get-TrustAuthorityCluster
```

Name	State	Id
----	-----	--
vTA Cluster	Disabled	TrustAuthorityCluster-domain-c8

```
PS C:\Users\Administrator.CORP> Set-TrustAuthorityCluster -TrustAuthorityCluster 'vTA Cluster' -
State Enabled
```

Confirmation

Setting TrustAuthorityCluster 'vTA Cluster' with new State 'Enabled'. Do you want to proceed?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): y

Name	State	Id
----	-----	--
vTA Cluster	Enabled	TrustAuthorityCluster-domain-c8

Continue with [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).

Collect Information About ESXi Hosts and vCenter Server to Be Trusted

To establish trust, the vSphere Trust Authority Cluster requires information about the Trusted Cluster's ESXi hosts and vCenter Server. You export this information as files for importing into the Trust Authority Cluster. You must ensure to keep these files confidential and transport them securely.

- Identify the ESXi versions and TPM hardware types that are in the Trusted Cluster, and whether you want to trust all TPM hardware types, only certain ones, or individual hosts.
- On the machine from which you run the PowerCLI cmdlets, create a local folder in which to save the information you export as files.
- [Enable the Trust Authority Administrator](#).
- [Enable the Trust Authority State](#).

You use vSphere Trust Authority PowerCLI cmdlets to export the following information as files from the ESXi hosts in the Trusted Cluster for the Trust Authority Cluster to know what software and hardware to trust.

- ESXi version
- TPM manufacturer (CA certificate)
- (Optional) Individual TPM (EK certificate)

Note: Store these exported files in a secure location, in case you must restore the vSphere Trust Authority configuration.

If you have hosts of the same type and vendor, and manufactured during the same timeframe and location, you might be able to trust all TPMs by obtaining the CA certificate of only one of the TPMs. To trust an individual TPM, you obtain the EK certificate of the TPM.

You must also obtain the principal information from the Trusted Cluster's vCenter Server. The principal information contains the vpxd solution user and its certificate chain. The principal information enables the Trusted Cluster's vCenter Server to discover the available trusted key providers configured on the Trust Authority Cluster.

To configure vSphere Trust Authority initially, you must collect the ESXi version and TPM information. Also, you must collect the ESXi version each time after you deploy a new version of ESXi, including when you upgrade or apply a patch.

You collect the vCenter Server principal information only one time per vCenter Server system.

1. In a PowerCLI session, run the following commands to disconnect any current connection and connect as the root user to one of the ESXi hosts in the Trusted Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server host_ip_address -User root -Password 'password'
```

2. Run the `Get-VMHost` cmdlet to confirm the ESXi host.

```
Get-VMHost
```

The host information is displayed.

3. Assign `Get-VMHost` to a variable.

For example:

```
$vmhost = Get-VMHost
```

4. Run the `Export-Tpm2CACertificate` cmdlet to export the CA certificate of a given TPM manufacturer.

- a) Assign `Get-Tpm2EndorsementKey -VMHost $vmhost` to a variable.

For example, this command assigns `Get-Tpm2EndorsementKey -VMHost $vmhost` to the variable `$tpm2`.

```
$tpm2 = Get-Tpm2EndorsementKey -VMHost $vmhost
```

- b) Run the `Export-Tpm2CACertificate` cmdlet.

For example, this command exports the TPM certificate to the `cacert.zip` file. Ensure that the destination directory exists before running this command.

```
Export-Tpm2CACertificate -Tpm2EndorsementKey $tpm2 -FilePath C:\vta\cacert.zip
```

The file is created.

- c) Repeat for each TPM hardware type in the cluster that you want to trust. Use a different file name for each TPM hardware type so that you do not overwrite a previously exported file.

5. Run the `Export-VMHostImageDb` cmdlet to export the ESXi host description of software (the ESXi image).

For example, this command exports the information to the `image.tgz` file. Ensure that the destination directory exists before running this command.

```
Export-VMHostImageDb -VMHost $vmhost -FilePath C:\vta\image.tgz
```

Note: The `Export-VMHostImageDb` cmdlet also works if you prefer to log in to the vCenter Server of the Trusted Cluster.

The file is created.

Repeat for each ESXi version in the cluster that you want to trust. Use a different file name for each version so that you do not overwrite a previously exported file.

6. Export the Trusted Cluster's vCenter Server principal information.

- a) Disconnect from the ESXi host.

```
Disconnect-VIServer -server * -Confirm:$false
```

- b) Connect to the vCenter Server of the Trusted Cluster using the Trust Authority administrator user. (Alternatively, you can use a user that has **Administrator** privileges.)

```
Connect-VIServer -server TrustedCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

- c) To export the Trusted Cluster's vCenter Server principal information, run the `Export-TrustedPrincipal` cmdlet. For example, this command exports the information to the `principal.json` file. Ensure that the destination directory exists before running this command.

```
Export-TrustedPrincipal -FilePath C:\vta\principal.json
```

The file is created.

7. Optional: If you want to trust an individual host, you must export the TPM EK public key certificate.

See [Export and Import a TPM Endorsement Key Certificate](#).

The following files are created:

- TPM CA certificate file (.zip file extension)
- ESXi image file (.tgz file extension)
- vCenter Server principal file (.json file extension)

Collecting Information About ESXi Hosts and vCenter Server to Be Trusted

This example shows how to use PowerCLI to export the ESXi host information and the vCenter Server Principal. The following table shows the example components and values that are used.

Table 264: Example vSphere Trust Authority Setup

Component	Value
ESXi host in Trusted Cluster	192.168.110.51
vCenter Server for Trusted Cluster	192.168.110.22
Variable \$vmhost	Get-VMHost
Variable \$tpm2	Get-Tpm2EndorsementKey -VMHost \$vmhost
Trust Authority administrator	trustedadmin@vsphere.local
Local directory to contain output files	C:\vta

```
PS C:\Users\Administrator.CORP> Connect-VIServer -server 192.168.110.51 -User root -Password 'VMware1!'
```

```
Name                               Port  User
----                               -
192.168.110.51                    443  root
```

```
PS C:\Users\Administrator.CORP> Get-VMHost
```

```
Name                               ConnectionState PowerState NumCpu CpuUsageMhz CpuTotalMhz MemoryUsageGB Memory-
TotalGB Version
----                               -
192.168.110.51                    Connected      PoweredOn    4      200      9576      1.614
7.999  7.0.0
```

```
PS C:\Users\Administrator.CORP> $vmhost = Get-VMHost
PS C:\Users\Administrator.CORP> $tpm2 = Get-Tpm2EndorsementKey -VMHost $vmhost
PS C:\> Export-Tpm2CACertificate -Tpm2EndorsementKey $tpm2 -FilePath C:\vta\cacert.zip
```

```
Mode                               LastWriteTime           Length Name
----                               -
-a----          10/8/2019   6:55 PM           1004 cacert.zip
```

```
PS C:\Users\Administrator.CORP> Export-VMHostImageDb -VMHost $vmhost -FilePath C:\vta\image.tgz
```

```
Mode                               LastWriteTime           Length Name
```

```

-----
-a----          10/8/2019  11:02 PM          2391 image.tgz

```

```

PS C:\Users\Administrator.CORP> Disconnect-VIServer -server * -Confirm:$false
PS C:\Users\Administrator.CORP> Connect-VIServer -server 192.168.110.22 -User trustedadmin@v-
sphere.local -Password 'VMware1!'

```

Name	Port	User
-----	----	----
192.168.110.22	443	VSPHERE.LOCAL\trustedadmin

```

PS C:\Users\Administrator.CORP> Export-TrustedPrincipal -FilePath C:\vta\principal.json

```

Mode	LastWriteTime	Length	Name
-----	-----	-----	----
-a----	10/8/2019 11:14 PM	1873	principal.json

Continue with [Import the Trusted Host Information to the Trust Authority Cluster](#).

Export and Import a TPM Endorsement Key Certificate

You can export a TPM endorsement key (EK) certificate from an ESXi host, and import it to the vSphere Trust Authority Cluster. You do so when you want to trust an individual ESXi host in the Trusted Cluster.

- [Enable the Trust Authority Administrator](#).
- [Enable the Trust Authority State](#).

To import a TPM EK certificate into the Trust Authority Cluster, you must change the Trust Authority Cluster's default attestation type to accept EK certificates. The default attestation type accepts TPM Certificate Authority (CA) certificates. Some TPMs do not include EK certificates. If you want to trust individual ESXi hosts, the TPM must include an EK certificate.

Note: Store the exported EK certificate files in a secure location, in case you must restore the vSphere Trust Authority configuration.

1. Ensure that you are connected as the Trust Authority administrator to the vCenter Server of the Trust Authority Cluster. For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trust Authority Cluster.

```

Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'

```

3. To change the Trust Authority Cluster's attestation type:
 - a) Run the `Get-TrustAuthorityCluster` cmdlet to show the clusters managed by this vCenter Server.

```

Get-TrustAuthorityCluster

```


The clusters are displayed.

- b) Assign the `Get-TrustAuthorityCluster` information to a variable.

For example, this command assigns the cluster named `vTA Cluster` to the variable `$vTA`.

```
$vTA = Get-TrustAuthorityCluster 'vTA Cluster'
```

- c) Assign the `Get-TrustAuthorityTpm2AttestationSettings` information to a variable.

For example, this command assigns the information to the variable `$tpm2Settings`.

```
$tpm2Settings = Get-TrustAuthorityTpm2AttestationSettings -TrustAuthorityCluster $vTA
```

- d) Run the `Set-TrustAuthorityTpm2AttestationSettings` cmdlet, specifying `RequireEndorsementKey`, or `RequireCertificateValidation`, or both.

For example, this command specifies `RequireEndorsementKey`.

```
Set-TrustAuthorityTpm2AttestationSettings -Tpm2AttestationSettings $tpm2Settings -RequireEndorsementKey
```

The system responds with a confirmation prompt similar to the following.

Confirmation

Configure the `Tpm2AttestationSettings` 'TrustAuthorityTpm2AttestationSettings-domain-c8' with the following parameters:

`RequireCertificateValidation: False`

`RequireEndorsementKey: True`

[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):

- e) At the confirmation prompt, press Enter. (The default is Y.)

The output shows a status of `True` for the setting specified. For example, this status shows `True` for `Require Endorsement Key`, and `False` for `Require Certificate Validation`.

Name	RequireEndorsementKey	RequireCertificateValidation
Health		
----	-----	-----

TrustAuthorityTpm2AttestationSettings...	True	False
Ok		

4. To export the TPM EK certificate:

- a) Disconnect from the vCenter Server of the Trust Authority Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
```

- b) Run the `Connect-VIServer` cmdlet to connect as the root user to one of the ESXi hosts in the Trusted Cluster.

```
Connect-VIServer -server host_ip_address -User root -Password 'password'
```

- c) Run the `Get-VMHost` cmdlet to confirm the ESXi host.

```
Get-VMHost
```

The host information is displayed.

- d) Assign `Get-VMHost` to a variable.

For example:

```
$vmhost = Get-VMHost
```

- e) Run the `Export-Tpm2EndorsementKey` cmdlet to export the EK certificate of the ESXi host.

For example, this command exports the EK certificate to the `tpm2ek.json` file.

```
Export-Tpm2EndorsementKey -VMHost $vmhost -FilePath C:\vta\tpm2ek.json
```

The file is created.

5. To import the TPM EK:

- a) Disconnect from the ESXi host in the Trusted Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
```

- b) Connect to the vCenter Server of the Trust Authority Cluster using the Trust Authority administrator user.

```
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

- c) Run the
- `Get-TrustAuthorityCluster`
- cmdlet.

```
Get-TrustAuthorityCluster
```

The clusters in the Trust Authority Cluster are displayed.

- d) Assign the
- `Get-TrustAuthorityCluster 'cluster'`
- information to a variable.

For example, this command assigns the information for cluster `vTA Cluster` to the variable `$vTA`.

```
$vTA = Get-TrustAuthorityCluster 'vTA Cluster'
```

- e) Run the
- `New-TrustAuthorityTpm2EndorsementKey`
- cmdlet.

For example, this command uses the `tpm2ek.json` file previously exported in Step 4.

```
New-TrustAuthorityTpm2EndorsementKey -TrustAuthorityCluster $vTA -FilePath C:\vta\tpm2ek.json
```

The imported endorsement key information is displayed.

The Trust Authority Cluster's attestation type is changed to accept EK certificates. The EK certificate is exported from the Trusted Cluster and imported to the Trust Authority Cluster.

Export and Import a TPM EK Certificate

This example shows how to use PowerCLI to change the Trust Authority Cluster's default attestation type to accept EK certificates, export the TPM EK certificate from the ESXi host in the Trusted Cluster, and import it to the Trust Authority Cluster. The following table shows the example components and values that are used.

Table 265: Example vSphere Trust Authority Setup

Component	Value
vCenter Server for Trust Authority Cluster	192.168.210.22
Variable <code>\$vTA</code>	<code>Get-TrustAuthorityCluster 'vTA Cluster'</code>
Variable <code>\$tpm2Settings</code>	<code>Get-TrustAuthorityTpm2AttestationSettings -TrustAuthorityCluster \$vTA</code>
Variable <code>\$vmhost</code>	<code>Get-VMHost</code>
ESXi host in Trusted Cluster	192.168.110.51
Trust Authority administrator	<code>trustedadmin@vsphere.local</code>
Local directory to contain output file	<code>C:\vta</code>

```
PS C:\Users\Administrator> Connect-VIServer -server 192.168.210.22 -User trustedadmin@vsphere.local -Password 'VMware1!'
```

```
Name                Port  User
----                -
192.168.210.22      443   VSPHERE.LOCAL\TrustedAdmin
```

```
PS C:\Users\Administrator> Get-TrustAuthorityCluster
```

```
Name                State      Id
```

```

----          -----          --
vTA Cluster      Enabled          TrustAuthorityCluster-domain-c8

PS C:\Users\Administrator> $vTA = Get-TrustAuthorityCluster 'vTA Cluster'

PS C:\Users\Administrator> $tpm2Settings = Get-TrustAuthorityTpm2AttestationSettings -TrustAuthori-
tyCluster $vTA

PS C:\Users\Administrator> Set-TrustAuthorityTpm2AttestationSettings -Tpm2AttestationSettings $tp-
m2Settings -RequireEndorsementKey

Confirmation
Configure the Tpm2AttestationSettings 'TrustAuthorityTpm2AttestationSettings-domain-c8' with the
following parameters:
RequireCertificateValidation: False
RequireEndorsementKey: True
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help (default is "Y"): y

Name                                     RequireEndorsementKey  RequireCertificateValida-
tion    Health
----
-----
TrustAuthorityTpm2AttestationSettings... True                  False
Ok

PS C:\Users\Administrator> Disconnect-VIServer -server * -Confirm:$false
PS C:\Users\Administrator> Connect-VIServer -server 192.168.110.51 -User root -Password 'VMware1!'

Name                                     Port  User
----
-----
192.168.110.51                          443   root

PS C:\Users\Administrator> Get-VMHost

Name                                     ConnectionState PowerState NumCpu CpuUsageMhz CpuTotalMhz  MemoryUsageGB
MemoryTotalGB Version
----
-----
192.168.110.51                          Connected      PoweredOn    4      55      9576      1.230
7.999  7.0.0

PS C:\Users\Administrator> $vmhost = Get-VMHost
PS C:\Users\Administrator> Export-Tpm2EndorsementKey -VMHost $vmhost -FilePath C:\vta\tpm2ek.json

Mode                                     LastWriteTime          Length Name
----
-----
-a----          12/3/2019 10:16 PM          2391 tpm2ek.json

PS C:\Users\Administrator> Disconnect-VIServer -server * -Confirm:$false
PS C:\Users\Administrator> Connect-VIServer -server 192.168.210.22 -User trustedadmin@vsphere.local
-Password 'VMware1!'

Name                                     Port  User

```

```

----
192.168.210.22          443   VSPHERE.LOCAL\TrustedAdmin

PS C:\Users\Administrator> Get-TrustAuthorityCluster

Name                State          Id
----                -
vTA Cluster         Enabled       TrustAuthorityCluster-domain-c8

PS C:\Users\Administrator> $vTA = Get-TrustAuthorityCluster 'vTA Cluster'
PS C:\Users\Administrator> New-TrustAuthorityTpm2EndorsementKey -TrustAuthorityCluster $vTA -
FilePath C:\vta\tpm2ek.json

TrustAuthorityClusterId          Name                Health
-----
TrustAuthorityCluster-domain-c8  1a520e42-4db8-1cbb-6dd7-f493fd921ccb  Ok

```

Continue with [Import the Trusted Host Information to the Trust Authority Cluster](#).

Import the Trusted Host Information to the Trust Authority Cluster

You import the exported ESXi host and vCenter Server information into the vSphere Trust Authority Cluster, so that the Trust Authority Cluster knows which hosts it can attest.

- [Enable the Trust Authority Administrator](#).
- [Enable the Trust Authority State](#).
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).

If you are following these tasks in order, you are still connected to the vCenter Server of the Trust Authority Cluster.

1. Ensure that you are connected as the Trust Authority administrator to the vCenter Server of the Trust Authority Cluster. For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trust Authority Cluster.

```

Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'

```

3. To show the clusters managed by this vCenter Server, run the `Get-TrustAuthorityCluster` cmdlet.

```
Get-TrustAuthorityCluster
```

The clusters are displayed.

4. Assign the `Get-TrustAuthorityCluster 'cluster'` information to a variable. For example, this command assigns the information for cluster `vTA Cluster` to the variable `$vTA`.


```
$vTA = Get-TrustAuthorityCluster 'vTA Cluster'
```
5. To import the vCenter Server principal information of the Trusted Cluster into the Trust Authority Cluster, run the `New-TrustAuthorityPrincipal` cmdlet. For example, the following command imports the `principal.json` file previously exported in [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).

```
New-TrustAuthorityPrincipal -TrustAuthorityCluster $vTA -FilePath C:\vta\principal.json
```

The `TrustAuthorityPrincipal` information is displayed.

6. To verify the import, run the `Get-TrustAuthorityPrincipal` cmdlet.

For example:

```
Get-TrustAuthorityPrincipal -TrustAuthorityCluster $vTA
```

The imported `TrustAuthorityPrincipal` information is displayed.

7. To import the Trusted Platform Module (TPM) CA certificate information, run the `New-TrustAuthorityTpm2CACertificate` cmdlet.

For example, the following command imports the TPM CA certificate information from the `cacert.zip` file previously exported in [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).

```
New-TrustAuthorityTpm2CACertificate -TrustAuthorityCluster $vTA -FilePath C:\vta\cacert.zip
```

The imported certificate information is displayed.

8. To import the ESXi host base image information, run the `New-TrustAuthorityVMHostBaseImage` cmdlet.

For example, the following command imports the image information from the `image.tgz` file previously exported in [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).

```
New-TrustAuthorityVMHostBaseImage -TrustAuthorityCluster $vTA -FilePath C:\vta\image.tgz
```

The imported image information is displayed.

The Trust Authority Cluster knows which ESXi hosts it can remotely attest, and so, which hosts it can trust.

Import the Trusted Host Information to the Trust Authority Cluster

This example shows how to use PowerCLI to import the vCenter Server principal information of the Trusted Cluster and the Trusted Host information files to the Trust Authority Cluster. It assumes that you are connected to the vCenter Server of the Trust Authority Cluster as the Trust Authority administrator. The following table shows the example components and values that are used.

Table 266: Example vSphere Trust Authority Setup

Component	Value
Variable <code>\$vTA</code>	<code>Get-TrustAuthorityCluster 'vTA Cluster1'</code>
vCenter Server for Trust Authority Cluster	192.168.210.22
Trust Authority Cluster names	vTA Cluster1 (Enabled) vTA Cluster2 (Disabled)
Principal information file	C:\vta\principal.json
TPM certificate file	C:\vta\cacert.cer
ESXi host base image file	C:\vta\image.tgz
Trust Authority administrator	trustedadmin@vsphere.local

```
PS C:\Users\Administrator> Disconnect-VIServer -server * -Confirm:$false
```

```
PS C:\Users\Administrator> Connect-VIServer -server 192.168.210.22 -User trustedadmin@vsphere.local  
-Password 'VMware1!'
```

```
Name                Port  User
----                -
192.168.210.22      443  VSPHERE.LOCAL\trustedadmin
```

```
PS C:\Users\Administrator> Get-TrustAuthorityCluster
```

```
Name                State      Id
```

```

----
vTA Cluster1      Enabled      TrustAuthorityCluster-domain-c8
vTA Cluster2      Disabled     TrustAuthorityCluster-domain-c26

```

```
PS C:\Users\Administrator> $vTA = Get-TrustAuthorityCluster 'vTA Cluster1'
```

```
PS C:\Users\Administrator.CORP> New-TrustAuthorityPrincipal -TrustAuthorityCluster $vTA -FilePath
C:\vta\principal.json
```

```

Name                                     Domain      Type      TrustAuthorityClusterId
----
vpxd-de207929-0601-43ef-9616-47d0cee0302f  vsphere.local STS_USER  TrustAuthorityCluster-do-
main-c8

```

```
PS C:\Users\Administrator.CORP> Get-TrustAuthorityPrincipal -TrustAuthorityCluster $vTA
```

```

Name                                     Domain      Type      TrustAuthorityClusterId
----
vpxd-de207929-0601-43ef-9616-47d0cee0302f  vsphere.local STS_USER  TrustAuthorityCluster-do-
main-c8

```

```
PS C:\Users\Administrator.CORP> New-TrustAuthorityTpm2CACertificate -TrustAuthorityCluster $vTA -
FilePath C:\vta\cacert.cer
```

```

TrustAuthorityClusterId                Name                                     Health
-----
TrustAuthorityCluster-domain-c8        52BDB7B4B2F55C925C047257DED4588A7767D961 Ok

```

```
PS C:\Users\Administrator.CORP> New-TrustAuthorityVMHostBaseImage -TrustAuthorityCluster $vTA -
FilePath C:\vta\image.tgz
```

```

TrustAuthorityClusterId                VMHostVersion                Health
-----
TrustAuthorityCluster-domain-c8        ESXi 7.0.0-0.0.14828939      Ok

```

Continue with [Create the Key Provider on the Trust Authority Cluster](#).

Create the Key Provider on the Trust Authority Cluster

For the Key Provider Service to connect to a key provider, you must create a trusted key provider then configure a trust setup between the vSphere Trust Authority Cluster and the key server (KMS). For most KMIP-compliant key servers, this configuration involves setting up client and server certificates.

- [Enable the Trust Authority Administrator](#).
- [Enable the Trust Authority State](#).
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).
- [Import the Trusted Host Information to the Trust Authority Cluster](#).
- Create and activate a key on the key server to be the primary key for the trusted key provider. This key wraps other keys and secrets used by this trusted key provider. See your key server vendor documentation for more information about creating keys.

What was previously called a KMS Cluster in vSphere 6.7 is now called a key provider in vSphere 7.0 and later. For more information about key providers, see [What Is the vSphere Trust Authority Key Provider Service](#).

In a production environment, you can create multiple key providers. By creating multiple key providers, you can address how to manage your deployment based on company organization, different business units or customers, and so on.

If you are following these tasks in order, you are still connected to the vCenter Server of the vSphere Trust Authority Cluster.

1. Ensure that you are connected to the vCenter Server of the Trust Authority Cluster. For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trust Authority Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

3. To create the trusted key provider, run the `New-TrustAuthorityKeyProvider` cmdlet. For example, this command uses 1 for the PrimaryKeyID and the name `clkp`. If you are following these tasks in order, you previously assigned `Get-TrustAuthorityCluster` information to a variable (for example, `$vTA = Get-TrustAuthorityCluster 'vTA Cluster'`).

```
New-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA -PrimaryKeyId 1 -Name clkp -KmpServerAddress ip_address
```

The PrimaryKeyID is normally a key ID that comes from the key server in the form of a UUID. Do not use the key name for PrimaryKeyID. The PrimaryKeyID value is vendor-dependent. See your key server documentation. The `New-TrustAuthorityKeyProvider` cmdlet can take other options, such as `KmpServerPort`, `ProxyAddress`, and `ProxyPort`. See the `New-TrustAuthorityKeyProvider` Help system for more information.

Each logical key provider, regardless of its type (Standard, Trusted, and Native Key Provider), must have a unique name across all vCenter Server systems.

For more information, see [Key Provider Naming](#).

Note: To add multiple key servers to the key provider, use the `Add-TrustAuthorityKeyProviderServer` cmdlet.

The key provider information is displayed.

4. Establish the trusted connection so that the key server trusts the trusted key provider. The exact process depends on the certificates that the key server accepts, and on your company policy. Select the option appropriate for your server and finish the steps.

Option	See
Upload Client Certificate	Upload the Client Certificate to Establish a Trusted Key Provider Trusted Connection.
Upload KMS certificate and private key	Upload the Certificate and Private Key to Establish a Trusted Key Provider Trusted Connection.
New Certificate Signing Request	Create a Certificate Signing Request to Establish a Trusted Key Provider Trusted Connection.

5. Finish the trust setup by uploading a key server certificate so that the trusted key provider trusts the key server.

- a) Assign the `Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA` information to a variable. For example:

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
```

This variable obtains the trusted key providers in the given Trust Authority Cluster, in this case, `$vTA`.

Note: If you have more than one trusted key provider, use commands similar to the following to select the one you want:

```
Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
<The trusted key providers listing is displayed.>
```

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA | Select-Object -Last 1
```

Using `Select-Object -Last 1` selects the last trusted key provider in the list.

- b) To get the key server server certificate, run the `Get-TrustAuthorityKeyProviderServerCertificate` command. For example:

```
Get-TrustAuthorityKeyProviderServerCertificate -KeyProviderServer $kp.KeyProviderServers
```

The server certificate information is displayed. Initially, the certificate is not trusted, so the `Trusted` state is `False`. If you have more than one key server configured, a list of certificates is returned. Verify and add each certificate using the following instructions.

- c) Before trusting the certificate, assign `Get-TrustAuthorityKeyProviderServerCertificate -KeyProviderServer $kp.KeyProviderServers` information to a variable (for example, `cert`), and run the `$cert.Certificate.ToString()` command and verify the output.

For example:

```
$cert = Get-TrustAuthorityKeyProviderServerCertificate -KeyProviderServer $kp.KeyProviderServers
$cert.Certificate.ToString()
```

The certificate information is displayed, including `Subject`, `Issuer`, and other information.

- d) To add the KMIP server certificate to the trusted key provider, run `Add-TrustAuthorityKeyProviderServerCertificate`.

For example:

```
Add-TrustAuthorityKeyProviderServerCertificate -ServerCertificate $cert
```

The certificate information is displayed and the `Trusted` state is now `True`.

6. Verify the status of the key provider.

- a) To refresh the key provider status, reassign the `$kp` variable.

For example:

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
```

Note: If you have more than one trusted key provider, use commands similar to the following to select the one you want:

```
Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
<The trusted key providers listing is displayed.>
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA | Select-Object -Last 1
```

Using `Select-Object -Last 1` selects the last trusted key provider in the list.

- b) Run the `$kp.Status` command to get the key provider status.

For example:

```
$kp.Status
```

Note: The status can take a few minutes to be refreshed. To view the status, reassign the `$kp` variable and rerun the `$kp.Status` command.

A Health status of `Ok` indicates that the key provider is running correctly.

The trusted key provider is created and has established trust with the key server.

Create the Key Provider on the Trust Authority Cluster

This example shows how to use PowerCLI to create the trusted key provider on the Trust Authority Cluster. It assumes that you are connected to the vCenter Server of the Trust Authority Cluster as the Trust Authority administrator. It also uses a certificate signed by the key server vendor after submitting a CSR to the vendor.

The following table shows the example components and values that are used.

Table 267: Example vSphere Trust Authority Setup

Component	Value
Variable \$vTA	Get-TrustAuthorityCluster 'vTA Cluster'
Variable \$kp	Get-TrustAuthorityKeyProvider -TrustAuthorityCluster \$vTA
Variable \$cert	Get-TrustAuthorityKeyProviderServerCertificate -KeyProviderServer \$kp.KeyProviderServers
vCenter Server for Trust Authority Cluster	192.168.210.22
KMIP-compliant key server	192.168.110.91
KMIP-compliant key server user	vcqekmip
Trust Authority Cluster name	vTA Cluster
Trust Authority administrator	trustedadmin@vsphere.local

```

PS C:\Users\Administrator.CORP> Disconnect-VIServer -server * -Confirm:$false
PS C:\Users\Administrator.CORP> Connect-VIServer -server 192.168.210.22 -User trustedadmin@v-
sphere.local -Password 'VMware!'

PS C:\Users\Administrator.CORP> New-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA -Prima-
ryKeyId 8 -Name clkp -KmipServerAddress 192.168.110.91
Name                PrimaryKeyId      Type          TrustAuthorityClusterId
----                -
clkp                8                KMIP          TrustAuthorityCluster-domain-c8

PS C:\Users\Administrator.CORP> New-TrustAuthorityKeyProviderClientCertificate -KeyProvider $kp
<Export the client certificate when you need to use it.>
PS C:\Users\Administrator.CORP> Export-TrustAuthorityKeyProviderClientCertificate -KeyProvider $kp
-FilePath clientcert.pem

PS C:\Users\Administrator.CORP> $kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
PS C:\Users\Administrator.CORP> Get-TrustAuthorityKeyProviderServerCertificate -KeyProviderServer
$kp.KeyProviderServers

Certificate                Trusted    KeyProviderServerId      KeyProviderId
-----                -
[Subject]...              False    domain-c8-clkp:192.16... domain-c8-clkp

PS C:\WINDOWS\system32> $cert.Certificate.ToString()
[Subject]
    E=<domain>, CN=<IP address>, OU=VMware Engineering, O=VMware, L=Palo Alto, S=California, C=US

[Issuer]
    O=<host>.eng.vmware.com, C=US, DC=local, DC=vsphere, CN=CA

[Serial Number]
    00CEF192BBF9D80C9F

[Not Before]

```

8/10/2015 4:16:12 PM

[Not After]

8/9/2020 4:16:12 PM

[Thumbprint]

C44068C124C057A3D07F51DCF18720E963604B70

```
PS C:\Users\Administrator.CORP> $cert = Get-TrustAuthorityKeyProviderServerCertificate -
KeyProviderServer $kp.KeyProviderServers
PS C:\Users\Administrator.CORP> Add-TrustAuthorityKeyProviderServerCertificate -ServerCertificate
$cert
```

Certificate	Trusted	KeyProviderServerId	KeyProviderId
-----	-----	-----	-----
[Subject]...	True		domain-c8-clkp

```
PS C:\Users\Administrator.CORP> $kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
PS C:\Users\Administrator.CORP> $kp.Status
```

KeyProviderId	Health	HealthDetails	ServerStatus
-----	-----	-----	-----
domain-c8-kp4	Ok	{}	{192.168.210.22}

Continue with [Export the Trust Authority Cluster Information](#).

Upload the Client Certificate to Establish a Trusted Key Provider Trusted Connection

Some key server (KMS) vendors require that you upload the trusted key provider's client certificate to the key server. After the upload, the key server accepts traffic that comes from the trusted key provider.

- [Enable the Trust Authority Administrator](#).
- [Enable the Trust Authority State](#).
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).
- [Import the Trusted Host Information to the Trust Authority Cluster](#).
- [Create the Key Provider on the Trust Authority Cluster](#).

1. Ensure that you are connected to the vCenter Server of the Trust Authority Cluster. For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trust Authority Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

3. Assign the `Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA` information to a variable. For example:

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
```

If you are following these tasks in order, you previously assigned `Get-TrustAuthorityCluster` information to a variable (for example, `$vTA = Get-TrustAuthorityCluster 'vTA Cluster'`).

This variable obtains the trusted key providers in the given Trust Authority Cluster, in this case, `$vTA`.

Note: If you have more than one trusted key provider, use commands similar to the following to select the one you want:

```
Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
<The trusted key providers listing is displayed.>
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA | Select-Object -Last 1
```

Using `Select-Object -Last 1` selects the last trusted key provider in the list.

4. To create the trusted key provider client certificate, run the `New-TrustAuthorityKeyProviderClientCertificate` cmdlet.

For example:

```
New-TrustAuthorityKeyProviderClientCertificate -KeyProvider $kp
```

The thumbprint is displayed.

5. To export the key provider client certificate, run the `Export-TrustAuthorityKeyProviderClientCertificate` cmdlet.

For example:

```
Export-TrustAuthorityKeyProviderClientCertificate -KeyProvider $kp -FilePath clientcert.pem
```

The certificate is exported to a file.

6. Upload the certificate file to the key server.

See your key server documentation for more information.

The trusted key provider has established trust with the key server.

Upload the Certificate and Private Key to Establish a Trusted Key Provider Trusted Connection

Some key server (KMS) vendors require that you configure the trusted key provider with the client certificate and private key provided by the key server. After you configure the trusted key provider, the key server accepts traffic from the trusted key provider.

- [Enable the Trust Authority Administrator.](#)
- [Enable the Trust Authority State.](#)
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted.](#)
- [Import the Trusted Host Information to the Trust Authority Cluster.](#)
- [Create the Key Provider on the Trust Authority Cluster.](#)
- Request a certificate and private key in PEM format from the key server vendor. If the certificate is returned in a format other than PEM, convert it to PEM. If the private key is protected with a password, create a PEM file with the password removed. You can use the `openssl` command for both operations. For example:
 - To convert a certificate from CRT to PEM format:


```
openssl x509 -in clientcert.crt -out clientcert.pem -outform PEM
```
 - To convert a certificate from DER to PEM format:


```
openssl x509 -inform DER -in clientcert.der -out clientcert.pem
```
 - To remove the password from a private key:


```
openssl rsa -in key.pem -out keynopassword.pem
```

```
Enter pass phrase for key.pem:
writing RSA key
```

1. Ensure that you are connected to the vCenter Server of the Trust Authority Cluster. For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trust Authority Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

3. Assign the `Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA` information to a variable. For example:

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
```

If you are following these tasks in order, you previously assigned `Get-TrustAuthorityCluster` information to a variable (for example, `$vTA = Get-TrustAuthorityCluster 'vTA Cluster'`).

The `$kp` variable obtains the trusted key providers in the given Trust Authority Cluster, in this case, `$vTA`.

Note: If you have more than one trusted key provider, use commands similar to the following to select the one you want:

```
Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
<The trusted key providers listing is displayed.>
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA | Select-Object -Last 1
```

Using `Select-Object -Last 1` selects the last trusted key provider in the list.

4. Upload the certificate and private key using the `Set-TrustAuthorityKeyProviderClientCertificate` command. For example:

```
Set-TrustAuthorityKeyProviderClientCertificate -KeyProvider $kp -CertificateFilePath <path/to/cert-file.pem> -PrivateKeyFilePath <path/to/privatekey.pem>
```

The trusted key provider has established trust with the key server.

Create a Certificate Signing Request to Establish a Trusted Key Provider Trusted Connection

Some key server (KMS) vendors require that you generate a Certificate Signing Request (CSR) and send that CSR to the key server vendor. The key server vendor signs the CSR and returns the signed certificate. After you configure this signed certificate as the trusted key provider's client certificate, the key server accepts traffic that comes from the trusted key provider.

- [Enable the Trust Authority Administrator.](#)
- [Enable the Trust Authority State.](#)
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted.](#)
- [Import the Trusted Host Information to the Trust Authority Cluster.](#)
- [Create the Key Provider on the Trust Authority Cluster.](#)

This task is a two-step process. First you generate the CSR and send it to your key server vendor. Then you upload the signed certificate that you receive from the key server vendor.

1. Ensure that you are connected to the vCenter Server of the Trust Authority Cluster. For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trust Authority Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
```

```
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

3. Assign the `Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA` information to a variable. For example:

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
```

If you are following these tasks in order, you previously assigned `Get-TrustAuthorityCluster` information to a variable (for example, `$vTA = Get-TrustAuthorityCluster 'vTA Cluster'`).

This variable obtains the trusted key providers in the given Trust Authority Cluster, in this case, `$vTA` .

Note: If you have more than one trusted key provider, use commands similar to the following to select the one you want:

```
Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
```

<The trusted key providers listing is displayed.>

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA | Select-Object -Last 1
```

Using `Select-Object -Last 1` selects the last trusted key provider in the list.

4. To generate a CSR, use the `New-TrustAuthorityKeyProviderClientCertificateCSR` cmdlet.

For example:

```
New-TrustAuthorityKeyProviderClientCertificateCSR -KeyProvider $kp
```

The CSR is displayed. You can also use the `Get-TrustAuthorityKeyProviderClientCertificateCSR -KeyProvider $kp` cmdlet to obtain the CSR.

5. To get a signed certificate, submit the CSR to your key server vendor.

The certificate must be in PEM format. If the certificate is returned in a format other than PEM, convert it to PEM by using the `openssl` command. For example:

- To convert a certificate from CRT to PEM format:

```
openssl x509 -in clientcert.crt -out clientcert.pem -outform PEM
```

- To convert a certificate from DER to PEM format:

```
openssl x509 -inform DER -in clientcert.der -out clientcert.pem
```

6. When you receive the signed certificate from the key server vendor, upload the certificate to the key server using the `Set-TrustAuthorityKeyProviderClientCertificate` cmdlet.

For example:

```
Set-TrustAuthorityKeyProviderClientCertificate -KeyProvider $kp -CertificateFilePath <path/tp/cert-file.pem>
```

The trusted key provider has established trust with the key server.

Export the Trust Authority Cluster Information

For the Trusted Cluster to connect to the vSphere Trust Authority Cluster, you must export the Trust Authority Cluster's service information in the form of a file then import that file to the Trusted Cluster. You must ensure to keep this file confidential and transport it securely.

- [Enable the Trust Authority Administrator.](#)
- [Enable the Trust Authority State.](#)
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted.](#)
- [Import the Trusted Host Information to the Trust Authority Cluster.](#)
- [Create the Key Provider on the Trust Authority Cluster.](#)

If you are following these tasks in order, you are still connected to the vCenter Server of the Trust Authority Cluster.

Note: Store the exported service information file in a secure location, in case you must restore the vSphere Trust Authority configuration.

1. Ensure that you are connected to the vCenter Server of the Trust Authority Cluster. For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trust Authority Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustAuthorityCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

3. To export the Trust Authority Cluster's Attestation Service and Key Provider Service information, run the `Export-TrustAuthorityServicesInfo` cmdlet.
For example, this command exports the service information to the `clsettings.json` file. If you are following these tasks in order, you previously assigned the `Get-TrustAuthorityCluster` information to a variable (for example, `$vTA = Get-TrustAuthorityCluster 'vTA Cluster'`).

```
Export-TrustAuthorityServicesInfo -TrustAuthorityCluster $vTA -FilePath C:\vta\clsettings.json
```

The file is created.

A file containing the Trust Authority Cluster information is created.

Export the Trust Authority Cluster Information

This example shows how to use PowerCLI to export the Trust Authority Cluster service information. The following table shows the example components and values that are used.

Table 268: Example vSphere Trust Authority Setup

Component	Value
Variable <code>\$vTA</code>	<code>Get-TrustAuthorityCluster 'vTA Cluster'</code>
vCenter Server for Trust Authority Cluster	192.168.210.22
Trust Authority administrator	trustedadmin@vsphere.local

```
PS C:\Users\Administrator.CORP> Disconnect-VIServer -server * -Confirm:$false
PS C:\Users\Administrator.CORP> Connect-VIServer -server 192.168.210.22 -User trustedadmin@vsphere.local -Password 'VMware1!'
```

```
PS C:\Users\Administrator.CORP> Export-TrustAuthorityServicesInfo -TrustAuthorityCluster $vTA -
FilePath C:\vta\clsettings.json
```

```
Mode                LastWriteTime         Length Name
----                -
-a----           10/16/2019   9:59 PM             8177 clsettings.json
```

Continue with [Import the Trust Authority Cluster Information to the Trusted Hosts](#).

Import the Trust Authority Cluster Information to the Trusted Hosts

After you have imported the vSphere Trust Authority Cluster information to the Trusted Cluster, the Trusted Hosts start the attestation process with the Trust Authority Cluster.

- [Enable the Trust Authority Administrator](#).
- [Enable the Trust Authority State](#).
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#).

- [Import the Trusted Host Information to the Trust Authority Cluster.](#)
- [Create the Key Provider on the Trust Authority Cluster.](#)
- [Export the Trust Authority Cluster Information.](#)

1. Ensure that you are connected as the Trust Authority administrator to the vCenter Server of the Trusted Cluster.
For example, you can enter `$global:defaultviservers` to show all the connected servers.
2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trusted Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustedCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

Note: Alternatively, you can start another PowerCLI session to connect to the vCenter Server of the Trusted Cluster.

3. Verify that the state of the Trusted Cluster is disabled.

```
Get-TrustedCluster
```

The State is shown as Disabled.

4. Assign the `Get-TrustedCluster` information to a variable.

For example, this command assigns information for the cluster `Trusted Cluster` to the variable `$TC`.

```
$TC = Get-TrustedCluster -Name 'Trusted Cluster'
```

5. Verify the value of the variable by echoing it.

For example:

```
$TC
```

The `Get-TrustedCluster` information is displayed.

6. To import the Trust Authority Cluster information to the vCenter Server, run the `Import-TrustAuthorityServicesInfo` cmdlet.

For example, this command imports the service information from the `clsettings.json` file previously exported in [Export the Trust Authority Cluster Information](#).

```
Import-TrustAuthorityServicesInfo -FilePath C:\vta\clsettings.json
```

The system responds with a confirmation prompt.

```
Confirmation
Importing the TrustAuthorityServicesInfo into Server 'ip_address'. Do you want to proceed?

[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help (default is "Y"):
```

7. At the confirmation prompt, press Enter. (The default is `Y`.)

The service information for the hosts in the Trust Authority Cluster is displayed.

8. To enable the Trusted Cluster, run the `Set-TrustedCluster` cmdlet.

For example:

```
Set-TrustedCluster -TrustedCluster $TC -State Enabled
```

The system responds with a confirmation prompt.

```
Confirmation
Setting TrustedCluster 'cluster' with new TrustedState 'Enabled'. Do you want to proceed?
[Y] Yes  [A] Yes to All  [N] No  [L] No to All  [S] Suspend  [?] Help (default is "Y"):
```

If the Trusted Cluster is not in a healthy state, the following warning message is displayed before the confirmation message:

```
WARNING: The TrustedCluster 'cluster' is not healthy in its TrustedClusterAppliedStatus. This cmdlet will
automatically remediate the TrustedCluster.
```

9. At the confirmation prompt, press Enter. (The default is Y.)

The Trusted Cluster is enabled.

Note: You can also enable the Trusted Cluster by enabling the Attestation Service and the Key Provider Service individually. Use the `Add-TrustedClusterAttestationServiceInfo` and `Add-TrustedClusterKeyProviderServiceInfo` commands. For example, the following commands enable the services one at a time for the cluster `Trusted Cluster` that has two Key Provider Services and two Attestation Services.

```
Add-TrustedClusterAttestationServiceInfo -TrustedCluster 'Trusted Cluster' -AttestationServiceInfo (Get-AttestationServiceInfo | Select-Object -index 0,1)
Add-TrustedClusterKeyProviderServiceInfo -TrustedCluster 'Trusted Cluster' -KeyProviderServiceInfo (Get-KeyProviderServiceInfo | Select-Object -index 0,1)
```

10. Verify that the Attestation Service and the Key Provider Service are configured in the Trusted Cluster.

- a) Assign the `Get-TrustedCluster` information to a variable.

For example, this command assigns information for the cluster `Trusted Cluster` to the variable `$TC`.

```
$TC = Get-TrustedCluster -Name 'Trusted Cluster'
```

- b) Verify that the Attestation Service is configured.

```
$tc.AttestationServiceInfo
```

The Attestation Service information is displayed.

- c) Verify that the Key Provider Service is configured.

```
$tc.KeyProviderServiceInfo
```

The Key Provider Service information is displayed.

The ESXi Trusted Hosts in the Trusted Cluster begin the attestation process with the Trust Authority Cluster.

Import the Trust Authority Cluster Information to the Trusted Hosts

This example shows how to import the Trust Authority Cluster service information to the Trusted Cluster. The following table shows the example components and values that are used.

Table 269: Example vSphere Trust Authority Setup

Component	Value
vCenter Server of the Trusted Cluster	192.168.110.22
Trust Authority administrator	trustedadmin@vsphere.local
Trusted Cluster name	Trusted Cluster
ESXi hosts in the Trust Authority Cluster	192.168.210.51 and 192.168.210.52
Variable <code>\$TC</code>	<code>Get-TrustedCluster -Name 'Trusted Cluster'</code>

```
PS C:\Users\Administrator.CORP> Disconnect-VIServer -server * -Confirm:$false
PS C:\Users\Administrator.CORP> Connect-VIServer -server 192.168.110.22 -User trustedadmin@vsphere.local -Password 'VMware!'
```

```
Name          Port  User
----          -
192.168.110.22 443   VSPHERE.LOCAL\trustedadmin
```

```
PS C:\Users\Administrator.CORP> Get-TrustedCluster
```


Name	State	Id
-----	-----	--
Trusted Cluster	Disabled	TrustedCluster-domain-c8

```
PS C:\Users\Administrator.CORP> $TC = Get-TrustedCluster -Name 'Trusted Cluster'
PS C:\Users\Administrator.CORP> $TC
```

Name	State	Id
-----	-----	--
Trusted Cluster	Disabled	TrustedCluster-domain-c8

```
PS C:\Users\Administrator.CORP> Import-TrustAuthorityServicesInfo -FilePath C:\vta\clsettings.json
```

Confirmation

Importing the TrustAuthorityServicesInfo into Server '192.168.110.22'. Do you want to proceed?

[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"): y

ServiceAddress	ServicePort	ServiceGroup
-----	-----	-----
192.168.210.51	443	host-13:86f7ab6c-ad6f-4606-...
192.168.210.52	443	host-16:86f7ab6c-ad6f-4606-...
192.168.210.51	443	host-13:86f7ab6c-ad6f-4606-...
192.168.210.52	443	host-16:86f7ab6c-ad6f-4606-...

```
PS C:\Users\Administrator.CORP> Set-TrustedCluster -TrustedCluster $TC -State Enabled
```

Confirmation

Setting TrustedCluster 'Trusted Cluster' with new TrustedState 'Enabled'. Do you want to proceed?

[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):

Name	State	Id
-----	-----	--
Trusted Cluster	Enabled	TrustedCluster-domain-c8

```
PS C:\Users\Administrator.CORP> $TC = Get-TrustedCluster -Name 'Trusted Cluster'
PS C:\Users\Administrator.CORP> $tc.AttestationServiceInfo
```

ServiceAddress	ServicePort	ServiceGroup
-----	-----	-----
192.168.210.51	443	host-13:dc825986-73d2-463c-...
192.168.210.52	443	host-16:dc825986-73d2-463c-...

```
PS C:\Users\Administrator.CORP> $tc.KeyProviderServiceInfo
```

ServiceAddress	ServicePort	ServiceGroup
-----	-----	-----
192.168.210.51	443	host-13:dc825986-73d2-463c-...
192.168.210.52	443	host-16:dc825986-73d2-463c-...

Continue with [Configure the Trusted Key Provider for Trusted Hosts Using the vSphere Client](#) or [Configure the Trusted Key Provider for Trusted Hosts Using the Command Line](#).

Configure the Trusted Key Provider for Trusted Hosts Using the vSphere Client

You can configure the trusted key provider by using the vSphere Client.

- [Enable the Trust Authority Administrator.](#)
 - [Enable the Trust Authority State.](#)
 - [Collect Information About ESXi Hosts and vCenter Server to Be Trusted.](#)
 - [Import the Trusted Host Information to the Trust Authority Cluster.](#)
 - [Create the Key Provider on the Trust Authority Cluster.](#)
 - [Export the Trust Authority Cluster Information.](#)
 - [Import the Trust Authority Cluster Information to the Trusted Hosts.](#)
1. Connect to vCenter Server of the Trusted Cluster by using the vSphere Client.
 2. Log in as the vCenter Server administrator, or an administrator that has the **Cryptographic operations > Manage key servers** privilege.
 3. Select the vCenter Server, then select **Configure**.
 4. Select **Key Providers** under **Security**.
 5. Select **Add Trusted Key Providers**.
The trusted key providers that are available are shown with a status of Connected.
 6. Select a trusted key provider and click **Add Key Providers**.
The trusted key provider is shown as Trusted and Connected. If this is the first trusted key provider that you add, it is marked as the default.

Note: It takes a while for all the hosts to be able to get the key provider, and for the vCenter Server to update its cache. Because of the way the information is propagated, you might have to wait for a few minutes to use the key provider for key operations on some of the hosts.

ESXi Trusted Hosts can now perform cryptographic operations, such as creating encrypted virtual machines.

Encrypting a virtual machine with a trusted key provider looks the same as the virtual machine encryption user experience that was first delivered in vSphere 6.5. See [Using Encryption in Your vSphere Environment](#).

Configure the Trusted Key Provider for Trusted Hosts Using the Command Line

You can configure trusted key providers by using the command line. You can configure the default trusted key provider for the vCenter Server, or at the cluster or the cluster folder level in the vCenter object hierarchy.

- [Enable the Trust Authority Administrator.](#)
- [Enable the Trust Authority State.](#)
- [Collect Information About ESXi Hosts and vCenter Server to Be Trusted.](#)
- [Import the Trusted Host Information to the Trust Authority Cluster.](#)
- [Create the Key Provider on the Trust Authority Cluster.](#)
- [Export the Trust Authority Cluster Information.](#)
- [Import the Trust Authority Cluster Information to the Trusted Hosts.](#)

On the Trusted Cluster, you must have a role that includes the **Cryptographic operations > Manage KMS** privilege.

1. Ensure that you are connected as an administrator to the vCenter Server of the Trusted Cluster.
For example, you can enter `$global:defaultviservers` to show all the connected servers.

2. Optional: If necessary, you can run the following commands to ensure that you are connected to the vCenter Server of the Trusted Cluster.

```
Disconnect-VIServer -server * -Confirm:$false
Connect-VIServer -server TrustedCluster_VC_ip_address -User admin_user -Password 'password'
```

3. Obtain the trusted key provider.

```
Get-KeyProvider
```

You can use the `-Name` *keyprovider* option to specify a single trusted key provider.

4. Assign the `Get-KeyProvider` trusted key provider information to a variable.
For example, this command assigns the information to the variable `$workload_kp`.

```
$workload_kp = Get-KeyProvider
```

If you have multiple trusted key providers, you can use `Select-Object` to select one of them.

```
$workload_kp = Get-KeyProvider | Select-Object -Index 0
```

5. Register the trusted key provider.

```
Register-KeyProvider -KeyProvider $workload_kp
```

To register additional trusted key providers, repeat Step 4 and Step 5.

Note: It takes a while for all the hosts to be able to get the key provider, and for the vCenter Server to update its cache. Because of the way the information is propagated, you might have to wait for a few minutes to use the key provider for key operations on some of the hosts.

6. Set the default trusted key provider to use.

- a) To set the default key provider at the vCenter Server level, run the following command.

```
Set-KeyProvider -KeyProvider $workload_kp -DefaultForSystem
```

- b) To set the key provider at the cluster level, run the following command.

For example, this command sets the key provider for the cluster `Trusted Cluster`.

```
Add-EntityDefaultKeyProvider -KeyProvider $workload_kp -Entity 'Trusted Cluster'
```

- c) To set the key provider at the cluster folder level, run the following command.

For example, this command sets the key provider for the cluster folder `TC Folder`, which was created on the `workLoad data center`.

```
Add-EntityDefaultKeyProvider -KeyProvider $workload_kp -Entity 'TC Folder'
```

Encrypting a virtual machine with a trusted key provider looks the same as the virtual machine encryption user experience that was first delivered in vSphere 6.5. See [Using Encryption in Your vSphere Environment](#).

Managing vSphere Trust Authority in Your vSphere Environment

After you configure vSphere Trust Authority, you can perform additional operations, such as stopping and starting services, adding hosts to clusters, and viewing the status of the Trust Authority Cluster.

You can perform tasks by using the vSphere Client, the API, and the PowerCLI cmdlets. See *vSphere Web Services SDK Programming Guide*, the *VMware PowerCLI* documentation, and the *VMware PowerCLI Cmdlets Reference* documentation.

Start, Stop, and Restart vSphere Trust Authority Services

You can start, stop, and restart vSphere Trust Authority services by using the vSphere Client.

The services that make up vSphere Trust Authority are the Attestation Service (attestd) and the Key Provider Service (kmtx).

1. Connect to the vCenter Server of the vSphere Trust Authority Cluster by using the vSphere Client.
2. Log in as an administrator.
3. Browse to an ESXi host in the Trust Authority Cluster.
4. Select **Configure**, then select **Services** under **System**.
5. Locate the attestd service and the kmtx service.
6. Select the **Restart**, **Start**, or **Stop** operation as appropriate.

View the Trust Authority Hosts

You can view the vSphere Trust Authority hosts configured for a Trusted Cluster by using the vSphere Client.

1. Connect to the vCenter Server of the Trusted Cluster by using the vSphere Client.
2. Log in as an administrator.
3. Select the vCenter Server instance.
4. Click the **Configure** tab and select **Trust Authority** under **Security**.

The ESXi hosts in the Trust Authority Cluster configured for the Trusted Cluster are displayed.

View the vSphere Trust Authority Cluster State

You can view the state of the vSphere Trust Authority Cluster by using the vSphere Client. The state is either enabled or disabled.

When the Trust Authority Cluster state is enabled, the Trusted Hosts in the Trusted Cluster can communicate with the Attestation Service and the Key Provider Service.

1. Connect to the vCenter Server of the Trust Authority Cluster by using the vSphere Client.
2. Log in as an administrator.
3. Select the Trust Authority Cluster in the object hierarchy.
4. Click the **Configure** tab and select **Trust Authority Cluster** under **Trust Authority**.

The status displays as Enabled or Disabled.

Restart the Trusted Host Service

You can restart the service that runs on your Trusted Hosts.

Access to the ESXi shell must be enabled. See [Activate Access to the ESXi Shell Using the vSphere Client](#).

The service, kmtx, runs on the ESXi Trusted Hosts.

1. Use SSH or another remote console connection to start a session on the ESXi Trusted Host.
2. Log in as root.
3. Run the following command.

```
/etc/init.d/kmtx restart
```

Adding and Removing vSphere Trust Authority Hosts

You add and remove ESXi hosts to a vSphere Trust Authority Cluster by using VMware-supplied scripts.

In vSphere 7.0, you add and remove ESXi hosts to and from an existing vSphere Trust Authority Cluster or a Trusted Cluster by using VMware-supplied scripts. In vSphere 7.0 Update 1 and later, you use the remediate function to add ESXi hosts to an existing Trusted Cluster. See [Add a Host to a Trusted Cluster Using the vSphere Client](#) and [Add a Host to a Trusted Cluster Using the Command Line](#).

In vSphere 7.0 Update 1 and later, you still must use scripts to add ESXi hosts to an existing Trust Authority Cluster. See the VMware knowledge base articles at <https://kb.vmware.com/s/article/77234> and <https://kb.vmware.com/s/article/77146>.

Add a Host to a Trusted Cluster Using the vSphere Client

You can add ESXi hosts to an existing Trusted Cluster using the vSphere Client.

The vCenter Server for the Trusted Cluster must be running vSphere 7.0 Update 1 or later.

If you are adding an ESXi host that has a different ESXi version, or a different TPM hardware type, than what you initially configured for the Trusted Cluster, additional steps are required. You must export and import this information to the vSphere Trust Authority Cluster. See [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#) and [Import the Trusted Host Information to the Trust Authority Cluster](#).

Required privileges: See the add hosts tasks in [Required vCenter Server Privileges for Common Tasks](#).

After you have initially configured a Trusted Cluster, you might want to add more ESXi hosts. However, when you add the host to a Trusted Cluster, you must take the additional step of remediation. When you remediate the Trusted Cluster, you ensure that its desired configuration state matches its applied configuration.

In the first version of vSphere Trust Authority released in vSphere 7.0, you run scripts to add a host to an existing Trusted Cluster. In vSphere 7.0 Update 1 and later, you use the remediate functionality to add a host to a Trusted Cluster. In vSphere 7.0 Update 1 and later, you still must use scripts to add a host to an existing Trust Authority Cluster. See [Adding and Removing vSphere Trust Authority Hosts](#).

1. Connect to the vCenter Server of the Trusted Cluster by using the vSphere Client.
2. Log in as a Trust Authority administrator.
3. Navigate to a Trusted Cluster.
4. On the **Configure** tab, select **Configuration > Quickstart**.
5. Click **Add** in the **Add hosts** card.
6. Follow the prompts.
7. On the **Trust Authority** tab, click **Remediate**.
8. To verify that the Trusted Cluster is healthy, click **Check Health**.

Add a Host to a Trusted Cluster Using the Command Line

You can add ESXi hosts to an existing Trusted Cluster using the command line.

- The vCenter Server for the Trusted Cluster must be running vSphere 7.0 Update 1 or later.
- PowerCLI 12.1.0 or later is required.
- Required privileges: See the add hosts tasks in [Required vCenter Server Privileges for Common Tasks](#).

After you have initially configured a Trusted Cluster, you might want to add more ESXi hosts. However, when you add the host to a Trusted Cluster, you must take the additional step of remediation. When you remediate the Trusted Cluster, you ensure that its desired configuration state matches its applied configuration.

In the first version of vSphere Trust Authority released in vSphere 7.0, you run scripts to add a host to an existing Trusted Cluster. In vSphere 7.0 Update 1 and later, you use the remediate functionality to add a Trusted Host. In vSphere 7.0 Update 1 and later, you still must use scripts to add a host to an existing Trust Authority Cluster. See [Adding and Removing vSphere Trust Authority Hosts](#).

1. Use whatever steps you normally do to add the ESXi host to the Trusted Cluster.
2. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as the Trust Authority administrator to the vCenter Server of the Trusted Cluster.

```
Connect-VIServer -server TrustedCluster_VC_ip_address -User trust_admin_user -Password 'password'
```

3. To check the Trusted Cluster's status, run the `Get-TrustedClusterAppliedStatus` PowerCLI cmdlet.

```
Get-TrustedClusterAppliedStatus -TrustedCluster 'TrustedCluster'
```

4. If the Trusted Cluster is not healthy, run the `Set-TrustedCluster` cmdlet with the `-Remediate` parameter.

```
Set-TrustedCluster -TrustedCluster 'TrustedCluster' -Remediate
```

5. To verify that the Trusted Cluster is healthy, rerun the `Get-TrustedClusterAppliedStatus` cmdlet.

```
Get-TrustedClusterAppliedStatus -TrustedCluster 'TrustedCluster'
```

Decommission Trusted Hosts from a Trusted Cluster

You can remove, or decommission, Trusted Hosts from a Trusted Cluster. You can decommission one or all Trusted Hosts from a Trusted Cluster, depending on the scenario.

- The vCenter Server for the Trusted Cluster must be running vSphere 7.0 Update 1 or later.
- If you use PowerCLI, version 12.1.0 or later is required.

When you decommission a Trusted Host, the remediate function sets the desired state of the Trusted Host to that of the non-Trusted Cluster where it is moved to. The decommissioned Trusted Host becomes a regular host. The Trusted Cluster (from where the Trusted Host was moved) continues to have its desired state configuration and functions still as a Trusted Cluster.

When you remove all the Trusted Hosts from a Trusted Cluster, you decommission the Trusted Cluster. You remove both the desired state configuration and applied configuration from the Trusted Hosts and the Trusted Cluster, then move all the Trusted Hosts to a non-Trusted Cluster.

You can reuse decommissioned Trusted Hosts in your environment. For example you can reuse the hosts in a non-trusted infrastructure capacity, or as vSphere Trust Authority Hosts. You can use the decommissioned hosts in the same vCenter Server or a different vCenter Server.

For more information about Trusted Cluster configuration and health, see [Checking and Remediating Trusted Cluster Health](#).

1. Connect to the vCenter Server of the Trusted Cluster by using the vSphere Client.
2. Log in as a Trust Authority administrator.
3. Navigate to a Trusted Cluster.
4. Decide how to decommission the Trusted Hosts from the Trusted Cluster.

Task	Steps
Keep the desired configuration state of the Trusted Cluster and the remaining Trusted Hosts	1. Put hosts into Maintenance mode and move them to a new, empty cluster (that is, the cluster does not contain any hosts). 2. Exit Maintenance mode on the hosts.

Task	Steps
	3. For the new, empty cluster (not the Trusted Cluster), on the Trust Authority tab, click Remediate. Remediation removes the Trusted configuration from the moved hosts. The Trusted Cluster retains its desired state configuration.
Remove the desired configuration state and applied configuration state of all the Trusted Hosts	1. In a PowerCLI session, run the <code>Connect-VIServer</code> cmdlet to connect as the Trust Authority administrator to the vCenter Server of the Trusted Cluster. <pre>Connect-VIServer -server TrustedCluster_VC_ip_address -User trust_admin_user -Password 'password'</pre> 2. Run the <code>Set-TrustedCluster</code> cmdlet, for example: <pre>Set-TrustedCluster -TrustedCluster 'TrustedCluster' -State Disabled</pre> The Trusted Infrastructure configuration is removed from all the Trusted Hosts, and the Trusted Cluster has its desired state configuration removed. 3. Put all hosts into Maintenance mode and move them to a different cluster. 4. Exit Maintenance mode on the hosts.

5. To verify that the Trusted Cluster is healthy, click **Check Health** on the **Trust Authority** tab for the Trusted Cluster.

If you no longer plan to attest the specific versions of ESXi or the TPM hardware from the decommissioned ESXi hosts, update the Trust Authority Cluster's configuration, for optimal security. See the VMware knowledge base article at <https://kb.vmware.com/s/article/77146>.

Backing Up the vSphere Trust Authority Configuration

Use the files you exported when configuring vSphere Trust Authority as your Trust Authority backup. You can use these files to restore a Trust Authority deployment. Keep these configuration files confidential and transport them securely.

Most vSphere Trust Authority configuration and state information is stored on the ESXi hosts in the ConfigStore database. The vCenter Server Management Interface that you use to back up a vCenter Server instance does not back up the configuration information for vSphere Trust Authority. If you save and securely store the configuration files that you exported when setting up your vSphere Trust Authority environment, then you have the necessary information to restore a vSphere Trust Authority configuration. See [Collect Information About ESXi Hosts and vCenter Server to Be Trusted](#) if you must generate this information.

Change the Primary Key of a Trusted Key Provider

You can change the primary key of a trusted key provider, for example, when you want to rotate the primary key that is used.

Create and activate a key on the key server (KMS) to be used as the new primary key for the trusted key provider. This key wraps other keys and secrets used by this trusted key provider. See your KMS vendor documentation for more information about creating keys.

See [Virtual Machine Encryption Best Practices](#) for guidance about key life cycle.

1. Run the `Set-TrustAuthorityKeyProvider` command.

For example:

```
Set-TrustAuthorityKeyProvider -MasterKeyId Key-ID
```

2. Verify the status of the key provider.

a) Assign `Get-TrustAuthorityCluster` information to a variable.

For example:

```
$vTA = Get-TrustAuthorityCluster 'vTA Cluster'
```

b) Assign the `Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA` information to a variable.

For example:

```
$kp = Get-TrustAuthorityKeyProvider -TrustAuthorityCluster $vTA
```

c) Verify the status of the key provider by running `$kp.Status`.

For example:

```
$kp.Status
```

```
KeyProviderId Health HealthDetails ServerStatus
-----
domain-c8-kp4    Ok {}                {IP_address}
```

A Health status of Ok indicates that the key provider is running correctly.

The new primary key is used for any new encryption operations. Data encrypted with the old primary key is still decrypted using the old key.

Trusted Host Attestation Reporting

In vSphere Trust Authority, vCenter Server verifies and reports on a Trusted Host's attestation status. You can use the vSphere Client to view the attestation status of Trusted Hosts.

What Is vSphere Trust Authority Attestation Reporting

vSphere Trust Authority uses remote attestation for Trusted Hosts to prove the authenticity of their booted software. Attestation verifies that the Trusted Hosts are running authentic VMware software, or VMware-signed partner software. The vCenter Server of the Trusted Cluster communicates with the Trusted Host to get an internal attestation report. The attestation report specifies if the Trusted Host has attested or not with the Attestation Service running on the Trust Authority Cluster. If the Trusted Host has not attested, the attestation report also specifies an error message. The vSphere Client displays the attestation status of a Trusted Host, and if vSphere Trust Authority or vCenter Server attested the host.

Passed Attestation Status

A status of Passed indicates that the Trusted Host has attested with a vSphere Trust Authority Attestation Service, and the internal attestation report is available to vCenter Server.

Failed Attestation Status

A status of Failed indicates that the Trusted Host was not able to attest with any vSphere Trust Authority Attestation Service. The vCenter Server internal attestation report contains the error reported by the Attestation Service that the Trusted Host tried to attest with.

Handling Unattested Trusted Hosts

When a Trusted Host is unattested, virtual machines, including encrypted virtual machines, that are running on the Trusted Host continue to be accessible. You cannot power on virtual machines on an unattested Trusted Host. However, you can still add unencrypted virtual machines. When a Trusted Host is unattested, take steps to resolve the attestation problem. See [Troubleshoot Trusted Host Attestation Problems](#).

Multiple Trust Authority Hosts and Attestation Reports

When you have configured multiple Trust Authority Hosts, there are potentially multiple attestation reports available from each host. When reporting status, the vSphere Client displays the status from the first "attested" report that it finds. If there are no "attested" reports, the vSphere Client displays the error from the first "unattested" report that it finds.

Even if you have configured multiple Trust Authority Hosts, the vSphere Client displays the status, and potentially an error message, from only one attestation report.

View the Trusted Cluster Attestation Status

You can view the attestation status of a Trusted Host by using the vSphere Client.

- Both the Trusted Hosts and the vSphere Trust Authority hosts must be running ESXi 7.0 Update 1 or later.
- The vCenter Server hosts for the respective clusters must be running vSphere 7.0 Update 1 or later.

1. Connect to the vCenter Server of the Trusted Cluster by using the vSphere Client.

2. Log in as an administrator.

You can log in as a Trust Authority administrator or vSphere administrator.

3. Navigate to a data center and click the **Monitor** tab.

4. Click **Security**.

5. Review the Trusted Host's status in the Attestation column and read the accompanying message in the Message column.

If there are errors, see [Troubleshoot Trusted Host Attestation Problems](#).

Troubleshoot Trusted Host Attestation Problems

The vSphere Trust Authority attestation reporting provides a starting point for troubleshooting Trusted Host attestation errors.

1. [View the Trusted Cluster Attestation Status](#).

2. Use the following table to troubleshoot and resolve errors.

Error	Cause and Solution
Attestation Services not configured.	Attestation Services have not been configured. Configure the Trusted Host to use attestation services by using the Remediate action. See Remediate a Trusted Cluster .
No TPM2 device available.	Install and configure the Trusted Host to use a Trusted Platform Module (TPM). See your vendor documentation.
TPM2 endorsement public key or certificate could not be retrieved.	Check that the TPM is supported, and that it has a valid endorsement key. You might need to contact VMware Support.
Attestation report is not available.	It is possible that the Trusted Host has not finished attestation. Wait a few minutes then recheck the attestation status.
Attestation Service version is incompatible with the request.	Update the Trust Authority host running the Attestation Service to vSphere 7.0 Update 1 or later.
Attestation failed because Secure Boot is not enabled.	Check that the Trusted Host is configured to use Secure Boot. See UEFI Secure Boot for ESXi Hosts .
Attestation failed to identify the remote software version.	Import the Trusted Host's base image information to the Attestation Service. See Import the Trusted Host Information to the Trust Authority Cluster .
Attestation failed because a TPM certificate is required.	Check that the TPM is supported. Alternatively, run the following PowerCLI cmdlet to modify the

Error	Cause and Solution
	com.vmware.esx.attestation.tpm2.settings to set <code>requireCertificateValidation</code> to <code>false</code>. <pre>Set-TrustAuthorityTpm2AttestationSettings -TrustAuthorityCluster TrustedCluster -RequireCertificateValidation:\$false -RequireEndorsementKey:\$true</pre>
Attestation failed due to an unknown TPM.	Import the TPM endorsement key to the Attestation Services. See Import the Trusted Host Information to the Trust Authority Cluster .
Error: vapi.send.failed.	The kmxa service might not be running on the Trusted Host or the kmxa service cannot contact the Attestation Service. Ensure that the kmxa service is started. Also, check that the Attestation Service is running. See Restart the Trusted Host Service .

Checking and Remediating Trusted Cluster Health

You can check and validate the health of a Trusted Cluster. If a Trusted Cluster's configuration is not healthy, you must resolve the configuration inconsistencies. You do so by remediating the Trusted Cluster. When you remediate a Trusted Cluster, you ensure that all the Trusted Hosts in the Trusted Cluster have the same trusted configuration.

A Trusted Cluster consists of a vCenter Server cluster of Trusted ESXi Hosts that are remotely attested by the Trust Authority Cluster. When you configure vSphere Trust Authority initially, you must import the Trust Authority Services information from your Trust Authority Cluster into the Trusted Cluster. The Trusted Cluster uses that configuration of components for contacting the Key Provider Service and the Attestation Service running on the Trust Authority Cluster. For more information about this aspect of configuring a Trusted Cluster, see [Import the Trust Authority Cluster Information to the Trusted Hosts](#). After you configure a Trusted Cluster, you can check and remediate its health.

Checking the Trusted Cluster Health

Checking the health of a Trusted Cluster depends upon the following.

Desired state configuration

The desired state configuration is based on the Trust Authority Services information that you import into the Trusted Cluster. The desired state configuration is the Trusted Cluster's "source of truth." Think of the desired state configuration as what is initially created when you set up the Trusted Cluster.

Applied configuration

The applied configuration is the registration of the specific Attestation Services and Key Provider Services for which you have configured the Trusted Cluster. The applied configuration is what the Trusted Cluster is running currently. You can think of the applied configuration as the "run-time" configuration. The desired state configuration should match the applied configuration. However, if the applied configuration is inconsistent with the desired state configuration, the Trusted Cluster is deemed "not healthy." A Trusted Cluster that is not healthy can experience degraded performance or not function at all.

This health check is not an indicator of the overall health for either a Trusted Cluster or the vSphere Trust Authority infrastructure. The health check only compares the Trusted Cluster's desired state configuration to the applied configuration.

Remediating the Trusted Cluster

Remediation is the process by which vSphere Trust Authority resolves an inconsistent configuration of a Trusted Cluster. A Trusted Cluster's configuration can become inconsistent over time or due to other operational errors.

Use remediation in the following way:

- Check the Trusted Cluster health.
- If the Trusted Cluster is unhealthy, remediate it.

You can use either the vSphere Client or the CLI to check the Trusted Cluster health. See [Check Trusted Cluster Health](#). You can also use either the vSphere Client or the CLI to remediate a Trusted Cluster. See [Remediate a Trusted Cluster](#).

Note: Remediation is also the appropriate process to use when you add a host to an existing Trusted Cluster. See [Add a Host to a Trusted Cluster Using the vSphere Client](#) and [Add a Host to a Trusted Cluster Using the Command Line](#).

Check Trusted Cluster Health

You can check the health status of a Trusted Cluster by using either the vSphere Client or the command line.

- The vCenter Server for the Trusted Cluster must be running vSphere 7.0 Update 1 or later.
- If you use PowerCLI, version 12.1.0 or later is required.

1. Check the Trusted Cluster health.

Tool	Steps
vSphere Client	1. Connect to the vCenter Server of the Trusted Cluster by using the vSphere Client. 2. Log in as a Trust Authority administrator. 3. Navigate to a Trusted Cluster, select Configure, then select Trust Authority. 4. Click Check Health.
CLI	1. In a PowerCLI session, run the <code>Connect-VIServer</code> cmdlet to connect as the Trust Authority administrator to the vCenter Server of the Trusted Cluster. <pre>Connect-VIServer -server TrustedCluster_VC_ip_address -User trust_admin_user -Password 'password'</pre> 2. Run the <code>Get-TrustedClusterAppliedStatus</code> cmdlet, for example: <pre>Get-TrustedClusterAppliedStatus -TrustedCluster 'TrustedCluster'</pre>

2. If there are errors, see [Remediate a Trusted Cluster](#).

Remediate a Trusted Cluster

You can remediate the configuration of a Trusted Cluster by using either the vSphere Client or the command line.

The vCenter Server for the Trusted Cluster must be running vSphere 7.0 Update 1 or later.

1. Connect to the vCenter Server of the Trusted Cluster.

Tool	Steps
vSphere Client	1. Connect to the vCenter Server of the Trusted Cluster by using the vSphere Client.

Tool	Steps
	2. Log in as a Trust Authority administrator.
CLI	In a PowerCLI session, run the <code>Connect-VIServer</code> cmdlet to connect as the Trust Authority administrator to the vCenter Server of the Trusted Cluster. <pre>Connect-VIServer -server TrustedCluster_VC_ip_address -User trust_admin_user -Password 'password'</pre>

2. Remediate the Trusted Cluster then recheck the Trusted Cluster health.

Tool	Steps
vSphere Client	1. Navigate to a Trusted Cluster. 2. Select Configure, then select Trust Authority. 3. Click Remediate. 4. Click Check Health.
CLI	1. Run the <code>Set-TrustedCluster</code> cmdlet with the <code>-Remediate</code> parameter, for example: <pre>Set-TrustedCluster -TrustedCluster 'Trusted-Cluster' -Remediate</pre> 2. Run the <code>Get-TrustedClusterAppliedStatus</code> cmdlet, for example: <pre>Get-TrustedClusterAppliedStatus -TrustedCluster 'TrustedCluster'</pre>

Using Encryption in Your vSphere Environment

Whether you use a standard key provider, trusted key provider, or vSphere Native Key Provider, using encryption in your vSphere environment requires some preparation.

See the following information to set up your environment to use a key provider:

- [Configuring and Managing a Standard Key Provider](#)
- [Configuring and Managing vSphere Native Key Provider](#)
- [Configuring vSphere Trust Authority](#)

After your environment is set up, you can use the vSphere Client to create encrypted virtual machines and virtual disks and encrypt existing virtual machines and disks.

You can perform additional tasks by using the API and by using the `crypto-util` CLI. See the *vSphere Web Services SDK Programming Guide* for the API documentation and the `crypto-util` command-line help for details about that tool.

Create an Encryption Storage Policy

Before you can create encrypted virtual machines, you must create an encryption storage policy. You create the storage policy once, and assign it each time you encrypt a virtual machine or virtual disk.

- Set up the connection to a key provider.
Although you can create a VM Encryption storage policy without the key provider connection in place, you cannot perform encryption tasks until trusted connection with the key provider is established.

- Required privileges: **Cryptographic operations > Manage encryption policies.**

If you want to use virtual machine encryption with other I/O filters, or to use the **Create VM Storage Policy** wizard in the vSphere Client, see the *vSphere Storage* documentation for details.

1. Log in to the vCenter Server by using the vSphere Client.
2. Select **Home**, click **Policies and Profiles**, then click **VM Storage Policies**.
3. Click **Create**.
4. Select the vCenter Server, enter a policy name, optionally enter a description, then click **Next**.
5. On the **Policy structure** page, check **Enable host based roles** then click **Next**.
6. On the **Host based services** page, select **Use storage policy component**, choose **Default encryption properties** from the drop-down menu, then click **Next**.
7. On the **Storage compatibility** page, leave **Compatible** selected, select a datastore, then click **Next**.
8. Review the information and click **Finish**.

The VM Encryption storage policy is added to the list, and is available for use when encrypting a virtual machine.

Activate Host Encryption Mode Explicitly

Host encryption mode must be set if you want to perform encryption tasks, such as creating an encrypted virtual machine, on an ESXi host. In most cases, host encryption mode is activated automatically when you perform an encryption task.

Required privilege: **Cryptographic operations.Register host**

Sometimes, turning on encryption mode explicitly is necessary. See [Prerequisites and Required Privileges for Virtual Machine Encryption Tasks](#).

1. Log in to the vCenter Server by using the vSphere Client.
2. Browse to the ESXi host and click **Configure**.
3. Under System, click **Security Profile**.
4. Click **Edit** in the Host Encryption Mode panel.
5. Select **Enabled** and click **OK**.

Deactivate Host Encryption Mode Using the API

Host encryption mode is activated automatically when a user performs an encryption task, if the user has sufficient privilege. After host encryption mode is activated, all core dumps are encrypted to avoid the release of sensitive information to support personnel. If you no longer use virtual machine encryption with an ESXi host, you can deactivate encryption mode.

After encryption mode is activated for an ESXi host, you might need to deactivate it. For example, you might need to deactivate encryption mode to generate an ESXi support bundle (using the `vm-support` command). Using the Host Encryption mode toggle (**Host > Configure > Security Profile > Edit Host Encryption Mode**) does not work when key material exists on the host.

You can use the API to deactivate host encryption mode by invoking the `CryptoManagerHostDisable` API method.

The crypto modes, or states, defined for an ESXi host are:

- `pendingIncapable`: The host is crypto deactivated, that is, the host cannot perform vSphere Virtual Machine Encryption operations.
- `incapable`: The host is not safe for receiving sensitive material.
- `prepared`: The host is prepared for receiving sensitive material but does not have a host key set yet.

- **safe**: The host is crypto safe (activated), and has a host key set, that is, vSphere Virtual Machine Encryption operations are possible.

After you invoke `CryptoManagerHostDisable` on a host, the crypto state of the host changes as follows:

- If the original host crypto state is incapable or prepared, the host crypto state is changed to incapable.
- If the original host crypto state is safe, the host crypto state is changed to pendingIncapable.
- If the host crypto state is pendingIncapable, the host crypto state is still pendingIncapable.

This task shows how to deactivate host encryption mode by using the vCenter Server Managed Object Browser (MOB). For more information about using the API, see the *vSphere Web Services API* documentation at <https://developer.vmware.com/apis/968/vsphere>.

1. Log in to the vCenter Server as an administrator.
2. Unregister all encrypted virtual machines from the ESXi host whose encryption mode you want to deactivate.
3. Access the MOB on the vCenter Server.

`https://vcenter_server/mob`

4. Invoke the `CryptoManagerHostDisable` method on a host.

- a) Under content name, click **content**.
- b) Under rootFolder, click **group-D1 (Datacenters)**.
- c) Under childEntity, click the appropriate datacenter.
- d) Under hostFolder, click the appropriate host.
- e) Under childEntity, click the appropriate cluster.
- f) Under host, click the appropriate host.
- g) Under configManager, click **configManager**.
- h) Under cryptoManager, click **CryptoManagerHost-*number***.
- i) Click **CryptoManagerHostDisable**.

The host crypto state is changed to either pendingIncapable or incapable, depending on its original crypto state.

5. Repeat step 4 for other hosts on which you want to deactivate encryption mode.
6. Reboot the hosts.

Once the host encryption mode is deactivated, you cannot perform encryption operations, such as adding encrypted virtual machines, unless you re-activate the host encryption mode.

Note: After you reboot an ESXi host on which you deactivated encryption mode, if the host crypto state was originally pendingIncapable, the host crypto state is still pendingIncapable. To re-activate host encryption mode, re-access the vCenter Server MOB and invoke the `ConfigureCryptoKey` API method. When re-activating host encryption mode, use the original host key ID if the host crypto state is pendingIncapable.

Create an Encrypted Virtual Machine

You can use the vSphere Client to create encrypted virtual machines.

- Establish a trusted connection with the KMS and select a default KMS.
- Create an encryption storage policy, or use the bundled sample, VM Encryption Policy.
- Ensure that the virtual machine is powered off.
- Verify that you have the required privileges:
 - **Cryptographic operations > Encrypt new**
 - If the host encryption mode is not Enabled, you also need **Cryptographic operations > Register host**.

The vSphere Client filters by virtual machine encryption storage policies, easing creation of encrypted virtual machines.

Note: Creating an encrypted virtual machine is faster and uses fewer storage resources than encrypting an existing virtual machine. If possible, encrypt virtual machines during the creation process.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, a host or a cluster.
3. Right-click the object and select **New Virtual Machine**.
4. Follow the prompts to create an encrypted virtual machine.

Option	Action
Select a creation type	Create a new virtual machine.
Select a name and folder	Specify a unique name and target location for the virtual machine.
Select a compute resource	Specify an object for which you have privileges to create encrypted virtual machines. See Prerequisites and Required Privileges for Virtual Machine Encryption Tasks .
Select storage	Select the Encrypt this virtual machine check box. Virtual machine storage policies that include encryption appear. Select a virtual machine storage policy (the bundled sample is VM Encryption Policy), and select a compatible datastore.
Select compatibility	Select the compatibility. You can migrate an encrypted virtual machine only to hosts with compatibility ESXi 6.5 and later.
Select a guest OS	Select a guest OS that you plan to install on the virtual machine later.
Customize hardware	Customize the hardware, for example, by changing disk size or CPU. (Optional) Select the VM Options tab, and expand Encryption . Select which disks to exclude from encryption. When you deselect a disk, only the VM Home and any other selected disks are encrypted. Any New Hard disk that you add is encrypted. You can change the storage policy for individual hard disks later.
Ready to complete	Review the information and click Finish .

Clone an Encrypted Virtual Machine

A cloned, encrypted virtual machine is encrypted with the same keys unless you change them. To change keys, you can use the vSphere Client, the PowerCLI, or the API. If you use the PowerCLI or the API, you can clone the encrypted virtual machine and change keys in one step.

- A key provider must be configured and enabled.
- Create an encryption storage policy, or use the bundled sample, VM Encryption Policy.
- Required privileges (applies to all key providers):
 - **Cryptographic operations > Clone**
 - **Cryptographic operations > Encrypt**
 - **Cryptographic operations > Decrypt**
 - **Cryptographic operations > Recrypt**
 - If the host encryption mode is not enabled, you also must have **Cryptographic operations > Register host** privileges.

You can perform the following operations during clone.

- Create an encrypted virtual machine from an unencrypted virtual machine or template virtual machine.

- Create a decrypted virtual machine from an encrypted virtual machine or template virtual machine.
- Recrypt the destination virtual machine with different keys from that of source virtual machine.
- In vSphere 8.0 and later, selecting the **Replace** option for a virtual machine with a vTPM starts with a new, blank vTPM, which gets its own secrets and identity.

Note: vSphere 8.0 and later includes the `vpxd.clone.tpmProvisionPolicy` advanced setting to make the default clone behavior for vTPMs to be "replace."

You can create an instant clone virtual machine from an encrypted virtual machine with the caution that the instant clone shares the same key with the source virtual machine. You cannot recrypt keys on either the source or the instant clone virtual machine.

To use the API to clone encrypted machines, see *vSphere Web Services SDK Programming Guide*.

1. Browse to the virtual machine in the vSphere Client inventory.
2. To create a clone of an encrypted machine, right-click the virtual machine, select **Clone > Clone to Virtual Machine**, and follow the prompts.
 - a) On the **Select a name and folder** page, specify a name and the target location for the clone.
 - b) On the **Select a compute resource** page, specify an object for which you have privileges.
 - c) Optional: Change the keys for the cloned vTPM.

Figure 122: Select TPM Provision Policy

VM-01 - Clone Existing Virtual Machine

- 1 Select a name and folder
- 2 Select a compute resource
- 3 Select TPM provision policy**
- 4 Select storage
- 5 Select clone options
- 6 Ready to complete

Select TPM provision policy

TPM Provision Policy

☒ Copy ☐ Replace

⚠ The virtual machine clone will be created with exact copy of source virtual machine's secrets. This may result in unintended compromise of secrets.

Cloning a virtual machine duplicates the entire virtual machine, including the vTPM and its secrets, which can be used to determine a system's identity. To change secrets on a vTPM, select **Replace** for **TPM Provision Policy**.

Note: When you replace the secrets of a vTPM, all keys, including workload-related keys, are replaced. As a best practice, ensure that your workloads no longer use a vTPM before you replace the keys. Otherwise, the workloads in the cloned virtual machine might not function correctly.

- d) On the **Select storage** page, select a datastore. You can change the storage policy as part of the clone operation. For example, changing from using an encryption policy to a non-encryption policy decrypts the disks.
 - e) On the **Select clone options** page, select clone options, as discussed in the *vSphere Virtual Machine Administration* documentation.
 - f) On the **Ready to complete** page, review the information and click **Finish**.
3. Optional: Change the keys for the cloned virtual machine.
- By default, the cloned virtual machine is created with the same keys as its parent. Best practice is to change the keys of the cloned virtual machine to ensure that multiple virtual machines do not have the same keys.
- a) Decide upon a shallow or deep recrypt.

To use a different DEK and KEK, perform a deep recrypt of the cloned virtual machine. To use a different KEK, perform a shallow recrypt of the cloned virtual machine. For a deep recrypt, you must power off the virtual machine. You can perform a shallow recrypt operation while the virtual machine is powered on, and if the virtual machine has snapshots present. Shallow recrypt of an encrypted virtual machine with snapshots is permitted only on a single snapshot branch (disk chain). Multiple snapshot branches are not supported. If the shallow recrypt fails before updating all links in the chain with the new KEK, you can still access the encrypted virtual machine if you have the old and new KEKs.
 - b) Perform a recrypt of the clone using the API. See *vSphere Web Services SDK Programming Guide*.

Encrypt an Existing Virtual Machine or Virtual Disk

You can encrypt an existing virtual machine or virtual disk by changing its storage policy. You can encrypt virtual disks only for encrypted virtual machines.

- Establish a trusted connection with the KMS and select a default KMS.
 - Create an encryption storage policy, or use the bundled sample, VM Encryption Policy.
 - Ensure that the virtual machine is powered off.
 - Verify that you have the required privileges:
 - **Cryptographic operations > Encrypt new**
 - If the host encryption mode is not Enabled, you also need **Cryptographic operations > Register host**.
1. Connect to vCenter Server by using the vSphere Client.
 2. Right-click the virtual machine that you want to change and select **VM Policies > Edit VM Storage Policies**.
You can set the storage policy for the virtual machine files, represented by VM home, and the storage policy for virtual disks.
 3. Select the storage policy.
 - To encrypt the VM and its hard disks, select an encryption storage policy and click **OK**.
 - To encrypt the VM but not the virtual disks, toggle on **Configure per disk**, select the encryption storage policy for VM Home and other storage policies for the virtual disks, and click **OK**.

You cannot encrypt the virtual disk of an unencrypted virtual machine. However, if you use the vSphere Client to encrypt the VM Home files, you can then reconfigure the unencrypted virtual machine with the encrypted disk.

4. If you prefer, you can encrypt the virtual machine, or both virtual machine and disks, from the **Edit Settings** menu in the vSphere Client.
 - a) Right-click the virtual machine and select **Edit Settings**.
 - b) Select the **VM Options** tab, and open **Encryption**. Choose an encryption policy. If you deselect all disks, only the VM home is encrypted.
 - c) Click **OK**.

Decrypt an Encrypted Virtual Machine or Virtual Disk

You can decrypt a virtual machine, its disks, or both, by changing the storage policy.

- The virtual machine must be encrypted.
- The virtual machine must be powered off or in maintenance mode.
- Required privileges: **Cryptographic operations > Decrypt**

This task describes how to decrypt an encrypted virtual machine using the vSphere Client.

All encrypted virtual machines require encrypted vMotion. During virtual machine decryption, the Encrypted vMotion setting remains. To change this setting so that Encrypted vMotion is no longer used, change the setting explicitly.

This task explains how to perform decryption using storage policies. For virtual disks, you can also perform decryption using the **Edit Settings** menu.

Note: In the Virtual Machine Details pane, a vTPM-enabled virtual machine displays both a lock icon and an "Encrypted with *key_provider*" message. To remove a vTPM from a virtual machine, see [Remove Virtual Trusted Platform Module from a Virtual Machine](#).

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine that you want to change and select **VM Policies > Edit VM Storage Policies**.
You can set the storage policy for the virtual machine files, represented by VM home, and the storage policy for virtual disks.
3. Select a storage policy.
 - To decrypt the VM and its hard disks, toggle off **Configure per disk**, select a storage policy from the drop-down menu, and click **OK**.
 - To decrypt a virtual disk but not the virtual machine, toggle on **Configure per disk**, select the encryption storage policy for VM Home and other storage policies for the virtual disks, and click **OK**.

You cannot decrypt the virtual machine and leave the disk encrypted.

4. If you prefer, you can use the vSphere Client to decrypt the virtual machine and disks from the **Edit Settings** menu.
 - a) Right-click the virtual machine and select **Edit Settings**.
 - b) Select the **VM Options** tab and expand **Encryption**.
 - c) To decrypt the VM and its hard disks, choose **None** from the **Encrypt VM** drop-down menu.
 - d) To decrypt a virtual disk but not the virtual machine, deselect the disk.
 - e) Click **OK**.
5. Optional: You can change the Encrypted vMotion setting.
 - a) Right-click the virtual machine and click **Edit Settings**.
 - b) Click **VM Options**, and open **Encryption**.
 - c) Set the **Encrypted vMotion** value.

Change the Encryption Policy for Virtual Disks

When you create an encrypted virtual machine from the vSphere Client, you can select which virtual disks that you add during virtual machine creation are encrypted. You can decrypt virtual disks that are encrypted by using the **Edit VM Storage Policies** option.

- You must have the **Cryptographic operations > Manage encryption policies** privilege.
- Ensure that the virtual machine is powered off.

Note: An encrypted virtual machine can have virtual disks that are not encrypted. However, an unencrypted virtual machine cannot have encrypted virtual disks.

See [Virtual Disk Encryption](#).

This task describes how to change the encryption policy using storage policies. You can also use the **Edit Settings** menu to make this change.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine and select **VM Policies > Edit VM Storage Policies**.
3. Change the storage policy.
 - To change the storage policy for the VM and its hard disks, select an encryption storage policy and click **OK**.
 - To encrypt the VM but not the virtual disks, toggle on **Configure per disk**, select the encryption storage policy for VM Home and other storage policies for the virtual disks, and click **OK**.

You cannot encrypt the virtual disk of an unencrypted VM.

4. If you prefer, you can change the storage policy from the **Edit Settings** menu.
 - a) Right-click the virtual machine and select **Edit Settings**.
 - b) Select the **Virtual Hardware** tab, expand a hard disk, and select an encryption policy from the drop-down menu.
 - c) Click **OK**.

Resolve Missing Encryption Key Issues

If the ESXi host cannot get the key (KEK) from vCenter Server for an encrypted virtual machine or an encrypted virtual disk, the encrypted virtual machine becomes locked. After you make the keys available on the key server (KMS), you can unlock a locked encrypted virtual machine.

Under certain circumstances when using a standard key provider, the ESXi host cannot get the key encryption key (KEK) for an encrypted virtual machine or an encrypted virtual disk from vCenter Server. In that case, you can still unregister or reload the virtual machine. However, you cannot perform other virtual machine operations such as powering on the virtual machine. After taking the necessary steps to make the required keys available on the key server, you can unlock a locked encrypted virtual machine by using the vSphere Client.

If the virtual machine key is not available, a vCenter Server alarm notifies you and the state of the virtual machine displays as invalid. The virtual machine cannot power on. If the virtual machine key is available, but a key for an encrypted disk is not available, the virtual machine state does not display as invalid. However, the virtual machine cannot power on and the following error results:

```
The disk [/path/to/the/disk.vmdk] is encrypted and a required key was not found.
```

Note: The following procedure illustrates the situations that can cause a virtual machine to become locked, the corresponding alarms and event logs that appear, and what to do in each case.

1. If the problem is the connection between the vCenter Server system and the key server, vCenter Server generates a virtual machine alarm. Also, an error message appears in the event log.

Restore the connection to the key server. When the key server and keys become available, unlock the locked virtual machines. See [Unlock Locked Virtual Machines](#). You can also reboot the host and re-register the virtual machine to unlock it after restoring the connection.

Losing the connection to the key server does not automatically lock the virtual machine. The virtual machine only enters a locked state if the following conditions are met:

- The key is not available on the ESXi host.
- vCenter Server cannot retrieve keys from the key server.

After each ESXi reboot, it is desirable, though not necessary, to have vCenter Server started first. vCenter Server requests the key with the corresponding ID from the key server and makes it available to ESXi.

Note: In vSphere 7.0 Update 2 and later, you can persist encryption keys across ESXi reboots. See [vSphere Key Persistence on ESXi Hosts](#).

If, after restoring connection to the key provider, the virtual machine remains locked, see [Unlock Locked Virtual Machines](#).

2. If the connection is restored, register the virtual machine. If an error results, or if the operation succeeds but the virtual machine is in a locked state, verify that you have the **Cryptographic operations > RegisterVM** privilege for the vCenter Server system.

This privilege is not required for powering on an encrypted virtual machine if the key is available. This privilege is required for registering the virtual machine if the key has to be retrieved.

3. If the key is no longer available on the key server, vCenter Server generates a virtual machine alarm. Also, an error message appears in the event log.

Ask the key server administrator to restore the key. You might encounter an inactive key if you are powering on a virtual machine that had been removed from the inventory and that had not been registered for a long time. It also happens if you reboot the ESXi host, and the key server is not available.

- a) Retrieve the key ID by using the Managed Object Browser (MOB) or the vSphere API.

Retrieve the `keyId` from `VirtualMachine.config.keyId.keyId`.

- b) Ask the key server administrator to reactivate the key that is associated with that key ID.
- c) After restoring the key, see [Unlock Locked Virtual Machines](#).

If the key can be restored on the key server, vCenter Server retrieves it and pushes it to the ESXi host the next time it is needed.

4. If the key server is accessible and the ESXi host is powered on, but the vCenter Server system is unavailable, follow these steps to unlock virtual machines.

- a) Restore the vCenter Server system, or set up a different vCenter Server system, then establish trust with the key server.

You must use the same key provider name, but the key server IP address can be different.

- b) Reregister all virtual machines that are locked.

The new vCenter Server instance retrieves the keys from the key server and the virtual machines are unlocked.

5. If the keys are missing only on the ESXi host, vCenter Server generates a virtual machine alarm and the following message appears in the event log:

`Virtual machine is locked because keys are missing on host.`

The vCenter Server system can retrieve the missing keys from the key provider. No manual recovery of keys is required. See [Unlock Locked Virtual Machines](#).

Unlock Locked Virtual Machines

A vCenter Server alarm notifies you when an encrypted virtual machine is in a locked state. You can unlock a locked encrypted virtual machine by using the vSphere Client after taking the necessary steps to make the required keys available on the key server.

- Verify that you have the required privileges: **Cryptographic operations.RegisterVM**
 - Other privileges might be required for optional tasks such as enabling host encryption.
 - Before unlocking a locked virtual machine, troubleshoot the cause of the lock and attempt to fix the problem manually. See [Resolve Missing Encryption Key Issues](#).
1. Connect to vCenter Server by using the vSphere Client.
 2. Navigate to the virtual machine's **Summary** tab.
When a virtual machine is locked, the Virtual Machine Locked alarm appears.
 3. Decide if you want to either acknowledge the alarm, or reset the alarm to green but not unlock the virtual machine now.
When you click either **Acknowledge** or **Reset to green**, the alarm goes away, but the virtual machine remains locked until you unlock it.
 4. Navigate to the **Monitor** tab of the virtual machine and click **Events**.
The **Events** pane displays information about why the virtual machine is locked.
 5. Perform suggested troubleshooting before you unlock the virtual machine.
 6. Navigate to the **Summary** tab of the virtual machine.
A Virtual Machine Locked Alarm appears.
 7. Select **Unlock VM** from the **Actions** drop-down menu on the right side.

Resolve ESXi Host Encryption Mode Issues

Under certain circumstances, the ESXi host's encryption mode can become deactivated.

An ESXi host requires that host encryption mode is activated if it contains any encrypted virtual machines. If the host detects it is missing its host key, or if the key provider is unavailable, the host might fail to activate the encryption mode. vCenter Server generates an alarm when the host encryption mode cannot be activated.

1. If the problem is the connection between the vCenter Server system and the key provider, an alarm is generated and an error message appears in the event log.
You must restore the connection to the key provider that contains the encryption keys in question.
2. If keys are missing, an alarm is generated and an error message appears in the event log.
You must ensure that the keys are present in the key provider. Consult the documentation for your key management vendor for information about restoring from backup.

If, after restoring connection to the key provider, or manually recovering keys to the key provider, the host's encryption mode remains deactivated, re-activate the host encryption mode. See [Re-Activate ESXi Host Encryption Mode](#).

Re-Activate ESXi Host Encryption Mode

In vSphere 6.7 and later, a vCenter Server alarm notifies you when an ESXi host's encryption mode has become deactivated. You can re-activate the host encryption mode if it has become deactivated.

- Verify that you have the required privileges: **Cryptographic operations > Register host**
 - Before re-activating encryption mode, troubleshoot the cause and attempt to fix the problem manually.
1. Connect to vCenter Server by using the vSphere Client.
 2. Navigate to the **Summary** tab for the ESXi host.
When the encryption mode is deactivated, the Host Requires Encryption Mode Enabled alarm appears.
 3. Decide if you want to either acknowledge the alarm, or reset the alarm to green but not re-activate the host encryption mode now.
When you click either **Acknowledge** or **Reset to green**, the alarm goes away, but the encryption mode for the host remains deactivated until you re-activate it.
 4. Navigate to the **Monitor** tab for the ESXi host and click **Events**.
More information is displayed about why encryption mode is deactivated. Perform suggested troubleshooting before you re-activate the encryption mode.
 5. On the **Summary** tab, click **Enable Host Encryption Mode** to re-activate host encryption.
A message appears, warning that encryption key data is transmitted to the host.
 6. Click **Yes**.

Set Key Server Certificate Expiration Threshold

By default, vCenter Server notifies you 30 days before your key server (KMS) certificates expire. You can change this default value.

Key server certificates have an expiration date. When the threshold for the expiration date is reached, an alarm notifies you.

vCenter Server and key servers exchange two types of certificates: server and client. The VMware Endpoint Certificate Store (VECS) on the vCenter Server system stores the server certificates and one client certificate per key provider. Because there are two certificate types, there are two alarms for each certificate type (one for client, one for server).

1. Log in to a vCenter Server system by using the vSphere Client.
2. Select the vCenter Server system in the object hierarchy.
3. Click **Configure**.
4. Under **Settings**, click **Advanced Settings**, and click **Edit Settings**.
5. Click the **Filter** icon and enter `vpzd.kmscert.threshold`, or scroll to the configuration parameter itself.
6. Enter your value in days and click **Save**.

vSphere Virtual Machine Encryption and Core Dumps

If your environment uses vSphere Virtual Machine Encryption, and if an error occurs on the ESXi host, the resulting core dump is encrypted to protect customer data. Core dumps that are included in the vm-support package are also encrypted.

Note: Core dumps can contain sensitive information. Follow your organization's data security and privacy policy when handling core dumps.

Core Dumps on ESXi Hosts

When an ESXi host, a user world, or a virtual machine fails, a core dump is generated, and the host reboots. If the ESXi host has encryption mode enabled, the core dump is encrypted using a key that is in the ESXi key cache. (Depending on the key provider in use, the key comes from an external key server, the Key Provider Service, or vCenter Server). See [How vSphere Virtual Machine Encryption Protects Your Environment](#) for background information.

When an ESXi host is cryptographically "safe," and a core dump is generated, an event is created. The event indicates that a core dump occurred along with the following information: world name, occurring times, keyID of the key used to encrypt the core dump, and core dump filename. You can view the event in the Events viewer under **Tasks and Events** for the vCenter Server.

The following table shows encryption keys used for each core dump type, by vSphere release.

Table 270: Core Dump Encryption Keys

Core Dump Type	Encryption Key (ESXi 6.5)	Encryption Key (ESXi 6.7 and Later)
ESXi Kernel	Host Key	Host Key
User World (hostd)	Host Key	Host Key
Encrypted Virtual Machine (VM)	Host Key	Virtual Machine Key

What you can do after an ESXi host reboot depends on several factors.

- In most cases, the key provider attempts to push the key to the ESXi host after reboot. If the operation is successful, you can generate the vm-support package and you can decrypt or re-encrypt the core dump. See [Decrypt or Re-Encrypt an Encrypted Core Dump](#).
- If vCenter Server cannot connect to the ESXi host, you might be able to retrieve the key. See [Resolve Missing Encryption Key Issues](#).
- If the host used a custom key, and that key differs from the key that vCenter Server pushes to the host, you cannot manipulate the core dump. Avoid using custom keys.

Core Dumps and vm-support Packages

When you contact VMware Technical Support because of a serious error, your support representative usually asks you to generate a vm-support package. The package includes log files and other information, including core dumps. If your support representatives cannot resolve the issues by looking at log files and other information, they might ask you to decrypt the core dumps and make relevant information available. To protect sensitive information such as keys, follow your organization's security and privacy policy. See [Collect a vm-support Package for an ESXi Host That Uses Encryption](#).

Core Dumps on vCenter Server Systems

A core dump on a vCenter Server system is not encrypted. vCenter Server already contains potentially sensitive information. At the minimum, ensure that the vCenter Server is protected. See [Securing vCenter Server Systems](#). You might also consider turning off core dumps for the vCenter Server system. Other information in log files can help determine the problem.

Collect a vm-support Package for an ESXi Host That Uses Encryption

If host encryption mode is enabled for the ESXi host, any core dumps in the `vm-support` package are encrypted. You can collect the package from the vSphere Client, and you can specify a password if you expect to decrypt the core dump later.

Inform your support representative that host encryption mode is enabled for the ESXi host. Your support representative might ask you to decrypt core dumps and extract relevant information.

Note: Core dumps can contain sensitive information. Follow your organization's security and privacy policy to protect sensitive information such as host keys.

The `vm-support` package includes log files, core dump files, and more.

1. Log in to the vCenter Server system with the vSphere Client.
2. Click **Hosts & Clusters**, and right-click the ESXi host.
3. Select **Export System Logs**.
4. In the dialog box, select **Password for encrypted core dumps**, and specify and confirm a password.
5. Leave the defaults for other options or make changes if requested by VMware Technical Support, and click **Export Logs**.

If you have not configured your browser to ask where to save files before downloading, the download starts. If you have configured your browser to ask where to save files, specify a location for the file.

6. If your support representative asked you to decrypt the core dump in the `vm-support` package, log in to any ESXi host and follow these steps.

- a) Log in to the ESXi host and connect to the directory where the `vm-support` package is located.

The filename follows the pattern `esx.date_and_time.tgz`.

- b) Make sure that the directory has enough space for the package, the uncompressed package, and the recompressed package, or move the package.
- c) Extract the package to the local directory.

```
vm-support -x *.tgz .
```

The resulting file hierarchy might contain core dump files for the ESXi host, usually in `/var/core`, and might contain multiple core dump files for virtual machines.

- d) Decrypt each encrypted core dump file separately.

```
crypto-util envelope extract --offset 4096 --keyfile vm-support-incident-key-file
--password encryptedZdumpdecryptedZdump
```

`vm-support-incident-key-file` is the incident key file that you find at the top level in the directory.

`encryptedZdump` is the name of the encrypted core dump file.

`decryptedZdump` is the name for the file that the command generates. Make the name similar to the `encryptedZdump` name.

- e) Provide the password that you specified when you created the `vm-support` package.
- f) Remove the encrypted core dumps, and compress the package again.

```
vm-support --reconstruct
```

7. Remove any files that contain confidential information.

Decrypt or Re-Encrypt an Encrypted Core Dump

You can decrypt or re-encrypt an encrypted core dump on your ESXi host by using the `crypto-util` CLI.

The key that was used to encrypt the core dump must be available on the ESXi host that generated the core dump.

You can decrypt and examine the core dumps in the `vm-support` package yourself. Core dumps might contain sensitive information. Follow the security and privacy policy for your organization to protect sensitive information such as keys.

For details about re-encrypting a core dump and other features of `crypto-util`, see the command-line help.

Note: `crypto-util` is for advanced users.

1. Log directly in to the ESXi host on which the core dump happened.
If the ESXi host is in lockdown mode, or if SSH access is deactivated, you might have to activate access first.
2. Determine whether the core dump is encrypted.

Option	Description
Monitor core dump	<code>crypto-util envelope describe vmmcores.ve</code>
zdump file	<code>crypto-util envelope describe --offset 4096 zdump-File</code>

3. Decrypt the core dump, depending on its type.

Option	Description
Monitor core dump	<code>crypto-util envelope extract vmmcores.ve vmmcores</code>
zdump file	<code>crypto-util envelope extract --offset 4096 zdumpEncryptedzdumpUnencrypted</code>

Activate and Deactivate Key Persistence on an ESXi Host

You must activate key persistence on an ESXi host. It is not activated by default.

Requirements to activate key persistence:

- ESXi 7.0 Update 2 or later
- ESXi host installed with TPM 2.0
- Have access to the ESXCLI command set. You can run ESXCLI commands remotely, or run them in the ESXi Shell.

Note: Key persistence is not necessary when using vSphere Native Key Provider. vSphere Native Key Provider is designed out-of-the-box to run without requiring access to a key server.

For additional security, the TPM can also use a sealing policy to prevent tampering during ESXi host boot. See [What Are TPM Sealing Policies](#).

For conceptual information about key persistence, see [vSphere Key Persistence on ESXi Hosts](#).

1. Start a session on the ESXi host by using SSH or another remote console connection.
2. Log in as root.
3. Verify that the ESXi host is in TPM mode.

```
esxcli system settings encryption get
```

If the Mode appears as NONE, you must enable the TPM in the firmware of the host, and set the mode by running the following command.

```
esxcli system settings encryption set --mode=TPM
```

4. Activate or deactivate key persistence.

- a) To activate key persistence:

```
esxcli system security keypersistence enable
```

- b) To deactivate persistence:

```
esxcli system security keypersistence disable --remove-all-stored-keys
```

Rekey an Encrypted Virtual Machine Using the vSphere Client

You can use the vSphere Client to perform a shallow rekey of an encrypted virtual machine. You might perform a rekey of an encrypted virtual machine for business or compliance reasons.

Required privilege: **Cryptographic operations > Recrypt**

A shallow rekey (also called recrypt) replaces only the Key Encryption Key (KEK). You do not need to power off the encrypted virtual machine to perform a shallow rekey. If you need to replace both the Disk Encryption Key (DEK) and the KEK, you must perform a deep rekey.

Note: Virtual machines configured with IDE controllers must be powered off to perform a shallow rekey operation.

For more conceptual information, see [How Do You Recrypt \(Rekey\) an Encrypted Virtual Machine](#).

1. Log in to the vCenter Server system with the vSphere Client.
2. Browse the inventory list and select the encrypted virtual machine.
3. Right-click the encrypted virtual machine and select **VM Policies**.
4. Select **Re-encrypt**.
5. Click **Yes**.

The encrypted virtual machine is rekeyed with the new KEK.

Note: If the rekey fails, the events subsystem posts the following event:

```
com.vmware.vc.vm.crypto.RekeyFail
```

Rekey an Encrypted Virtual Machine Using the CLI

You can use the CLI to perform a shallow rekey of an encrypted virtual machine. You might perform a rekey of an encrypted virtual machine for business or compliance reasons.

Required privilege: **Cryptographic operations > Recrypt**

Note: Virtual machines configured with IDE controllers must be powered off to perform a shallow rekey operation.

A shallow key (also called recrypt) replaces only the Key Encryption Key (KEK). You do not need to power off the encrypted virtual machine to perform a shallow rekey. If you need to replace both the Disk Encryption Key (DEK) and the KEK, you must perform a deep rekey.

This task shows how to perform a shallow rekey on an encrypted virtual machine using the currently assigned key provider.

For more conceptual information, see [How Do You Recrypt \(Rekey\) an Encrypted Virtual Machine](#).

1. In a PowerCLI session, run the `Connect-VIServer` cmdlet to connect as an administrator to the vCenter Server host.
2. Assign the current key provider to a variable.

```
$kp = Get-KeyProvider keyprovider_name
```

3. Assign the encrypted virtual machine to a variable.

```
$vm = Get-VM encrypted_vm_name
```

4. Check the security information for the encrypted virtual machine.

```
Get-SecurityInfo -Entity $vm
```

Note the EncryptionKeyId.

5. Perform the shallow rekey of the encrypted virtual machine.

```
Set-VM -vm $vm -KeyProvider $kp
```

Type `Y` to confirm the rekey.

6. To verify that the EncryptionKeyId is changed, check the security information for the encrypted virtual machine.

```
Get-SecurityInfo -Entity $vm
```

Set the Default Key Provider Using the vSphere Client

You must set the default key provider if you do not make the first key provider the default, or if your environment uses multiple key providers and you remove the default one. You can use the vSphere Client to set the default key provider at the vCenter Server level.

As a best practice, verify that the Connection Status in the Key Providers tab shows Active and a green check mark.

1. Log in using the vSphere Client.
2. Navigate to the .
3. Click **Configure** and select **Key Providers** under **Security**.
4. Select the key provider.
5. Click **Set as Default**.

A confirmation dialog box appears.

6. Click **Set as Default**.

The key provider displays as the current default.

Set the Default Key Provider Using the Command Line

You must set the default key provider if you do not make the first key provider the default, or if your environment uses multiple key providers and you remove the default one. You can use PowerCLI to set the default key provider at the vCenter Server level, the cluster level, or the cluster folder level.

As a best practice, verify that the Connection Status in the Key Providers tab shows Active and a green check mark.

You must have a role that includes the **Cryptographic operations > Manage KMS** privilege. In vSphere Trust Authority, the role must be applied to the Trusted Cluster.

1. Ensure that you are connected as an administrator to the vCenter Server where you created the key provider.

Note: In vSphere Trust Authority, connect to the vCenter Server of the Trusted Cluster.

```
Connect-VIServer -server VC_ip_address -User admin_user -Password 'password'
```

2. Obtain the key provider.

```
Get-KeyProvider
```

You can use the `-Namekeyprovider` option to specify a single key provider.

3. Assign the `Get-KeyProvider` key provider information to a variable.

For example, this command assigns the information to the variable `$kp`.

```
$kp = Get-KeyProvider
```

If you have multiple key providers, you can use `Select-Object` to select one of them.

```
$kp = Get-KeyProvider | Select-Object -Index 0
```

4. Use one of the following PowerCLI commands.

Where to set the default	Command
vCenter Server level	<code>Set-KeyProvider -KeyProvider \$kp -DefaultForSystem</code>
Cluster level	This example command sets the key provider for the cluster CL-01 . <pre>Add-EntityDefaultKeyProvider -KeyProvider \$kp -Entity 'CL-01'</pre>

Where to set the default	Command
Cluster Folder level	This example command sets the key provider for the cluster folder <code>Cluster-Folder-01</code>. <pre>Add-EntityDefaultKeyProvider -KeyProvider \$kp -Entity 'Cluster-Folder-01'</pre>

Securing Virtual Machines with Virtual Trusted Platform Module

With the Virtual Trusted Platform Module (vTPM) feature, you can add a TPM 2.0 virtual cryptoprocessor to a virtual machine.

A vTPM is a software-based representation of a physical Trusted Platform Module 2.0 chip. A vTPM acts as any other virtual device. You can add a vTPM to a virtual machine in the same way you add virtual CPUs, memory, disk controllers, or network controllers. A vTPM does not require a hardware Trusted Platform Module chip.

What Is a Virtual Trusted Platform Module

A virtual Trusted Platform Module (vTPM) is a software-based representation of a physical Trusted Platform Module 2.0 chip. A vTPM acts as any other virtual device.

vTPMs provide hardware-based, security-related functions such as random number generation, attestation, key generation, and more. When added to a virtual machine, a vTPM enables the guest operating system to create and store keys that are private. These keys are not exposed to the guest operating system itself. Therefore, the virtual machine attack surface is reduced. Usually, compromising the guest operating system compromises its secrets, but enabling a vTPM greatly reduces this risk. These keys can be used only by the guest operating system for encryption or signing. With an attached vTPM, a client can remotely attest the identity of the virtual machine, and verify the software that it is running.

Note: By default, no storage policy is associated with a virtual machine that has been enabled with a vTPM. Only the virtual machine files (VM Home) are encrypted. If you prefer, you can choose to add encryption explicitly for the virtual machine and its disks, but the virtual machine files would have already been encrypted.

How Do You Configure a vTPM for a Virtual Machine

From the perspective of the virtual machine, a vTPM is a virtual device. You can add a vTPM to either a new or an existing virtual machine. A vTPM depends on virtual machine encryption to secure vital TPM data, thus, it requires that you configure a key provider. When you configure a vTPM, the virtual machine files are encrypted but not the disks. You can choose to add encryption explicitly for the virtual machine and its disks.

When you back up a virtual machine enabled with a vTPM, the backup must include all virtual machine data, including the `*.nvram` file. If your backup does not include the `*.nvram` file, you cannot restore a virtual machine with a vTPM. Also, because the VM home files of a vTPM-enabled virtual machine are encrypted, ensure that the encryption keys are available at the time of a restore.

In vSphere 8.0 and later, when cloning a virtual machine with a vTPM, selecting the **Replace** option for a virtual machine with a vTPM starts with a new, blank vTPM, which gets its own secrets and identity. When you replace the secrets of a vTPM, all keys, including workload-related keys, are replaced. As a best practice, ensure that your workloads no longer use a vTPM before you replace the keys. Otherwise, the workloads in the cloned virtual machine might not function correctly.

vSphere Requirements for vTPMs

To use a vTPM, your vSphere environment must meet these requirements:

- Virtual machine requirements:
 - EFI firmware
 - Hardware version 14 and later

- Component requirements:
 - vCenter Server 6.7 and later for Windows virtual machines, vCenter Server 7.0 Update 2 and later for Linux virtual machines.
 - Virtual machine encryption (to encrypt the virtual machine home files).
- Guest OS support:
 - Linux
 - Windows Server 2008 and later
 - Windows 7 and later

Differences Between a Hardware TPM and a Virtual TPM

You use a hardware Trusted Platform Module (TPM) to provide secure storage of credentials or keys. A vTPM performs the same functions as a TPM, but it performs cryptographic coprocessor capabilities in software. A vTPM uses the `.nvram` file, which is encrypted using virtual machine encryption, as its secure storage.

A hardware TPM includes a preloaded key called the Endorsement Key (EK). The EK has a private and public key. The EK provides the TPM with a unique identity. For a vTPM, this key is provided either by the VMware Certificate Authority (VMCA) or by a third-party Certificate Authority (CA). After the vTPM uses a key, it is typically not changed because doing so invalidates sensitive information stored in the vTPM. The vTPM does not contact the third-party CA at any time.

Create a Virtual Machine with a Virtual Trusted Platform Module

You can add a Virtual Trusted Platform Module (vTPM) when you create a virtual machine to provide enhanced security to the guest operating system. You must create a key provider before you can add a vTPM.

- The guest OS you use can be Windows Server 2008 and later, Windows 7 and later, or Linux.
- The ESXi hosts running in your environment must be ESXi 6.7 or later (Windows guest OS), or 7.0 Update 2 (Linux guest OS).
- The virtual machine must use EFI firmware.
- Verify that you have the required privileges:
 - **Cryptographic operations > Clone**
 - **Cryptographic operations > Encrypt**
 - **Cryptographic operations > Encrypt new**
 - **Cryptographic operations > Migrate**
 - **Cryptographic operations > Register VM**

Note: After creating a virtual machine with a vTPM, the **Cryptographic operations > Direct Access** privilege is required to open a console session.

The VMware virtual TPM is compatible with TPM 2.0 and creates a TPM-enabled virtual chip for use by the virtual machine and the guest OS it hosts.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, an host or a cluster.
3. Right-click the object, select **New Virtual Machine**, and follow the prompts to create a virtual machine.

Option	Action
Select a creation type	Create a new virtual machine.
Select a name and folder	Specify a name and target location.
Select a compute resource	
Select storage	Select a compatible datastore.
Select compatibility	You must select ESXi 6.7 and later for Windows guest OS, or ESXi 7.0 U2 and later for Linux guest OS.

Option	Action
Select a guest OS	Select Windows or Linux for use as the guest OS.
Customize hardware	Click Add New Device and select Trusted Platform Module . You can further customize the hardware, for example, by changing disk size or CPU.
Ready to complete	Review the information and click Finish .

The vTPM-enabled virtual machine appears in your inventory as specified.

Add Virtual Trusted Platform Module to an Existing Virtual Machine

You can add a Virtual Trusted Platform Module (vTPM) to an existing virtual machine to provide enhanced security to the guest operating system. You must create a key provider before you can add a vTPM.

- The guest OS you use can be Windows Server 2008 and later, Windows 7 and later, or Linux.
- Verify that the virtual machine is turned off.
- The ESXi hosts running in your environment must be ESXi 6.7 and later (Windows guest OS), or 7.0 Update 2 and later (Linux guest OS).
- The virtual machine must use EFI firmware.
- Verify that you have the required privileges:
 - **Cryptographic operations > Clone**
 - **Cryptographic operations > Encrypt**
 - **Cryptographic operations > Encrypt new**
 - **Cryptographic operations > Migrate**
 - **Cryptographic operations > Register VM**
 - **Virtual machine > Change Configuration > Add or remove device**

Note: After adding a vTPM to a virtual machine, the **Cryptographic operations > Direct Access** privilege is required to open a console session.

The VMware virtual TPM is compatible with TPM 2.0, and creates a TPM-enabled virtual chip for use by the virtual machine and the guest OS it hosts.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. In the **Edit Settings** dialog box, click **Add New Device** and select **Trusted Platform Module**.
4. Click **OK**.

The **Virtual Machine Details** pane reflects that encryption has been applied to the virtual machine.

Remove Virtual Trusted Platform Module from a Virtual Machine

You can remove Virtual Trusted Platform Module (vTPM) security from a virtual machine.

- Ensure that the virtual machine is powered off.
- Verify that you have the required privileges: **Virtual machine > Change Configuration > Add or remove device** and **Cryptographic operations > Decrypt**

Removing a vTPM device causes all encrypted information on the virtual machine to become unrecoverable. Before removing a vTPM from a virtual machine, deactivate any applications in the Guest OS that use the vTPM device, such

as BitLocker. Failure to do so can cause the virtual machine not to boot. Also, you cannot remove a vTPM from a virtual machine that contains snapshots.

1. Connect to vCenter Server by using the vSphere Client.
2. Right-click the virtual machine in the inventory that you want to modify and select **Edit Settings**.
3. In the **Virtual Hardware** tab, expand **Security Devices**.
4. Click the ellipses icon for Virtual TPM.
5. Click **Remove device**.
6. Click **Delete** to confirm you want to remove the vTPM.
The vTPM device is marked for removal.
7. Click **OK**.

Identify Virtual Trusted Platform Module Enabled Virtual Machines

You can identify which of your virtual machines are enabled to use a Virtual Trusted Platform Module (vTPM).

You can generate a list of all virtual machines in your inventory showing virtual machine name, operating system, and vTPM status. You can also export this list to a CSV file for use in compliance audits.

1. Connect to vCenter Server by using the vSphere Client.
2. Select a vCenter Server instance, a host, or a cluster.
3. Click the **VMs** tab and click **Virtual Machines**.
4. To view all virtual machines on which a TPM is enabled, click **Manage Columns** and select **TPM**.
The TPM column displays "Present" for virtual machines on which a TPM is enabled. Virtual machines without a TPM are listed as "Not present."
5. To export the contents of an inventory list view to a CSV file, click **Export**.

View Virtual Trusted Platform Module Device Certificates

Virtual Trusted Platform Module (vTPM) devices are pre-configured with default certificates, which you can review.

You must have a vTPM-enabled virtual machine in your environment.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, a host or a cluster.
3. To identify vTPM-enabled virtual machines, click **VMs** and click **Virtual Machines**.
If necessary, click **Manage Columns** and select **TPM** to display virtual machines with a TPM "Present."
4. Select the vTPM-enabled virtual machine whose certificate information you want to view.
5. Click the virtual machine's **Configure** tab.
6. Under **TPM**, select **Certificates**.
7. Select the certificate and view its information.
8. (Optional) To export the certificate information, click **Export**.
The certificate is saved to disk.

You can replace the default certificate with a certificate issued by a third-party certificate authority (CA). See [Export and Replace Virtual Trusted Platform Module Device Certificates](#).

Export and Replace Virtual Trusted Platform Module Device Certificates

You can replace the default certificate that comes with a Virtual Trusted Platform Module (vTPM) device.

You must have a vTPM-enabled virtual machine in your environment.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, an host or a cluster.
3. Select the vTPM-enabled virtual machine in the inventory whose certificate information you want to replace.
4. Click the **Configure** tab.
5. Under **TPM** select **Signing Requests**.
6. Select a certificate.
7. To export the certificate information, click **Export**.
The certificate is saved to disk.
8. Get a certificate issued by a third-party certificate authority (CA) against the certificate signing request (CSR) you exported.
You can use any CA that you might have in your IT environment.
9. When you have the new certificate, replace the existing certificate.
 - a) Right-click the virtual machine in the inventory whose certificate you want to replace and select **Edit Settings**.
 - b) In the **Edit Settings** dialog box, expand **Security Devices**, then expand **Virtual TPM**.
The certificates appear.
 - c) Click **Replace** for the certificate you want to replace.
The **File Upload** dialog box appears.
 - d) On your local machine, locate the new certificate and upload it.
The new certificate replaces the default certificate that came with your vTPM device.
 - e) The certificate information is updated in the Certificates list.

Securing Windows Guest Operating Systems with Virtualization-based Security

In vSphere 6.7 and later, you can enable Microsoft virtualization-based security (VBS) on supported Windows guest operating systems.

Microsoft VBS, a feature introduced in Windows 10 and Windows Server 2016 operating systems, uses hardware and software virtualization to enhance system security by creating an isolated, hypervisor-restricted, specialized subsystem.

VBS permits you to use the following Windows security features to harden your system and isolate key system and user secrets from being compromised:

- **Credential Guard:** Aims to isolate and harden key system and user secrets against compromise.
- **Device Guard:** Provides a set of features designed to work together to prevent and eliminate malware from running on a Windows system.
- **Configurable Code Integrity:** Ensures that only trusted code runs from the boot loader onwards.

See the topic on virtualization-based security in the Microsoft documentation for more information.

After you enable VBS for a virtual machine through vCenter Server, you enable VBS within the Windows guest operating system.

vSphere Virtualization-based Security Best Practices

Follow best practices for virtualization-based security (VBS) to maximize security and manageability of your Windows guest operating system environment.

Avoid problems by following these best practices.

VBS Hardware Requirements

Use the following hardware for VBS:

- Intel
 - Haswell CPU or later. For best performance, use the Skylake-EP CPU or later.
 - The Ivy Bridge CPU is acceptable.
 - The Sandy Bridge CPU might cause some slow performance.
- AMD
 - Zen 2 series CPUs (Rome) or later.
 - Older CPUs might cause some slow performance.

The mitigations for the Machine Check Exception on Page Size Change Intel CPU vulnerability can impact guest OS performance negatively when VBS is in use. For more information, see the VMware KB article at <https://kb.vmware.com/kb/76050>.

VBS and Windows Guest OS Compatibility

On Intel, VBS is supported for Windows 10 and Windows Server 2016 and later virtual machines, although Windows Server 2016 versions 1607 and 1703 require patches. Check the Microsoft documentation for ESXi host hardware compatibility. Using Intel CPUs for VBS requires vSphere 6.7 or later and hardware version 14 or later.

On AMD, VBS is supported on Windows 10, version 1809, and Windows 2019 and later virtual machines. Using AMD CPUs for VBS requires vSphere 7.0 Update 2 or later and hardware version 19 or later.

Initially, Windows 10 required that you enable Hyper-V for VBS. Enabling Hyper-V is not required for Windows 10. The same applies to Windows Server 2016 and later. Consult the current Microsoft documentation and the *VMware vSphere Release Notes* for more information.

Unsupported VMware Features on VBS

The following features are not supported in a virtual machine when VBS is enabled:

- Fault tolerance
- PCI passthrough
- Hot add of CPU or memory

Installation and Upgrade Caveats with VBS

Before you configure VBS, understand the following installation and upgrade caveats:

- New virtual machines configured for Windows 10 and Windows Server 2016 and later on virtual hardware versions less than version 14 are created using Legacy BIOS by default. You must reinstall the guest operating system after changing the virtual machine's firmware type from Legacy BIOS to UEFI.
- If you plan to migrate your virtual machines from previous vSphere releases to vSphere 6.7 or later, and enable VBS on your virtual machines, use UEFI to avoid having to reinstall the operating system.

Activate Virtualization-based Security on a Virtual Machine

You can activate Microsoft virtualization-based security (VBS) for supported Windows guest operating systems at the same time you create a virtual machine.

See [vSphere Virtualization-based Security Best Practices](#) for acceptable CPUs.

Using Intel CPUs for VBS requires vSphere 6.7 or later. Create a virtual machine that uses hardware version 14 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit) or later releases
- Windows Server 2016 (64 bit) or later releases

Using AMD CPUs for VBS requires vSphere 7.0 Update 2 or later. Create a virtual machine that uses hardware version 19 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit), version 1809 or later releases
- Windows Server 2019 (64 bit) or later releases

Ensure that you install the latest patches for Windows 10, version 1809, and Windows Server 2019, before activating VBS.

For more information about activating VBS on virtual machines on AMD platforms, see the VMware KB article at <https://kb.vmware.com/s/article/89880>.

Configuring VBS is a process that involves first activating VBS in the virtual machine then activating VBS in the Windows guest OS.

1. Connect to vCenter Server by using the vSphere Client.
2. Select an object in the inventory that is a valid parent object of a virtual machine, for example, an ESXi host or a cluster.
3. Right-click the object, select **New Virtual Machine**, and follow the prompts to create a virtual machine.

Option	Action
Select a creation type	Create a virtual machine.
Select a name and folder	Specify a name and target location.
Select a compute resource	Specify an object for which you have privileges to create virtual machines.
Select storage	In the VM storage policy, select the storage policy. Select a compatible datastore.
Select compatibility	Intel CPU: Ensure that ESXi 6.7 and later is selected. AMD CPU: Ensure that ESXi 7.0 U2 and later is selected.
Select a guest OS	Select the Windows guest operating system option that best corresponds to operating system release. Select the Enable Windows Virtualization Based Security check box.
Customize hardware	Customize the hardware, for example, by changing disk size or CPU.
Ready to complete	Review the information and click Finish .

The Virtual Machine Details tile under the **Summary** tab displays "Virtualization Based Security - Enable".

See [Activate Virtualization-based Security on the Guest Operating System](#).

Activate Virtualization-based Security on an Existing Virtual Machine

You can activate Microsoft virtualization-based security (VBS) on existing virtual machines for supported Windows guest operating systems.

Using Intel CPUs for VBS requires vSphere 6.7 or later. The virtual machine must have been created using hardware version 14 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit) or later releases
- Windows Server 2016 (64 bit) or later releases

Using AMD CPUs for VBS requires vSphere 7.0 Update 2 or later. The virtual machine must have been created using hardware version 19 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit), version 1809 or later releases
- Windows Server 2019 (64 bit) or later releases

Ensure that you install the latest patches for Windows 10, version 1809, and Windows Server 2019, before activating VBS.

For more information about activating VBS on virtual machines on AMD platforms, see the VMware KB article at <https://kb.vmware.com/s/article/89880>.

Configuring VBS is a process that involves first activating VBS in the virtual machine then activating VBS in the guest OS.

Note: New virtual machines configured for Windows 10, Windows Server 2016, and Windows Server 2019 on hardware versions less than version 14 are created using Legacy BIOS by default. If you change the virtual machine's firmware type from Legacy BIOS to UEFI, you must reinstall the guest operating system.

1. In the vSphere Client, browse to the virtual machine.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click the **VM Options** tab.
4. Select the **Enable** check box for Virtualization Based Security.
5. Click **OK**.

The Virtual Machine Details tile under the **Summary** tab displays "Virtualization Based Security - Enable".

See [Activate Virtualization-based Security on the Guest Operating System](#).

Activate Virtualization-based Security on the Guest Operating System

You can activate Microsoft virtualization-based security (VBS) for supported Windows guest operating systems.

- Ensure that virtualization-based security has been activated on the virtual machine.

You activate VBS from within the Windows Guest OS. Windows configures and enforces VBS through a Group Policy Object (GPO). The GPO gives you the ability to turn off and on the various services, such as Secure Boot, Device Guard, and Credential Guard, that VBS offers. Certain Windows versions also require you to perform the additional step of enabling the Hyper-V platform.

See Microsoft's documentation about deploying Device Guard to activate virtualization-based security for details.

1. In Microsoft Windows, edit the group policy to turn on VBS and choose other VBS-related security options.
2. Optional: For Microsoft Windows versions less than Redstone 4, in the Windows Features control panel, enable the Hyper-V platform.
3. Reboot the guest operating system.

Deactivate Virtualization-based Security

If you no longer use virtualization-based security (VBS) with a virtual machine, you can deactivate VBS. When you deactivate VBS for the virtual machine, the Windows VBS options remain unchanged but might induce performance issues. Before deactivating VBS on the virtual machine, deactivate VBS options within Windows.

Ensure that the virtual machine is powered off.

1. In the vSphere Client, browse to the virtual machine that is using VBS.
See [Identify VBS-Enabled Virtual Machines](#) for help in locating virtual machines that use VBS.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click **VM Options**.
4. Deselect the **Enable** check box for Virtualization Based Security.
A message reminds you to deactivate VBS in the guest OS.
5. Click **OK**.
6. Verify that the virtual machine's **Summary** tab no longer displays "Virtualization Based Security - Enable" in the Guest OS description.

Identify VBS-Enabled Virtual Machines

You can identify which of your virtual machines have VBS enabled, for reporting and compliance purposes.

1. Connect to vCenter Server by using the vSphere Client.
2. Select a vCenter Server instance, a data center, or a host in the inventory.
3. Click the **VMs** tab and click **Virtual Machines**.
4. To show the **VBS** column, click **Manage Columns** and select the **VBS** check box.
5. Scan for "Present" in the **VBS** column.

Securing vSphere Networking

Securing vSphere networking is an essential part of protecting your environment. You secure different vSphere components in different ways. See the *vSphere Networking* documentation for detailed information about networking in the vSphere environment.

Network security in the vSphere environment shares many characteristics of securing a physical network environment, but also includes some characteristics that apply only to virtual machines.

Using Firewalls

Add firewall protection to your virtual network by installing and configuring host-based firewalls on some or all its virtual machines.

For efficiency, you can set up private virtual machine Ethernet networks or virtual networks. With virtual networks, you install a host-based firewall on a virtual machine at the head of the virtual network. This firewall serves as a protective buffer between the physical network adapter and the remaining virtual machines in the virtual network.

Host-based firewalls can slow performance. Balance your security needs against performance goals before you install host-based firewalls on virtual machines elsewhere in the virtual network.

See [Securing the Network with Firewalls](#).

Using Network Segmentation

Keep different virtual machine zones within a host on different network segments. If you isolate each virtual machine zone on its own network segment, you minimize the risk of data leakage from one zone to the next. Segmentation prevents various threats, including Address Resolution Protocol (ARP) spoofing. With ARP spoofing, an attacker manipulates the ARP table to remap MAC and IP addresses, and gains access to network traffic to and from a host. Attackers use ARP spoofing to generate man in the middle (MITM) attacks, perform denial of service (DoS) attacks, hijack the target system, and otherwise disrupt the virtual network.

Planning segmentation carefully lowers the chances of packet transmissions between virtual machine zones. Segmentation therefore prevents sniffing attacks that require sending network traffic to the victim. Also, an attacker cannot use a nonsecure service in one virtual machine zone to access other virtual machine zones in the host. You can implement segmentation by using one of two approaches.

- Use separate physical network adapters for virtual machine zones to ensure that the zones are isolated. Maintaining separate physical network adapters for virtual machine zones is probably the most secure method. After the initial segment creation. This approach is less prone to misconfiguration.
- Set up virtual local area networks (VLANs) to help safeguard your network. VLANs provide almost all the security benefits inherent in implementing physically separate networks without the hardware overhead. VLANs can save you the cost of deploying and maintaining additional devices, cabling, and so on. See [Securing Virtual Machines with VLANs](#).

Preventing Unauthorized Access to Virtual Machines

Requirements for securing virtual machines are often the same as requirements for securing physical machines.

- If a virtual machine network is connected to a physical network, it can be subject to breaches like a network that consists of physical machines.
- Even if you do not connect a virtual machine to the physical network, the virtual machine can be attacked by other virtual machines.

Virtual machines are isolated from each other. One virtual machine cannot read or write another virtual machine's memory, access its data, use its applications, and so forth. However, within the network, any virtual machine or group of virtual machines can still be the target of unauthorized access from other virtual machines. Protect your virtual machines from such unauthorized access.

For additional information about protecting virtual machines, see the NIST document titled "Secure Virtual Network Configuration for Virtual Machine (VM) Protection" at:

<https://csrc.nist.gov/publications/detail/sp/800-125b/final>

Securing the Network with Firewalls

Security administrators use firewalls to safeguard the network or selected components in the network from intrusion.

Firewalls control access to devices within their perimeter by closing all ports except for ports that the administrator explicitly or implicitly designates as authorized. The ports that administrators open allow traffic between devices on different sides of the firewall.

Important: The ESXi firewall in ESXi 5.5 and later does not allow per-network filtering of vMotion traffic. Therefore, you must install rules on your external firewall to ensure that no incoming connections can be made to the vMotion socket.

In a virtual machine environment, you can plan the layout for firewalls between components.

- Firewalls between physical machines such as vCenter Server systems and ESXi hosts.
- Firewalls between one virtual machine and another, for example, between a virtual machine acting as an external Web server and a virtual machine connected to your company's internal network.
- Firewalls between a physical machine and a virtual machine, such as when you place a firewall between a physical network adapter card and a virtual machine.

How you use firewalls in your ESXi configuration is based on how you plan to use the network and how secure any given component has to be. For example, if you create a virtual network where each virtual machine is dedicated to running a different benchmark test suite for the same department, the risk of unwanted access from one virtual machine to the next is minimal. Therefore, a configuration where firewalls are present between the virtual machines is not necessary. However, to prevent interruption of a test run from an outside host, you can configure a firewall at the entry point of the virtual network to protect the entire set of virtual machines.

For the list of all supported ports and protocols in VMware products, including vSphere and , see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com/>. You can search ports by VMware product, create a customized list of ports, and print or save port lists.

Firewalls for Configurations with vCenter Server

If you access ESXi hosts through vCenter Server, you typically protect vCenter Server using a firewall.

Firewalls must be present at entry points. A firewall might lie between the clients and vCenter Server or vCenter Server and the clients can both be behind a firewall.

For the list of all supported ports and protocols in VMware products, including vSphere and , see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com/>. You can search ports by VMware product, create a customized list of ports, and print or save port lists.

Networks configured with vCenter Server can receive communications through the vSphere Client, other UI clients, or clients that use the vSphere API. During normal operation, vCenter Server listens for data from its managed hosts and clients on designated ports. vCenter Server also assumes that its managed hosts listen for data from vCenter Server on designated ports. If a firewall is present between any of these elements, you must ensure that the firewall has open ports to support data transfer.

You might also include firewalls at other access points in the network, depending on the network usage and on the level of security that clients require. Select the locations for your firewalls based on the security risks for your network configuration. The following firewall locations are commonly used.

- Between the vSphere Client or a third-party network-management client and vCenter Server.
- If your users access virtual machines through a Web browser, between the Web browser and the ESXi host.
- If your users access virtual machines through the vSphere Client, between the vSphere Client and the ESXi host. This connection is in addition to the connection between the vSphere Client and vCenter Server, and it requires a different port.
- Between vCenter Server and the ESXi hosts.
- Between the ESXi hosts in your network. Although traffic between hosts is usually considered trusted, you can add firewalls between them if you are concerned about security breaches from machine to machine.
If you add firewalls between ESXi hosts and plan to migrate virtual machines between them, open ports in any firewall that divides the source host from the target hosts.
- Between the ESXi hosts and network storage such as NFS or iSCSI storage. These ports are not specific to VMware. Configure them according to the specifications for your network.

Connecting to vCenter Server Through a Firewall

Open TCP port 443 in the firewall to enable vCenter Server to receive data.

By default vCenter Server uses TCP port 443 to listen for data from its clients. If you have a firewall between vCenter Server and its clients, you must configure a connection through which vCenter Server can receive data from the clients. Firewall configuration depends on what is used at your site, ask your local firewall system administrator for information.

Connecting ESXi Hosts Through Firewalls

If you have a firewall between your ESXi hosts and vCenter Server, ensure that the managed hosts can receive data.

To configure a connection for receiving data, open ports for traffic from services such as vSphere High Availability, vMotion, and vSphere Fault Tolerance. See [Configuring the ESXi Firewall](#) for a discussion of configuration files, vSphere Client access, and firewall commands. For a list of ports, see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com>.

Firewalls for Configurations Without vCenter Server

If your environment does not include vCenter Server, clients can connect directly to the ESXi network.

You can connect to a standalone ESXi host in several ways.

- VMware Host Client
- ESXCLI interface
- vSphere Web Services SDK or vSphere Automation SDKs
- Third-party clients

The firewall requirements for standalone hosts are similar to requirements when a vCenter Server is present.

- Use a firewall to protect your ESXi layer or, depending on your configuration, your clients, and the ESXi layer. This firewall provides basic protection for your network.
- Licensing in this type of configuration is part of the ESXi package that you install on each of the hosts. Because licensing is resident to ESXi, a separate License Server with a firewall is not required.

You can configure firewall ports using ESXCLI or using the VMware Host Client. See *vSphere Single Host Management - VMware Host Client*.

Connecting to the Virtual Machine Console Through a Firewall

Certain ports must be open for user and administrator communication with the virtual machine console. Which ports must be open depends on the type of virtual machine console, and on whether you connect through vCenter Server with the vSphere Client or directly to the ESXi host from the VMware Host Client.

For more information about ports, purpose, and classification (incoming, outgoing, or bidirectional), see the VMware Ports and Protocols Tool™ at <https://ports.vmware.com>.

Connecting to a Browser-Based Virtual Machine Console Through the vSphere Client

When you are connecting with the vSphere Client, you always connect to the vCenter Server system that manages the ESXi host, and access the virtual machine console from there.

If you are using the vSphere Client and connecting to a browser-based virtual machine console, the following access must be possible:

- The firewall must allow vSphere Client to access vCenter Server on port 443.
- The firewall must allow vCenter Server to access the ESXi host on port 902.

Connecting to a VMware Remote Console Through the vSphere Client

If you are using the vSphere Client and connecting to a VMware Remote Console (VMRC), the following access must be possible:

- The firewall must allow the vSphere Client to access vCenter Server on port 443.
- The firewall must allow the VMRC to access vCenter Server on port 443 and to access the ESXi host on port 902 for VMRC versions before 11.0, and port 443 for VMRC version 11.0 and greater. For more information about VMRC version 11.0 and ESXi port requirements, see the VMware knowledge base article at <https://kb.vmware.com/s/article/76672>.

Connecting to ESXi Hosts Directly with the VMware Host Client

You can use the VMware Host Client virtual machine console if you connect directly to an ESXi host.

Note: Do not use the VMware Host Client to connect directly to hosts that are managed by a vCenter Server system. If you make changes to such hosts from the VMware Host Client, instability in your environment results.

The firewall must allow access to the ESXi host on ports 443 and 902.

The VMware Host Client uses port 902 to provide a connection for guest operating system MKS activities on virtual machines. It is through this port that users interact with the guest operating systems and applications of the virtual machine. VMware does not support configuring a different port for this function.

Secure the Physical Switch on ESXi Hosts

Secure the physical switch on each ESXi host to prevent attackers from gaining access to the host and its virtual machines.

For best protection of your hosts, ensure that physical switch ports are configured with spanning tree deactivated and ensure that the non-negotiate option is configured for trunk links between external physical switches and virtual switches in Virtual Switch Tagging (VST) mode.

1. Log in to the physical switch and ensure that spanning tree protocol is deactivated or that Port Fast is configured for all physical switch ports that are connected to ESXi hosts.
2. For virtual machines that perform bridging or routing, check periodically that the first upstream physical switch port is configured with BPDU Guard and Port Fast deactivated and with spanning tree protocol activated.
To prevent the physical switch from potential Denial of Service (DoS) attacks, you can turn on the guest BPDU filter on the ESXi hosts.
3. Log in to the physical switch and ensure that Dynamic Trunking Protocol (DTP) is not activated on the physical switch ports that are connected to the ESXi hosts.
4. Routinely check physical switch ports to ensure that they are properly configured as trunk ports if connected to virtual switch VLAN trunking ports.

Securing Standard Switch Ports with Security Policies

The VMkernel port group or virtual machine port group on a standard switch has a configurable security policy. The security policy determines how strongly you enforce protection against impersonation and interception attacks on VMs.

Just like physical network adapters, virtual machine network adapters can impersonate another VM. Impersonation is a security risk.

- A VM can send frames that appear to be from a different machine so that it can receive network frames that are intended for that machine.
- A virtual machine network adapter can be configured so that it receives frames targeted for other machines

When you add a VMkernel port group or virtual machine port group to a standard switch, ESXi configures a security policy for the ports in the group. You can use this security policy to ensure that the host prevents the guest operating systems of its VMs from impersonating other machines on the network. The guest operating system that might attempt impersonation does not detect that the impersonation was prevented.

The security policy determines how strongly you enforce protection against impersonation and interception attacks on VMs. To correctly use the settings in the security profile, see the Security Policy section in the *vSphere Networking* publication. This section explains:

- How VM network adapters control transmissions.
- How attacks are staged at this level

Securing vSphere Standard Switches

You can secure standard switch traffic against Layer 2 attacks by restricting some of the MAC address modes of the VM network adapters.

Each VM network adapter has an initial MAC address and an effective MAC address.

Initial MAC Address

The initial MAC address is assigned when the adapter is created. Although the initial MAC address can be reconfigured from outside the guest operating system, it cannot be changed by the guest operating system.

Effective MAC Address

Each adapter has an effective MAC address that filters out incoming network traffic with a destination MAC address that is different from the effective MAC address. The guest operating system is responsible for setting the effective MAC address and typically matches the effective MAC address to the initial MAC address.

What Happens When You Create a Virtual Machine Network Adapter

Upon creating a virtual machine network adapter, the effective MAC address and initial MAC address are the same. The guest operating system can alter the effective MAC address to another value at any time. If an operating system changes the effective MAC address, its network adapter receives network traffic that is destined for the new MAC address.

When sending packets through a network adapter, the guest operating system typically places its own adapter effective MAC address in the source MAC address field of the Ethernet frames. It places the MAC address for the receiving network adapter in the destination MAC address field. The receiving adapter accepts packets only if the destination MAC address in the packet matches its own effective MAC address.

An operating system can send frames with an impersonated source MAC address. An operating system can therefore impersonate a network adapter that the receiving network authorizes, and stage malicious attacks on the devices in a network.

Using Security Policies to Protect Ports and Groups

Protect virtual traffic against impersonation and interception Layer 2 attacks by configuring a security policy on port groups or ports.

The security policy on distributed port groups and ports includes the following options:

- MAC address changes (see [MAC Address Changes](#))
- Promiscuous mode (see [Promiscuous Mode Operation](#))
- Forged transmits (see [Forged Transmits](#))

You can view and change the default settings by selecting the virtual switch associated with the host from the vSphere Client. See *vSphere Networking* documentation.

MAC Address Changes

The security policy of a virtual switch includes a **MAC address changes** option. This option allows virtual machines to receive frames with a Mac Address that is different from the one configured in the VMX.

When the **Mac address changes** option is set to **Accept**, ESXi accepts requests to change the effective MAC address of a virtual machine to a different address than the initial MAC address.

When the **Mac address changes** option is set to **Reject**, ESXi does not honor requests to change the effective MAC address of a virtual machine to a different address than the initial MAC address. This setting protects the host against MAC impersonation. The port that the virtual machine adapter used to send the request is deactivated and the virtual machine adapter does not receive any more frames until the effective MAC address matches the initial MAC address. The guest operating system does not detect that the MAC address change request was not honored.

Note: The iSCSI initiator relies on being able to get MAC address changes from certain types of storage. If you are using ESXi iSCSI with iSCSI storage, set the **MAC address changes** option to **Accept**.

In some situations, you can have a legitimate need for more than one adapter to have the same MAC address on a network, for example, if you are using Microsoft Network Load Balancing in unicast mode. When Microsoft Network Load Balancing is used in the standard multicast mode, adapters do not share MAC addresses.

Note: Starting in vSphere 7.0, the defaults for **Forged transmits** and **MAC address changes** have been changed to Reject instead of Accept. Contact your storage vendor to validate.

Forged Transmits

The **Forged transmits** option affects traffic that is transmitted from a virtual machine.

When the **Forged transmits** option is set to **Accept**, ESXi does not compare source and effective MAC addresses.

To protect against MAC impersonation, you can set the **Forged transmits** option to **Reject**. If you do, the host compares the source MAC address being transmitted by the guest operating system with the effective MAC address for its virtual machine adapter to see if they match. If the addresses do not match, the ESXi host drops the packet.

The guest operating system does not detect that its virtual machine adapter cannot send packets by using the impersonated MAC address. The ESXi host intercepts any packets with impersonated addresses before they are delivered, and the guest operating system might assume that the packets are dropped.

Note: Starting in vSphere 7.0, the defaults for **Forged transmits** and **MAC address changes** have been changed to Reject instead of Accept.

Promiscuous Mode Operation

Promiscuous mode eliminates any reception filtering that the virtual machine adapter performs so that the guest operating system receives all traffic observed on the wire. By default, the virtual machine adapter cannot operate in promiscuous mode.

Although promiscuous mode can be useful for tracking network activity, it is an insecure mode of operation, because any adapter in promiscuous mode has access to the packets even if some of the packets are received only by a particular network adapter. This means that an administrator or root user within a virtual machine can potentially view traffic destined for other guest or host operating systems.

See the topic on configuring the security policy for a vSphere Standard Switch or Standard Port Group in the *vSphere Networking* documentation for information about configuring the virtual machine adapter for promiscuous mode.

Note: In some situations, you might have a legitimate reason to configure a standard or a distributed virtual switch to operate in promiscuous mode, for example, if you are running network intrusion detection software or a packet sniffer.

Standard Switch Protection and VLANs

VMware standard switches provide safeguards against certain threats to VLAN security. Because of the way that standard switches are designed, they protect VLANs against a variety of attacks, many of which involve VLAN hopping.

Having this protection does not guarantee that your virtual machine configuration is invulnerable to other types of attacks. For example, standard switches do not protect the physical network against these attacks; they protect only the virtual network.

Standard switches and VLANs can protect against the following types of attacks.

Because new security threats develop over time, do not consider this an exhaustive list of attacks. Regularly check VMware security resources on the Web to learn about security, recent security alerts, and VMware security tactics.

MAC Flooding

MAC flooding floods a switch with packets that contain MAC addresses tagged as having come from different sources. Many switches use a content-addressable memory table to learn and store the source address for each packet. When the table is full, the switch can enter a fully open state in which every incoming packet is broadcast on all ports, letting the attacker see all of the traffic on the switch. This state might result in packet leakage across VLANs.

Although VMware standard switches store a MAC address table, they do not get the MAC addresses from observable traffic and are not vulnerable to this type of attack.

802.1q and ISL Tagging Attacks

802.1q and ISL tagging attacks force a switch to redirect frames from one VLAN to another by tricking the switch into acting as a trunk and broadcasting the traffic to other VLANs.

VMware standard switches do not perform the dynamic trunking required for this type of attack and, therefore, are not vulnerable.

Double-encapsulation Attacks

Double-encapsulation attacks occur when an attacker creates a double-encapsulated packet in which the VLAN identifier in the inner tag is different from the VLAN identifier in the outer tag. For backward compatibility, native VLANs strip the outer tag from transmitted packets unless configured to do otherwise. When a native VLAN switch strips the outer tag, only the inner tag is left, and that inner tag routes the packet to a different VLAN than the one identified in the now-missing outer tag.

VMware standard switches drop any double-encapsulated frames that a virtual machine attempts to send on a port configured for a specific VLAN. Therefore, they are not vulnerable to this type of attack.

Multicast Brute-force Attacks

Involve sending large numbers of multicast frames to a known VLAN almost simultaneously to overload the switch so that it mistakenly allows some of the frames to broadcast to other VLANs.

VMware standard switches do not allow frames to leave their correct broadcast domain (VLAN) and are not vulnerable to this type of attack.

Spanning-tree Attacks

Spanning-tree attacks target Spanning-Tree Protocol (STP), which is used to control bridging between parts of the LAN. The attacker sends Bridge Protocol Data Unit (BPDU) packets that attempt to change the network topology, establishing themselves as the root bridge. As the root bridge, the attacker can sniff the contents of transmitted frames.

VMware standard switches do not support STP and are not vulnerable to this type of attack.

Random Frame Attacks

Random frame attacks involve sending large numbers of packets in which the source and destination addresses stay the same, but in which fields are randomly changed in length, type, or content. The goal of this attack is to force packets to be mistakenly rerouted to a different VLAN.

VMware standard switches are not vulnerable to this type of attack.

Secure vSphere Distributed Switches and Distributed Port Groups

Administrators have several options for securing vSphere Distributed Switches in their vSphere environment.

The same rules apply for VLANs in a vSphere Distributed Switch as they do in a standard switch. For more information, see [Standard Switch Protection and VLANs](#).

1. For distributed port groups with static binding, deactivate the Auto Expand feature.

Auto Expand is activated by default.

To deactivate Auto Expand, configure the `autoExpand` property under the distributed port group with the vSphere Web Services SDK or with a command-line interface. See the *vSphere Web Services SDK* documentation.

2. Ensure that all private VLAN IDs of any vSphere Distributed Switch are fully documented.
3. If you are using VLAN tagging on a dvPortgroup, VLAN IDs must correspond to the IDs on external VLAN-aware upstream switches. If VLAN IDs are not tracked correctly, mistaken reuse of IDs might allow unintended traffic. Similarly, wrong or missing VLAN IDs might lead to traffic not passing between physical and virtual machines.
4. Ensure that no unused ports exist on a virtual port group associated with a vSphere Distributed Switch.
5. Label all vSphere Distributed Switches.
vSphere Distributed Switches associated with an ESXi host require a text box for the name of the switch. This label serves as a functional descriptor for the switch, just like the host name associated with a physical switch. The label on the vSphere Distributed Switch indicates the function or the IP subnet of the switch. For example, you can label the switch as internal to indicate that it is only for internal networking on a virtual machine's private virtual switch. No traffic goes over physical network adapters.
6. Deactivate network health check for your vSphere Distributed Switches if you are not actively using it.
Network health check is deactivated by default. Once activated, the health check packets contain information about the host, switch, and port that an attacker can potentially use. Use network health check only for troubleshooting, and turn it off when troubleshooting is finished.

7. Protect virtual traffic against impersonation and interception Layer 2 attacks by configuring a security policy on port groups or ports.

The security policy on distributed port groups and ports includes the following options:

- MAC address changes (see [MAC Address Changes](#))
- Promiscuous mode (see [Promiscuous Mode Operation](#))
- Forged transmits (see [Forged Transmits](#))


You can view and change the current settings by selecting **Manage Distributed Port Groups** from the right-button menu of the distributed switch and selecting **Security** in the wizard. See the *vSphere Networking* documentation.

Securing Virtual Machines with VLANs

The network can be one of the most vulnerable parts of any system. Your virtual machine network requires as much protection as your physical network. Using VLANs can improve networking security in your environment.

VLANs are an IEEE standard networking scheme with specific tagging methods that allow routing of packets to only those ports that are part of the VLAN. When properly configured, VLANs provide a dependable means for you to protect a set of virtual machines from accidental or malicious intrusions.

VLANs let you segment a physical network so that two machines in the network are unable to transmit packets back and forth unless they are part of the same VLAN. For example, accounting records and transactions are among a company's most sensitive internal information. In a company whose sales, shipping, and accounting employees all use virtual machines in the same physical network, you might protect the virtual machines for the accounting department by setting up VLANs.

Figure 123: Sample VLAN Layout

In this configuration, all employees in the accounting department use virtual machines in VLAN A and the employees in sales use virtual machines in VLAN B.

The router forwards packets containing accounting data to the switches. These packets are tagged for distribution to VLAN A only. Therefore, the data is confined to Broadcast Domain A and cannot be routed to Broadcast Domain B unless the router is configured to do so.

This VLAN configuration prevents the sales force from intercepting packets destined for the accounting department. It also prevents the accounting department from receiving packets intended for the sales group. The virtual machines serviced by a single virtual switch can be in different VLANs.

Security Considerations for VLANs

The way you set up VLANs to secure parts of a network depends on factors such as the guest operating system and the way your network equipment is configured.

ESXi features a complete IEEE 802.1q-compliant VLAN implementation. VMware cannot make specific recommendations on how to set up VLANs, but there are factors to consider when using a VLAN deployment as part of your security enforcement policy.

Secure VLANs

Administrators have several options for securing the VLANs in their vSphere environment.

1. Ensure that port groups are not configured to VLAN values that are reserved by upstream physical switches
Do not set VLAN IDs to values reserved for the physical switch.
2. Ensure that port groups are not configured to VLAN 4095 unless you are using for Virtual Guest Tagging (VGT).
Three types of VLAN tagging exist in vSphere:
 - External Switch Tagging (EST)
 - Virtual Switch Tagging (VST) - The virtual switch tags with the configured VLAN ID the traffic that is incoming to the attached virtual machines and removes the VLAN tag from the traffic that is leaving them. To set up VST mode, assign a VLAN ID between 1 and 4094.
 - Virtual Guest Tagging (VGT) - Virtual machines handle VLAN traffic. To activate VGT mode, set the VLAN ID to 4095. On a distributed switch, you can also allow virtual machine traffic based on its VLAN by using the **VLAN Trunking** option.


On a standard switch you can configure VLAN networking mode at switch or port group level, and on a distributed switch at distributed port group or port level.

3. Ensure that all VLANs on each virtual switch are fully documented and that each virtual switch has all required VLANs and only required VLANs.

Creating Multiple Networks Within a Single ESXi Host

The ESXi system is designed so that you can connect some groups of virtual machines to the internal network, others to the external network, and still others to both—all on the same host. This capability is an outgrowth of basic virtual machine isolation coupled with a well-planned use of virtual networking features.

Figure 124: External Networks, Internal Networks, and a DMZ Configured on a Single ESXi Host



In the figure, the system administrator configured a host into three distinct virtual machine zones: FTP server, internal virtual machines, and DMZ. Each zone serves a unique function.

FTP Server Zone

Virtual Machine 1 is configured with FTP software and acts as a holding area for data sent to and from outside resources such as forms and collateral localized by a vendor.

This virtual machine is associated with an external network only. It has its own virtual switch and physical network adapter that connect it to External Network 1. This network is dedicated to servers that the company uses to receive data from outside sources. For example, the company uses External Network 1 to receive FTP traffic from vendors and allow vendors access to data stored on externally available servers through FTP. In addition to servicing Virtual Machine 1, External Network 1 services FTP servers configured on different ESXi hosts throughout the site.

Because Virtual Machine 1 does not share a virtual switch or physical network adapter with any virtual machines in the host, the other resident virtual machines cannot transmit packets to or receive packets from the Virtual Machine 1 network. This restriction prevents sniffing attacks, which require sending network traffic to the victim. More importantly, an attacker cannot use the natural vulnerability of FTP to access any of the host's other virtual machines.

Internal Network Zone

Virtual Machines 2 through 5 are reserved for internal use. These virtual machines process and store company-private data such as medical records, legal settlements, and fraud investigations. As a result, the system administrators must ensure the highest level of protection for these virtual machines.

These virtual machines connect to Internal Network 2 through their own virtual switch and network adapter. Internal Network 2 is reserved for internal use by personnel such as claims processors, in-house lawyers, or adjusters.

Virtual Machines 2 through 5 can communicate with one another through the virtual switch and with internal virtual machines elsewhere on Internal Network 2 through the physical network adapter. They cannot communicate with externally facing machines. As with the FTP server, these virtual machines cannot send packets to or receive packets from the other virtual machines' networks. Similarly, the host's other virtual machines cannot send packets to or receive packets from Virtual Machines 2 through 5.

DMZ Zone

Virtual Machines 6 through 8 are configured as a DMZ that the marketing group uses to publish the company's external website.

This group of virtual machines is associated with External Network 2 and Internal Network 1. The company uses External Network 2 to support the Web servers that use the marketing and financial department to host the corporate website and other Web facilities that it hosts to outside users. Internal Network 1 is the conduit that the marketing department uses to publish its content to the corporate Web site, post downloads, and maintain services like user forums.

Because these networks are separate from External Network 1 and Internal Network 2, and the virtual machines have no shared points of contact (switches or adapters), there is no risk of attack to or from the FTP server or the internal virtual machine group.

Advantages of Using Virtual Machine Zones

By capitalizing on virtual machine isolation, correctly configuring virtual switches, and maintaining network separation, you can house all three virtual machine zones on the same ESXi host and be confident that there will be no data or resource breaches.

The company enforces isolation among the virtual machine groups by using multiple internal and external networks and making sure that the virtual switches and physical network adapters for each group are separate from those of other groups.

Because none of the virtual switches straddle virtual machine zones, you succeed in eliminating the risk of packet leakage from one zone to another. A virtual switch, by design, cannot leak packets directly to another virtual switch. The only way for packets to travel from one virtual switch to another is under the following circumstances:

- The virtual switches are connected to the same physical LAN.
- The virtual switches connect to a common virtual machine, which might be used to transmit packets.

Neither of these conditions occur in the sample configuration. If you want to verify that no common virtual switch paths exist, you can check for possible shared points of contact by reviewing the network switch layout in the vSphere Client.

To safeguard the resources of virtual machines, configure a resource reservation and a limit for each virtual machine, which lowers the risk of DoS and DDoS attacks. You can further protect the ESXi host and virtual machines by installing software firewalls at the front and back ends of the DMZ. Finally, ensure that the host is behind a physical firewall, and configure the networked storage resources so that each has its own virtual switch.

Using Internet Protocol Security on ESXi Hosts

Internet Protocol Security (IPsec) secures IP communications coming from and arriving at a host. ESXi hosts support IPsec using IPv6.

When you set up IPsec on an ESXi host, you enable authentication and encryption of incoming and outgoing packets. When and how IP traffic is encrypted depends on how you set up the system's security associations and security policies.

A security association determines how the system encrypts traffic. When you create a security association, you specify the source and destination, encryption parameters, and a name for the security association.

A security policy determines when the system should encrypt traffic. The security policy includes source and destination information, the protocol and direction of traffic to be encrypted, the mode (transport or tunnel) and the security association to use.

List Available Security Associations on ESXi Hosts

ESXi can provide a list of all security associations available for use by security policies. The list includes both user created security associations and any security associations the VMkernel installed using Internet Key Exchange.

You can get a list of available security associations using the `esxcli` command.

At the command prompt, enter the command `esxcli network ip ipsec sa list`.

ESXi displays a list of all available security associations.

Add an IPsec Security Association to an ESXi Host

Add a security association to specify encryption parameters for associated IP traffic.

You can add a security association using the `esxcli` command.

At the command prompt, enter the command `esxcli network ip ipsec sa add` with one or more of the following options.

Option	Description
<code>--sa-source= <i>source address</i></code>	Required. Specify the source address.
<code>--sa-destination= <i>destination address</i></code>	Required. Specify the destination address.
<code>--sa-mode= <i>mode</i></code>	Required. Specify the mode, either <code>transport</code> or <code>tunnel</code> .
<code>--sa-spi= <i>security parameter index</i></code>	Required. Specify the security parameter index. The security parameter index identifies the security association to the host. It must be a hexadecimal with a 0x prefix. Each security association

Option	Description
	you create must have a unique combination of protocol and security parameter index.
--encryption-algorithm= <i>encryption algorithm</i>	Required. Specify the encryption algorithm using one of the following parameters. 3des-cbc - aes128-cbc - null (provides no encryption)
--encryption-key= <i>encryption key</i>	Required when you specify an encryption algorithm. Specify the encryption key. You can enter keys as ASCII text or as a hexadecimal with a 0x prefix.
--integrity-algorithm= <i>authentication algorithm</i>	Required. Specify the authentication algorithm, either hmac-sha1 or hmac-sha2-256.
--integrity-key= <i>authentication key</i>	Required. Specify the authentication key. You can enter keys as ASCII text or as a hexadecimal with a 0x prefix.
--sa-name= <i>name</i>	Required. Provide a name for the security association.

New Security Association Command

The following example contains extra line breaks for readability.

```
esxcli network ip ipsec sa add
--sa-source 3ffe:501:ffff:0::a
--sa-destination 3ffe:501:ffff:0001:0000:0000:0000:0001
--sa-mode transport
--sa-spi 0x1000
--encryption-algorithm 3des-cbc
--encryption-key 0x6970763672656164796c6f676f336465736362636f757432
--integrity-algorithm hmac-sha1
--integrity-key 0x6970763672656164796c6f67736861316f757432
--sa-name sal
```

Remove an IPsec Security Association from an ESXi Host

You can remove a security association using the ESXCLI command.

Verify that the security association you want to use is not currently in use. If you try to remove a security association that is in use, the removal operation fails.

At the command prompt, enter the command `esxcli network ip ipsec sa remove --sa-name security_association_name`.

List Available IPsec Security Policies on an ESXi Host

You can list available security policies using the ESXCLI command.

At the command prompt, enter the command `esxcli network ip ipsec sp list`.

The host displays a list of all available security policies.

Create an IPSec Security Policy on an ESXi Host

Create a security policy to determine when to use the authentication and encryption parameters set in a security association. You can add a security policy using the ESXCLI command.

Before creating a security policy, add a security association with the appropriate authentication and encryption parameters as described in [Add an IPsec Security Association to an ESXi Host](#).

At the command prompt, enter the command `esxcli network ip ipsec sp add` with one or more of the following options.

Option	Description
<code>--sp-source= source address</code>	Required. Specify the source IP address and prefix length.
<code>--sp-destination= destination address</code>	Required. Specify the destination address and prefix length.
<code>--source-port= port</code>	Required. Specify the source port. The source port must be a number between 0 and 65535.
<code>--destination-port= port</code>	Required. Specify the destination port. The source port must be a number between 0 and 65535.
<code>--upper-layer-protocol= protocol</code>	Specify the upper layer protocol using one of the following parameters. - tcp - udp - icmp6 - any
<code>--flow-direction= direction</code>	Specify the direction in which you want to monitor traffic using either <code>in</code> or <code>out</code> .
<code>--action= action</code>	Specify the action to take when traffic with the specified parameters is encountered using one of the following parameters. <code>none</code>: Take no action. <code>discard</code>: Do not allow data in or out. <code>ipsec</code>: Use the authentication and encryption information supplied in the security association to determine whether the data comes from a trusted source.
<code>--sp-mode= mode</code>	Specify the mode, either <code>tunnel</code> or <code>transport</code> .
<code>--sa-name=security association name</code>	Required. Provide the name of the security association for the security policy to use.
<code>--sp-name=name</code>	Required. Provide a name for the security policy.

New Security Policy Command

The following example includes extra line breaks for readability.

```
esxcli network ip ipsec add
--sp-source=2001:db8:1::/64
--sp-destination=2002:db8:1::/64
--source-port=23
--destination-port=25
--upper-layer-protocol=tcp
--flow-direction=out
--action=ipsec
--sp-mode=transport
--sa-name=sal
--sp-name=sp1
```

Remove an IPsec Security Policy from an ESXi Host

You can remove a security policy from the ESXi host using the ESXCLI command.

Verify that the security policy you want to use is not currently in use. If you try to remove a security policy that is in use, the removal operation fails.

At the command prompt, enter the command `esxcli network ip ipsec sp remove --sa-name security policy name`.

To remove all security policies, enter the command `esxcli network ip ipsec sp remove --remove-all`.

Ensure Proper SNMP Configuration on ESXi Hosts

If SNMP is not properly configured, monitoring information can be sent to a malicious host. The malicious host can then use this information to plan an attack.

ESXi includes an SNMP agent that can send notifications (traps and informs) and receive `GET`, `GETBULK`, and `GETNEXT` requests. SNMP is not activated by default. SNMP must be configured on each ESXi host. You can use ESXCLI, PowerCLI, or the vSphere Web Services SDK for configuration.

See the *vSphere Monitoring and Performance* documentation for detailed information about configuring SNMP, including SNMP v3. SNMP v3 provides stronger security than SNMP v1 or SNMP v2c, including key authentication and encryption. See *ESXCLI Reference* for more information about the `esxcli system snmp` command options.

1. To determine whether SNMP is used, run the following command.

```
esxcli system snmp get
```

2. To activate SNMP, run the following command.

```
esxcli system snmp set --enable true
```

3. To deactivate SNMP, run the following command.

```
esxcli system snmp set --enable false
```

vSphere Networking Security Best Practices

Following networking security best practices helps ensure the integrity of your vSphere deployment.

General vSphere Networking Security Recommendations

Following general network security recommendations is the first step in securing your vSphere networking environment. You can then move on to special areas, such as securing the network with firewalls or using IPsec.

Recommendations to Secure a vSphere Networking Environment

- Spanning Tree Protocol (STP) detects and prevents loops from forming in the network topology. VMware virtual switches prevent loops in other ways, but do not support STP directly. When network topology changes occur, some time is required (30–50 seconds) while the network relearns the topology. During that time, no traffic is allowed to pass. To avoid these problems, network vendors have created features for switch ports to continue forwarding traffic. For more information, see the VMware knowledge base article at <https://kb.vmware.com/s/article/1003804>. Consult your network vendor documentation for the proper network and networking hardware configurations.
- Ensure that Netflow traffic for a Distributed Virtual Switch is only sent to authorized collector IP addresses. Netflow exports are not encrypted and can contain information about the virtual network. This information increases the potential for sensitive information to be viewed and captured in transit by attackers. If Netflow export is required, verify that all Netflow target IP addresses are correct.
- Ensure that only authorized administrators have access to virtual networking components by using the role-based access controls. For example, give virtual machine administrators only access to port groups in which their virtual machines reside. Give network administrators access to all virtual networking components but no access to virtual

machines. Limiting access reduces the risk of misconfiguration, whether accidental or malicious, and enforces key security concepts of separation of duties and least privilege.

- Ensure that port groups are not configured to the value of the native VLAN. Physical switches are often configured with a native VLAN, and that native VLAN is often VLAN 1 by default. ESXi does not have a native VLAN. Frames with VLAN specified in the port group have a tag, but frames with VLAN not specified in the port group are not tagged. This can cause a problem because virtual machines that are tagged with a 1 end up belonging to native VLAN of the physical switch.

For example, frames on VLAN 1 from a Cisco physical switch are untagged because VLAN 1 is the native VLAN on that physical switch. However, frames from the ESXi host that are specified as VLAN 1 are tagged with a 1. As a result, traffic from the ESXi host that is destined for the native VLAN is not routed correctly because it is tagged with a 1 instead of being untagged. Traffic from the physical switch that is coming from the native VLAN is not visible because it is not tagged. If the ESXi virtual switch port group uses the native VLAN ID, traffic from virtual machines on that port is not visible to the native VLAN on the switch because the switch is expecting untagged traffic.

- Ensure that port groups are not configured to VLAN values reserved by upstream physical switches. Physical switches reserve certain VLAN IDs for internal purposes and often disallow traffic configured to these values. For example, Cisco Catalyst switches typically reserve VLANs 1001–1024 and 4094. Using a reserved VLAN might result in a denial of service on the network.
- Ensure that port groups are not configured to VLAN 4095 except for Virtual Guest Tagging (VGT). Setting a port group to VLAN 4095 activates VGT mode. In this mode, the virtual switch passes all network frames to the virtual machine without modifying the VLAN tags, leaving it to the virtual machine to deal with them.
- Restrict port-level configuration overrides on a distributed virtual switch. Port-level configuration overrides are deactivated by default. When overrides are activated, you can use different security settings for a virtual machine than the port-group level settings. Certain virtual machines require unique configurations, but monitoring is essential. If overrides are not monitored, anyone who gains access to a virtual machine with a less secure distributed virtual switch configuration might attempt to exploit that access.
- Ensure that distributed virtual switch port mirror traffic is sent only to authorized collector ports or VLANs. A vSphere Distributed Switch can mirror traffic from one port to another to allow packet capture devices to collect specific traffic flows. Port mirroring sends a copy of all specified traffic in unencrypted format. This mirrored traffic contains the full data in the packets captured and can result in total compromise of that data if misdirected. If port mirroring is required, verify that all port mirror destination VLAN, port, and uplink IDs are correct.

Labeling vSphere Networking Components

Identifying the different components of your vSphere networking architecture is critical and helps ensure that no errors are introduced as your network expands.

Follow these best practices:

- Ensure that port groups are configured with a clear network label. These labels serve as a functional descriptor for the port group and help you identify each port group's function as the network becomes more complex.
- Ensure that each vSphere Distributed Switch has a clear network label that indicates the function or IP subnet of the switch. This label serves as a functional descriptor for the switch, just as physical switches require a host name. For example, you can label the switch as internal to show that it is for internal networking. You cannot change the label for a standard virtual switch.

Document and Check the vSphere VLAN Environment

Check your VLAN environment regularly to avoid addressing problems. Fully document the VLAN environment and ensure that VLAN IDs are used only once. Your documentation can help with troubleshooting and is essential when you want to expand the environment.

1. Ensure that all vSwitch and VLANs IDs are fully documented
If you are using VLAN tagging on a virtual switch, the IDs must correspond to the IDs on external VLAN-aware upstream switches. If VLAN IDs are not tracked completely, mistaken reuse of IDs might allow for traffic between the wrong physical and virtual machines. Similarly, if VLAN IDs are wrong or missing, traffic between physical and virtual machines might be blocked where you want traffic to pass.
2. Ensure that VLAN IDs for all distributed virtual port groups (dvPortgroup instances) are fully documented.
If you are using VLAN tagging on a dvPortgroup the IDs must correspond to the IDs on external VLAN-aware upstream switches. If VLAN IDs are not tracked completely, mistaken reuse of IDs might allow for traffic between the wrong physical and virtual machines. Similarly, if VLAN IDs are wrong or missing, traffic between physical and virtual machines might be blocked where you want traffic to pass.
3. Ensure that private VLAN IDs for all distributed virtual switches are fully documented.
Private VLANs (PVLANS) for distributed virtual switches require primary and secondary VLAN IDs. These IDs must correspond to the IDs on external PVLAN-aware upstream switches. If VLAN IDs are not tracked completely, mistaken reuse of IDs might allow for traffic between the wrong physical and virtual machines. Similarly, if PVLAN IDs are wrong or missing, traffic between physical and virtual machines might be blocked where you want traffic to pass.
4. Verify that VLAN trunk links are connected only to physical switch ports that function as trunk links.
When connecting a virtual switch to a VLAN trunk port, you must properly configure both the virtual switch and the physical switch at the uplink port. If the physical switch is not properly configured, frames with the VLAN 802.1q header are forwarded to a switch that not expecting their arrival.

Adopting Network Isolation Practices in vSphere

Network isolation practices bolster network security in your vSphere environment.

Isolate the vSphere Management Network

The vSphere management network provides access to the vSphere management interface on each component. Services running on the management interface provide an opportunity for an attacker to gain privileged access to the systems. Remote attacks are likely to begin with gaining access to this network. If an attacker gains access to the management network, it provides the staging ground for further intrusion.

Strictly control access to management network by protecting it at the security level of the most secure VM running on an ESXi host or cluster. No matter how the management network is restricted, administrators must have access to this network to configure the ESXi hosts and vCenter Server system.

Place the vSphere management port group in a dedicated VLAN on a common standard switch. Production (VM) traffic can share the standard switch if the vSphere management port group's VLAN is not used by production VMs.

Check that the network segment is not routed, except to networks where other management-related entities are found. Routing a network segment might make sense for vSphere Replication. In particular, make sure that production VM traffic cannot be routed to this network.

Strictly control access to management functionality by using one of the following approaches.

- To access the management network in especially sensitive environments, configure a controlled gateway or other controlled method. For example, require that administrators connect to the management network through a VPN. Allow access to the management network only to trusted administrators.
- Configure bastion hosts that run management clients.

Isolate Storage Traffic

Ensure that IP-based storage traffic is isolated. IP-based storage includes iSCSI and NFS. Virtual machines might share virtual switches and VLANs with the IP-based storage configurations. This type of configuration might expose IP-based storage traffic to unauthorized virtual machine users.

IP-based storage frequently is not encrypted. Anyone with access to this network can view IP-based storage traffic. To restrict unauthorized users from viewing IP-based storage traffic, logically separate the IP-based storage network traffic from the production traffic. Configure the IP-based storage adapters on separate VLANs or network segments from the VMkernel management network to limit unauthorized users from viewing the traffic.

Isolate vMotion Traffic

vMotion migration information is transmitted in plain text. Anyone with access to the network over which this information flows can view it. Potential attackers can intercept vMotion traffic to obtain the memory contents of a VM. They might also stage an MITM attack in which the contents are modified during migration.

Separate vMotion traffic from production traffic on an isolated network. Set up the network to be nonroutable, that is, make sure that no layer-3 router is spanning this and other networks, to prevent outside access to the network.

Use a dedicated VLAN on a common standard switch for the vMotion port group. Production (VM) traffic can use the same standard switch if the vMotion port group's VLAN is not used by production virtual machines.

Isolate vSAN Traffic

When configuring your vSAN network, isolate vSAN traffic on its own Layer 2 network segment. You can perform this isolation by using dedicated switches or ports, or by using a VLAN.

Use Virtual Switches with the vSphere Network Appliance API Only If Required

Do not configure your host to send network information to a virtual machine unless you are using products that use the vSphere Network Appliance API (DvFilter). If the vSphere Network Appliance API is enabled, an attacker might attempt to connect a virtual machine to the filter. This connection might provide access to the network of other virtual machines on the host.

If you are using a product that uses this API, verify that the host is configured correctly. See the sections on DvFilter in the *Developing and Deploying vSphere Solutions, vServices, and ESX Agents* documentation. If your host is set up to use the API, make sure that the value of the `Net.DVFilterBindIpAddress` parameter matches the product that uses the API.

1. Browse to the host in the vSphere Client inventory.
2. Click **Configure**.
3. Under System, click **Advanced System Settings**.
4. Scroll down to `Net.DVFilterBindIpAddress` and verify that the parameter has an empty value.
The order of parameters is not strictly alphabetical. Enter `DVFilter` in the Filter text box to display all related parameters.
5. Verify the setting.
 - If you are not using DvFilter settings, make sure that the value is blank.
 - If you are using DvFilter settings, make sure that the value of the parameter is correct. The value must match the value that the product that uses the DvFilter is using.

Best Practices Involving Multiple vSphere Components

Some security best practices, such as setting up PTP or NTP in your environment, affect more than one vSphere component. Consider these recommendations when configuring your environment.

See [Securing ESXi Hosts](#) and [Securing Virtual Machines](#) for related information.

Synchronizing Clocks on the vSphere Network

Verify that all components on the vSphere network have their clocks synchronized. If the clocks on the physical machines in your vSphere network are not synchronized, SSL certificates and SAML tokens, which are time-sensitive, might not be recognized as valid in communications between network machines.

Unsynchronized clocks can result in authentication problems, which can cause the installation to fail or prevent the vCenter Server `vmware-vpxd` service from starting.

Time inconsistencies in vSphere can cause the first boot of a component in your environment to fail at different services depending on where in the environment time is not accurate and when the time is synchronized. Problems most commonly occur when the target ESXi host for the destination vCenter Server is not synchronized with NTP or PTP. Similarly, issues can arise if the destination vCenter Server migrates to an ESXi host set to a different time due to fully automated DRS.

To avoid time synchronization issues, ensure that the following is correct before installing, migrating, or upgrading a vCenter Server instance.

- The target ESXi host where the destination vCenter Server is to be deployed is synchronized to NTP or PTP.
- The ESXi host running the source vCenter Server is synchronized to NTP or PTP.
- When upgrading or migrating from vSphere 6.7 to vSphere 8.0, if the vCenter Server appliance is connected to an external Platform Services Controller, ensure the ESXi host running the external Platform Services Controller is synchronized to NTP or PTP.
- If you are upgrading or migrating from vSphere 6.7 to vSphere 8.0, verify that the source vCenter Server or vCenter Server appliance and external Platform Services Controller have the correct time.

Verify that any Windows host machine on which vCenter Server runs is synchronized with the Network Time Server (NTP) server.

To synchronize ESXi clocks with an NTP or a PTP server, you can use the VMware Host Client. For information about editing the time configuration of an ESXi host, see topic *Edit the Time Configuration of an ESXi Host in the VMware Host Client* in the *vSphere Single Host Management - VMware Host Client* documentation.

To learn how to change time synchronization settings for vCenter Server, see topic *Configure the System Time Zone and Time Synchronization Settings* in the *vCenter Server Configuration* documentation.

To learn how to edit the time configuration for a host by using the vSphere Client, see topic *Editing the Time Configuration Settings of a Host* in the *vCenter Server and Host Management* documentation.

Synchronize ESXi Clocks with a Network Time Server

Before you install vCenter Server, make sure all machines on your vSphere network have their clocks synchronized.

This task explains how to set up NTP from the VMware Host Client.

1. Start the VMware Host Client, and connect to the ESXi host.
2. Click **Manage**.
3. Under **System**, click **Time & date**, and click **Edit settings**.
4. Select **Use Network Time Protocol (enable NTP client)**.
5. In the NTP servers text box, enter the IP address or fully qualified domain name of one or more NTP servers to synchronize with.
6. From the **NTP Service Start-up Policy** drop-down menu, select **Start and stop with host**.
7. Click **Save**.

The host synchronizes with the NTP server.

Related Links

[Configuring Time Synchronization Settings in vCenter Server on page 1779](#)

You can change the time synchronization settings in vCenter Server after deployment.

Configuring Time Synchronization Settings in vCenter Server

You can change the time synchronization settings in vCenter Server after deployment.

When you deploy vCenter Server, you can choose the time synchronization method to be either by using an NTP server or by using VMware Tools. In case the time settings in your vSphere network change, you can edit the vCenter Server and configure the time synchronization settings by using the commands in the appliance shell.

When you enable periodic time synchronization, VMware Tools sets the time of the guest operating system to be the same as the time of the host.

After time synchronization occurs, VMware Tools checks once every minute to determine whether the clocks on the guest operating system and the host still match. If not, the clock on the guest operating system is synchronized to match the clock on the host.

Native time synchronization software, such as Network Time Protocol (NTP), is typically more accurate than VMware Tools periodic time synchronization and is therefore preferred. You can use only one form of periodic time synchronization in vCenter Server. If you decide to use native time synchronization software, vCenter Server VMware Tools periodic time synchronization is deactivated.

Related Links

[Synchronize ESXi Clocks with a Network Time Server on page 1347](#)

Before you install vCenter Server, make sure all machines on your vSphere network have their clocks synchronized.

Use VMware Tools Time Synchronization

You can set up vCenter Server to use VMware Tools time synchronization.

1. Access the appliance shell and log in as a user who has the administrator or super administrator role.
The default user with super administrator role is root.
2. Run the command to enable VMware Tools time synchronization.

```
timesync.set --mode host
```

- Optional: Run the command to verify that you successfully applied the VMware Tools time synchronization.

```
timesync.get
```

The command returns that the time synchronization is in host mode.

The time of the appliance is synchronized with the time of the ESXi host.

Add or Replace NTP Servers in the vCenter Server Configuration

To set up the vCenter Server to use NTP-based time synchronization, you must add the NTP servers to the vCenter Server configuration.

- Access the appliance shell and log in as a user who has the administrator or super administrator role.

The default user with super administrator role is root.

- Add NTP servers to the vCenter Server configuration by running the following `ntp.set` command.

```
ntp.set --servers IP-addresses-or-host-names
```

In this command, *IP-addresses-or-host-names* is a comma-separated list of IP addresses or host names of the NTP servers.

This command removes the current NTP servers (if any) and adds the new NTP servers to the configuration. If the time synchronization is based on an NTP server, then the NTP daemon is restarted to reload the new NTP servers. Otherwise, this command replaces the current NTP servers in the NTP configuration with the new NTP servers you specify.

- Optional: To verify that you successfully applied the new NTP configuration settings, run the following command.

```
ntp.get
```

The command returns a space-separated list of the servers configured for NTP synchronization. If the NTP synchronization is activated, the command returns that the NTP configuration is in Up status. If the NTP synchronization is deactivated, the command returns that the NTP configuration is in Down status.

- Optional: To verify if the NTP server is reachable, run the following command.

```
ntp.test --servers IP-addresses-or-host-names
```

The command returns the status of the NTP servers.

If the NTP synchronization is deactivated, you can configure the time synchronization settings in the vCenter Server to be based on an NTP server. See [Synchronize the Time in vCenter Server with an NTP Server](#).

Synchronize the Time in vCenter Server with an NTP Server

You can configure the time synchronization settings in the vCenter Server to be based on an NTP server.

Set up one or more Network Time Protocol (NTP) servers in the vCenter Server configuration. See [Add or Replace NTP Servers in the vCenter Server Configuration](#).

- Access the appliance shell and log in as a user who has the administrator or super administrator role.

The default user with super administrator role is root.

- Run the command to enable NTP-based time synchronization.

```
timesync.set --mode NTP
```

- Optional: Run the command to verify that you successfully applied the NTP synchronization.

```
timesync.get
```

The command returns that the time synchronization is in NTP mode.

Storage Security Best Practices

Follow best practices for storage security, as outlined by your storage security provider. You can also take advantage of CHAP and mutual CHAP to secure iSCSI storage, mask and zone SAN resources, and configure Kerberos credentials for NFS 4.1.

See also the *Administering VMware vSAN* documentation.

Securing iSCSI Storage

The storage you configure for a host might include one or more storage area networks (SANs) that use iSCSI. When you configure iSCSI on a host, you can take measures to minimize security risks.

iSCSI supports accessing SCSI devices and exchanging data by using TCP/IP over a network port rather than through a direct connection to a SCSI device. An iSCSI transaction encapsulates blocks of raw SCSI data in iSCSI records and transmits the data to the requesting device or user.

iSCSI SANs support efficient use of the existing Ethernet infrastructure to provide hosts access to storage resources that they can dynamically share. iSCSI SANs are an economical storage solution for environments that rely on a common storage pool to serve many users. As with any networked system, your iSCSI SANs can be subject to security breaches.

Note: The requirements and procedures for securing an iSCSI SAN are similar for hardware iSCSI adapters associated with hosts and for iSCSI configured directly through the host.

Securing iSCSI Devices

To secure iSCSI devices, require that the ESXi host, or initiator, can authenticate to the iSCSI device, or target, whenever the host attempts to access data on the target LUN.

Authentication ensures that the initiator has the right to access a target. You grant this right when you configure authentication on the iSCSI device.

ESXi does not support Secure Remote Protocol (SRP), or public-key authentication methods for iSCSI. You can use Kerberos only with NFS 4.1.

ESXi supports both CHAP and Mutual CHAP authentication. The *vSphere Storage* documentation explains how to select the best authentication method for your iSCSI device and how to set up CHAP.

Ensure uniqueness of CHAP secrets. Set up a different mutual authentication secret for each host. If possible, set up a different secret for each client that to the ESXi host. Unique secrets ensure that an attacker cannot create another arbitrary host and authenticate to the storage device even if one host is compromised. With a shared secret, compromise of one host might allow an attacker to authenticate to the storage device.

Protecting an iSCSI SAN

When you plan your iSCSI configuration, take measures to improve the overall security of the iSCSI SAN. Your iSCSI configuration is only as secure as your IP network, so by enforcing good security standards when you set up your network, you help safeguard your iSCSI storage.

The following are some specific suggestions for enforcing good security standards.

Protect Transmitted Data

A primary security risk in iSCSI SANs is that an attacker might sniff transmitted storage data.

Take additional measures to prevent attackers from easily seeing iSCSI data. Neither the hardware iSCSI adapter nor ESXi iSCSI initiator encrypts the data that they transmit to and from the targets, making the data more vulnerable to sniffing attacks.

Allowing your virtual machines to share standard switches and VLANs with your iSCSI configuration potentially exposes iSCSI traffic to misuse by a virtual machine attacker. To help ensure that intruders cannot listen to iSCSI transmissions, make sure that none of your virtual machines can see the iSCSI storage network.

If you use a hardware iSCSI adapter, you can accomplish this by making sure that the iSCSI adapter and ESXi physical network adapter are not inadvertently connected outside the host by virtue of sharing a switch or some other means. If you configure iSCSI directly through the ESXi host, you can accomplish this by configuring iSCSI storage through a different standard switch than the one used by your virtual machines.

In addition to protecting the iSCSI SAN by giving it a dedicated standard switch, you can configure your iSCSI SAN on its own VLAN to improve performance and security. Placing your iSCSI configuration on a separate VLAN ensures that no devices other than the iSCSI adapter have visibility into transmissions within the iSCSI SAN. Also, network congestion from other sources cannot interfere with iSCSI traffic.

Secure iSCSI Ports

When you run iSCSI devices, ESXi does not open any ports that listen for network connections. This measure reduces the chances that an intruder can break into ESXi through spare ports and gain control over the host. Therefore, running iSCSI does not present any additional security risks at the ESXi end of the connection.

Any iSCSI target device that you run must have one or more open TCP ports to listen for iSCSI connections. If any security vulnerabilities exist in the iSCSI device software, your data can be at risk through no fault of ESXi. To lower this risk, install all security patches that your storage equipment manufacturer provides and limit the devices connected to the iSCSI network.

Masking and Zoning SAN Resources

You can use zoning and LUN masking to separate SAN activity and restrict access to storage devices.

You can protect access to storage in your vSphere environment by using zoning and LUN masking with your SAN resources. For example, you might manage zones defined for testing independently within the SAN so they do not interfere with activity in the production zones. Similarly, you might set up different zones for different departments.

When you set up zones, take into account any host groups that are set up on the SAN device.

Zoning and masking capabilities for each SAN switch and disk array and the tools for managing LUN masking are vendor specific.

See your SAN vendor's documentation and the *vSphere Storage* documentation.

Using Kerberos for NFS 4.1

With NFS version 4.1, ESXi supports the Kerberos authentication mechanism.

The RPCSEC_GSS Kerberos mechanism is an authentication service. It allows an NFS 4.1 client installed on ESXi to prove its identity to an NFS server before mounting an NFS share. The Kerberos security uses cryptography to work across an insecure network connection.

The ESXi implementation of Kerberos for NFS 4.1 provides two security models, krb5 and krb5i, that offer different levels of security.

- Kerberos for authentication only (krb5) supports identity verification.
- Kerberos for authentication and data integrity (krb5i), in addition to identity verification, provides data integrity services. These services help to protect the NFS traffic from tampering by checking data packets for any potential modifications.

Kerberos supports cryptographic algorithms that prevent unauthorized users from gaining access to NFS traffic. The NFS 4.1 client on ESXi attempts to use either the AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 algorithm to access a share on the NAS server. Before using your NFS 4.1 datastores, make sure that AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-SHA1-96 are enabled on the NAS server.

The following table compares Kerberos security levels that ESXi supports.

Table 271: Types of Kerberos Security

		ESXi 6.0	ESXi 6.5 and later
Kerberos for authentication only (krb5)	Integrity checksum for RPC header	Yes with DES	Yes with AES
	Integrity checksum for RPC data	No	No
Kerberos for authentication and data integrity (krb5i)	Integrity checksum for RPC header	No krb5i	Yes with AES
	Integrity checksum for RPC data		Yes with AES

When you use Kerberos authentication, the following considerations apply:

- ESXi uses Kerberos with the Active Directory domain.
- As a vSphere administrator, you specify Active Directory credentials to provide access to NFS 4.1 Kerberos datastores for an NFS user. A single set of credentials is used to access all Kerberos datastores mounted on that host.
- When multiple ESXi hosts share the NFS 4.1 datastore, you must use the same Active Directory credentials for all hosts that access the shared datastore. To automate the assignment process, set the user in host profiles and apply the profile to all ESXi hosts.
- You cannot use two security mechanisms, AUTH_SYS and Kerberos, for the same NFS 4.1 datastore shared by multiple hosts.

Verify That Sending Host Performance Data to Guests Is Deactivated

vSphere includes virtual machine performance counters on Windows operating systems where VMware Tools is installed. Performance counters allow virtual machine owners to do accurate performance analysis within the guest operating system. By default, vSphere does not expose host information to the guest virtual machine.

By default, the capability to send host performance data to a virtual machine is deactivated. This default setting prevents a virtual machine from obtaining detailed information about the physical host. If a security breach of the virtual machine occurs, the setting does not make host data available to the attacker.

Note: The following procedure illustrates the basic process. Consider using ESXCLI or VMware PowerCLI commands for performing this task on all hosts simultaneously.

1. On the ESXi system that hosts the virtual machine, browse to the VMX file.
Virtual machine configuration files are located in the `/vmfs/volumes/datastore` directory, where *datastore* is the name of the storage device where the virtual machine files are stored.
2. In the VMX file, verify that the following parameter is set.
`tools.guestlib.enableHostInfo=FALSE`
3. Save and close the file.

You cannot retrieve performance information about the host from inside the guest virtual machine.

Setting Timeouts for the ESXi Shell and the vSphere Client

To prevent intruders from using an idle session, set timeouts for the ESXi Shell and the vSphere Client.

ESXi Shell Timeout

For the ESXi Shell, you can set the following timeouts from the vSphere Client and from the Direct Console User Interface (DCUI).

Availability Timeout

The availability timeout setting is the amount of time that can elapse before you must log in after the shell is activated. After the timeout period, the service is deactivated and users are not allowed to log in.

Idle Timeout

The idle timeout is the amount of time that can elapse before the user is logged out of an idle interactive session. Changes to the idle timeout apply the next time a user logs in to the shell. Changes do not affect existing sessions.

Change the vSphere Client Timeout

You can configure the timeout intervals for vCenter Server operations. These intervals specify the amount of time after which the vSphere Client times out.

Prerequisites:

- Required privilege: **Global > Settings**
- 1. In the vSphere Client, navigate to the vCenter Server instance.
- 2. Select the **Configure** tab.
- 3. Under **Settings**, select **Advanced Settings**.
- 4. In **client.timeout.normal**, enter the timeout interval in seconds for normal operations.
The value cannot be less than five (5).
- 5. In **client.timeout.long**, enter the timeout interval in minutes for long operations.
The value cannot be less than five (5).
- 6. Click **Save**.
- 7. Restart the vCenter Server system for the changes to take effect.

vSphere TLS Configuration

Starting in 8.0 Update 3, vSphere supports TLS 1.3 and 1.2 through the use of TLS profiles. TLS profiles simplify the job of administering TLS parameters and also improve supportability.

vSphere 8.0 Update 3 activates the default TLS profile, named COMPATIBLE, on ESXi and vCenter Server hosts. The COMPATIBLE profile supports TLS 1.3, and some TLS 1.2 connections.

You can manage TLS profiles on ESXi hosts either by using vSphere Configuration Profiles or `esxcli` commands. On vCenter Server hosts, you can manage TLS profiles by using the APIs. For example, you can use the Developer Center in the vSphere Client. See *vSphere Automation SDKs Programming Guide* and *vSphere Automation REST API Programming Guide*.

vCenter Server and Envoy

vCenter Server runs two reverse proxy services:

- VMware reverse proxy service, `rhttpproxy`
- Envoy

Envoy is an open source edge and service proxy. Envoy owns port 443, and all incoming vCenter Server requests are routed through Envoy. The `rhttpproxy` serves as a configuration management server for Envoy. As a result, the TLS configuration is applied to `rhttpproxy`, which in turn sends the configuration to Envoy.

How vSphere Implements TLS Using TLS Profiles

vSphere 8.0 Update 3 implements TLS 1.3 by grouping parameters, including protocol versions, groups (also called curves), and ciphers, into a single TLS profile. This TLS profile is applied system-wide. Using a single TLS profile eases administrative overhead of your hosts. You no longer need to manually configure individual TLS parameters, though that capability is still available if required. TLS profiles also significantly improve supportability. The grouping of parameters into TLS profiles simplifies the set of VMware verified TLS solutions from which to choose. On ESXi, TLS profiles are integrated with vSphere Configuration Profiles.

The following ESXi TLS profiles are provided:

- **COMPATIBLE:** The default profile. The exact mapping of the parameters in this profile can change from release to release but the profile is guaranteed to be compatible with all products and versions supported (currently N-2 versions). That is, an ESXi host from release N using the COMPATIBLE profile can communicate with a host from release N-2.
- **NIST_2024:** A more restrictive profile that specifically supports the NIST 2024 standard. The exact mapping of parameters in this profile is guaranteed to satisfy the NIST 2024 standard across releases. This profile is guaranteed to be compatible only with current or newer releases, and not older releases.
- **MANUAL:** Use this profile to create and test an ad hoc configuration in which you manually supply the TLS parameters. It is not guaranteed that a MANUAL profile functions error-free, because Broadcom does not test all possible combinations of parameters that are possible in the MANUAL profile. You must test a MANUAL profile, including across software upgrades. When you choose to use the MANUAL profile, the system behavior first defaults to the profile previously selected (COMPATIBLE or NIST_2024), and remains as such until you make changes. You must use `esxcli` commands to manage the MANUAL TLS profile. See the help text that ships with `esxcli` for more information about changing the parameters in a MANUAL TLS profile.

When configuring the TLS profile to the desired state, you must reboot the ESXi host or remediate the vLCM cluster in which the ESXi host resides to apply changes.

The following tables show the details of TLS profiles for ESXi and vCenter Server in vSphere 8.0 Update 3. The Cipher List column shows the TLS ciphers for TLS 1.2 and below protocols. The Cipher Suites column shows the ciphers for the TLS 1.3 protocol.

Table 272: ESXi TLS 1.3 Profiles

TLS Profile Name	TLS Protocol Versions	Cipher List	Cipher Suites	Curves	VMware Supported?
COMPATIBLE	TLS 1.3 and TLS 1.2	ECDHE+AES GCM:ECDHE+AES	TLS_AES_256_GCM_SHA384; TLS_AES_128_GCM_SHA256	prime256v1:secp384r1:secp521r1	Yes
NIST_2024	TLS 1.3 and TLS 1.2	ECDHE+AES GCM	TLS_AES_256_GCM_SHA384; TLS_AES_128_GCM_SHA256	prime256v1:secp384r1:secp521r1	Yes
MANUAL	Any	Any	Any	Any	No

Notes:

- Supported settings (protocols, cipher list, cipher suites, and curves) represent at most what is supported.

- The NIST_2024 profile applies to inbound connections only.
- The BoringSSL cryptographic module used in vSphere 8.0 Update 3 has not yet reached FIPS certification for TLS 1.3 usage. As a result, both on ESXi and vCenter Server, port 443 (Reverse Proxy) communicates using TLS 1.2. The COMPATIBLE and NIST_2024 TLS profiles do not use non-FIPS TLS 1.3.

The following vCenter Server TLS 1.3 profiles are provided:

- COMPATIBLE: The default profile. The exact mapping of the parameters in this profile can change from release to release but the profile is guaranteed to be compatible with all products and versions supported (currently N-2 versions).
- NIST_2024: A more restrictive profile that specifically supports the NIST 2024 standard. The exact mapping of parameters in this profile is guaranteed to satisfy the NIST 2024 standard across releases. This profile is guaranteed to be compatible only with current or newer releases, and not older releases.
- COMPATIBLE-NON-FIPS: A modified profile that allows a non-FIPS TLS 1.3 connection from the Envoy proxy. FIPS is not enabled.

Table 273: vCenter Server TLS 1.3 Profiles

TLS Profile Name	TLS Protocol Versions	Cipher Suites	Curves	FIPS Enabled?	VMware Supported?
COMPATIBLE	TLS 1.3	TLS_AES_256_GCM_SHA384; TLS_AES_128_GCM_SHA256	prime256v1:secp384r1:secp521r1	Yes	Yes
	TLS 1.2	ECDHE-RSA-AES256-GCM-SHA384 ECDHE-RSA-AES128-GCM-SHA256 ECDHE-ECDSA-AES256-GCM-SHA384 ECDHE-ECDSA-AES128-GCM-SHA256 AES256-GCM-SHA384 AES128-GCM-SHA256 ECDHE-RSA-AES256-SHA ECDHE-RSA-AES128-SHA ECDHE-ECDSA-AES256-SHA ECDHE-ECDSA-AES128-SHA AES256-SHA AES128-SHA			
NIST_2024	TLS 1.3	TLS_AES_256_GCM_SHA384 TLS_AES_128_GCM_SHA256	prime256v1:secp384r1:secp521r1	Yes	Yes
	TLS 1.2	ECDHE-RSA-AES256-GCM-SHA384 ECDHE-RSA-AES128-GCM-SHA256 ECDHE-ECDSA-AES256-GCM-SHA384 ECDHE-ECDSA-AES128-GCM-SHA256			
COMPATIBLE-NON-FIPS	TLS 1.3	TLS_AES_256_GCM_SHA384 TLS_AES_128_GCM_SHA256	prime256v1:secp384r1:secp521r1	No	Yes
	TLS 1.2	ECDHE-RSA-AES256-GCM-SHA384 ECDHE-RSA-AES128-GCM-SHA256 ECDHE-ECDSA-AES256-GCM-SHA384 ECDHE-ECDSA-AES128-GCM-SHA256 AES256-GCM-SHA384 AES128-GCM-SHA256 ECDHE-RSA-AES256-SHA ECDHE-RSA-AES128-SHA ECDHE-ECDSA-AES256-SHA ECDHE-ECDSA-AES128-SHA			

TLS Profile Name	TLS Protocol Versions	Cipher Suites	Curves	FIPS Enabled?	VMware Supported?
		AES256-SHA AES128-SHA			

Note: You can switch back and forth between TLS profiles. For example, if you change from using the COMPATIBLE profile to using the NIST_2024 profile, you can switch back to using the COMPATIBLE profile. There is no limitation on switching back and forth between TLS profiles. For more information, see [Managing vSphere TLS](#).

TLS and Inbound and Outbound Connections in ESXi and vCenter Server

ESXi 8.0 Update 3 supports TLS 1.3 on both inbound (server) and outbound (client) connections. The ESXi inbound (server) connections are of most concern, and where the more restrictive NIST_2024 profile applies.

For ESXi, you can use the COMPATIBLE, NIST_2024, and MANUAL settings on inbound (server) connections. You can use the COMPATIBLE and MANUAL settings on outbound (client) connections.

vCenter Server TLS profiles apply their settings to both inbound and outbound connections.

Some vSphere services expose ports that accept TLS connections, while most services use the Reverse Proxy. All inbound connections accept TLS 1.2 and TLS 1.3 by default. Currently, port 443 (Reverse Proxy) has TLS 1.3 disabled and communicates by using TLS 1.2. Outbound connections support TLS 1.2 and TLS 1.3. For more information, see [TLS 1.3 on Port 443 in ESXi and FIPS](#).

TLS and Lifecycle Management

Upgrading or migrating an ESXi host or vCenter Server host to 8.0 Update 3 enables the COMPATIBLE TLS profile by default. vSphere 8.0 Update 3 supports TLS 1.3, and TLS 1.2 for bare minimum interoperability out of the box. In the future, upgrading to a later version of ESXi or vCenter Server retains the current TLS profile in use as long as that profile has not been retired.

When upgrading to a new version, as a recommended best practice, first set the TLS profile to COMPATIBLE.

If you make local service-level edits before upgrading to vSphere 8.0 Update 3, after the upgrade, the host is assigned the COMPATIBLE profile, which does not reflect those changes. To have the host reflect those changes, switch to using the MANUAL profile. See [Change the TLS Profile of an ESXi Host Using the vSphere Client](#) or [Change the TLS Profile of an ESXi Host Using the CLI](#).

Warning: The MANUAL TLS profile is not guaranteed to work error-free across upgrades. You must verify that an edited MANUAL TLS profile functions from one release to another, or switch to the COMPATIBLE TLS profile prior to upgrading.

TLS 1.3 on Port 443 in ESXi and FIPS

Currently, vSphere disables TLS 1.3 on port 443. The version of Boring SSL cryptographic module used in vSphere 8.0 Update 3 is not FIPS certified for TLS 1.3. When using the COMPATIBLE or NIST_2024 TLS profile, all ports except 443 communicate by TLS 1.3. For now, because of this issue, port 443 uses TLS 1.2.

To enable non-FIPS TLS 1.3 on port 443, see the VMware knowledge base article at <https://kb.vmware.com/s/article/92473>.

Managing vSphere TLS

Starting in vSphere 8.0 Update 3, you can manage TLS profiles for ESXi by using the vSphere Client, `esxcli` commands, or the APIs. For vCenter Server, you manage TLS profiles by using the APIs.

If you use vSphere Configuration Profiles, you can manage the TLS setting for ESXi hosts at the vLCM cluster level. You can change the TLS setting for the cluster and remediate the cluster against this new configuration. For more information, see the chapter on managing vSphere Configuration Profiles in the *Managing Host and Cluster Lifecycle* documentation.

For standalone ESXi hosts, and non-vLCM clusters, you must manage the TLS profile by using `esxcli` commands. See the *ESXCLI Concepts and Examples* documentation and the `esxcli` online help.

Currently, you can only manage vCenter Server TLS profiles by using the APIs. See *vSphere Automation SDKs Programming Guide* and *vSphere Automation REST API Programming Guide*.

View the TLS Profile of an ESXi Host Using the vSphere Client

You can use the vSphere Client to view the TLS profile of an ESXi host that is part of a vLCM cluster.

You have enabled vSphere Configuration Profiles and created a draft configuration for the cluster. See the *Managing Host and Cluster Lifecycle* documentation.

In vSphere Configuration Profiles, settings that are not explicitly configured use the default values from the appropriate profile. For TLS profiles, the default is COMPATIBLE.

To view the TLS profile of a stand-alone or non-vLCM cluster ESXi host, see [View the TLS Profile of an ESXi Host Using the CLI](#).

1. In the vSphere Client, navigate to the vLCM cluster that you manage with a single image.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Settings** tab, click **system**.
4. Click **tls_client** or **tls_server** to view which TLS profile is defined in the current desired configuration document.

View the TLS Profile of an ESXi Host Using the CLI

You can use the CLI to view the currently configured TLS profile of an ESXi host.

Enable either SSH or the ESXi Shell on the ESXi host.

For standalone ESXi hosts, and non-vLCM clusters, you must manage the TLS profile by using `esxcli` commands. For more information, see *ESXCLI Reference*. For ESXi hosts in a vLCM cluster, you can use either vSphere Configuration Profiles or `esxcli` commands.

1. Connect to the ESXi host.
You can use SSH or the ESXi Shell.
2. To view the currently configured TLS profile, run the following command.
3. To view the parameters in the currently configured TLS profile, run the following command:

```
esxcli system tls [client | server] get
```

```
esxcli system tls [client | server] get --show-profile-defaults
```

Change the TLS Profile of an ESXi Host Using the vSphere Client

You can change the TLS profile of an ESXi host. The default TLS profile is COMPATIBLE.

You have enabled vSphere Configuration Profiles and created a draft configuration for the cluster. See the *Managing Host and Cluster Lifecycle* documentation.

1. In the vSphere Client, navigate to a cluster that you manage with a single image.
2. On the **Configure** tab, click **Desired State > Configuration**.
3. On the **Settings** tab, click **system**.
4. Click either **tls_client** or **tls_server**.
Depending on if the setting has been previously changed, click either **Configure Settings** or **Edit**.

5. Select the TLS profile from the drop-down.
6. Click **Save**.
7. Remediate the cluster against the draft configuration.
 - a) To remediate the cluster against the draft configuration, on the **Draft** tab, click **Apply Changes**.
 - b) Follow the steps in the **Remediate** wizard. For more information, see the *Managing Host and Cluster Lifecycle* documentation.

All the ESXi hosts in the cluster are compliant with the desired configuration.

Change the TLS Profile of an ESXi Host Using the CLI

You can change the TLS profile of an ESXi host. The default TLS profile is COMPATIBLE.

Enable either SSH or the ESXi Shell on the ESXi host.

For standalone ESXi hosts, and non-vLCM clusters, you must manage the TLS profile by using `esxcli` commands. For more information, see *ESXCLI Reference*. For ESXi hosts in a vLCM cluster, you can use either vSphere Configuration Profiles or `esxcli` commands.

1. Connect to the ESXi host.
You can use either SSH or the ESXi Shell.
2. Put the ESXi host into maintenance mode.
3. To change the TLS profile, run the following command.

```
esxcli system tls [client | server] set --profile [COMPATIBLE | NIST_2024 | MANUAL]
```

Note: If you want to make changes to TLS parameters (either at the system-level or service-level), select the MANUAL profile.

4. Reboot the ESXi host for the change to take effect.
5. After the ESXi host reboots, take it out of maintenance mode.

Edit the Parameters in the MANUAL TLS Profile Using the CLI

You can edit the set of parameters in the MANUAL TLS profile. To change TLS parameters such as cipher list and cipher suite, you must first set the TLS profile to MANUAL.

Enable either SSH or the ESXi Shell on the ESXi host.

Change the TLS profile to MANUAL. See either [Change the TLS Profile of an ESXi Host Using the vSphere Client](#) or [Change the TLS Profile of an ESXi Host Using the CLI](#).

Warning:

Broadcom does not test all possible combinations of TLS settings that can be configured in the MANUAL TLS profile. Before enabling custom settings using the MANUAL profile in your production environment, you should test these settings to ensure that your TLS servers and clients operate correctly, and that lifecycle operations such as upgrades can be completed successfully with these settings.

You must administer parameters in the MANUAL TLS profile by using `esxcli` commands. Administering the MANUAL TLS profile parameters is not integrated with vSphere Configuration Profiles.

You cannot set TLS parameters for individual vSphere services. The changes that you make by using the MANUAL TLS profile are applied at the system level.

1. Connect to the ESXi host.

You can use either SSH or the ESXi Shell.

2. Put the ESXi host into maintenance mode.

3. Ensure that the TLS profile is MANUAL.

```
esxcli system tls [client | server] get
```

4. To change the parameters, run any of the following commands.

```
esxcli system tls [client | server] set --cipher-list=str
esxcli system tls [client | server] set --cipher-suite=str
esxcli system tls [client | server] set --groups=str
esxcli system tls [client | server] set --protocol-versions=str
```

where *str* is a string in OpenSSL-style that is colon-, comma-, or space-delimited. For example: `--cipher-list=ECDHE+AESGCM:ECDHE+AES`

For more information, run the following command:

```
esxcli system tls [client | server] set --help
```

5. Reboot the ESXi host for the change to take effect.

6. After the ESXi host reboots, take it out of maintenance mode.

The following example first sets the TLS profile to MANUAL then sets a more restrictive set of curves (groups). A reboot would be required to put the changes into effect.

```
[root@host1] esxcli system tls server get
Profile: COMPATIBLE
Cipher List: <profile default>
Cipher Suite: <profile default>
Groups: <profile default>
Protocol Versions: <profile default>
Reboot Required: false
[root@host1] esxcli system tls server set --profile MANUAL
[root@host1] esxcli system tls server get
Profile: MANUAL
Cipher List: ECDHE+AESGCM:ECDHE+AES
Cipher Suite: TLS_AES_128_GCM_SHA256:TLS_AES_256_GCM_SHA384
Groups: prime256v1:secp384r1:secp521r1
Protocol Versions: tls1.2,tls1.3
Reboot Required: true
[root@host1] esxcli system tls server set --groups=prime256v1:secp384r1
[root@host1] esxcli system tls server get
Profile: MANUAL
Cipher List: TLS_AES_128_GCM_SHA256
Cipher Suite: TLS_AES_128_GCM_SHA256:TLS_AES_256_GCM_SHA384
Groups: prime256v1:secp384r1
Protocol Versions: tls1.2,tls1.3
Reboot Required: true
```

Manage the TLS Profile of a vCenter Server Host

You use the APIs to view and change the TLS profile for a vCenter Server host.

You can use various ways to execute HTTP requests. This task shows how to use the Developer Center in the vSphere Client to manage TLS profiles. See *VMware vCenter Server Management Programming Guide* for more information about using APIs to manage the vCenter Server Appliance.

1. Log in to the vCenter Server system with the vSphere Client.
2. From the Menu, select **Developer Center**.
3. Click **API Explorer**.
4. From the **Select API** drop-down, select **appliance**.

The following API categories and actions are available.

Table 274: vCenter Server TLS APIs

Option	API Category	Associated Action
Gets the list of all TLS profiles and their configuration.	tls/profiles/	GET
Gets the parameters of a specific TLS profile.	tls/profiles/{id}	GET
Gets the name of the current TLS profile configured globally.	tls/profiles/global/	GET
Sets one of the standard profiles that you specify globally.	tls/profiles/global/	PUT Note: This action restarts the vCenter Server services.
Gets the parameters of the current TLS profile configured globally.	tls/manual-parameters/global	GET

Note: Currently, you cannot change the parameters of a vCenter Server TLS profile.

5. Execute the desired command.

Defined Privileges

The following tables list the default privileges that, when selected for a role, can be paired with a user and assigned to an object.

vCenter Server allows fine-grained control over authorization with permissions and roles. When you assign a permission to an object in the vCenter Server object hierarchy, you specify which user or group has which privileges on that object. To specify the privileges, you use roles, which are sets of privileges. For more information, see:

- [Managing Permissions for vCenter Server Components](#)
- [Using vCenter Server Roles to Assign Privileges](#)

When setting permissions, verify all the object types are set with appropriate privileges for each particular action. Some operations require access permission at the root folder or parent folder in addition to access to the object being manipulated. Some operations require access or performance permission at a parent folder and a related object. See also [Privilege Recorder](#).

vCenter Server extensions might define additional privileges not listed here. Refer to the documentation for the extension for more information on those privileges.

Alarms Privileges

Alarms privileges control the ability to create, modify, and respond to alarms on inventory objects.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 275: Alarms Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Acknowledge alarm	Allows suppression of all alarm actions on all triggered alarms.	Object on which an alarm is defined	Alarm.Acknowledge
Create alarm	Allows creation of a new alarm. When creating alarms with a custom action, privilege to perform the action is verified when the user creates the alarm.	Object on which an alarm is defined	Alarm.Create
Disable alarm action	Allows stopping an alarm action from occurring after an alarm has been triggered. The alarm itself is not deactivated.	Object on which an alarm is defined	Alarm.DisableActions
Disable or enable alarm on entity	Allows activating or deactivating a particular alarm on a particular target type.	Object on which the alarm can trigger	Alarm.ToggleEnableOnEntity
Modify alarm	Allows changing the properties of an alarm.	Object on which an alarm is defined	Alarm.Edit
Remove alarm	Allows deletion of an alarm.	Object on which an alarm is defined	Alarm.Delete
Set alarm status	Allows changing the status of the configured event alarm. The status can change to Normal , Warning , or Alert .	Object on which an alarm is defined	Alarm.SetStatus

Auto Deploy and Image Profile Privileges

Auto Deploy privileges control who can perform different tasks on Auto Deploy rules, and who can associate a host. Auto Deploy privileges also allow you to control who can create or edit an image profile.

The table describes privileges that determine who can manage Auto Deploy rules and rule sets and who can create and edit image profiles. For more information about Auto Deploy, see the *VMware ESXi Installation and Setup* documentation.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 276: Auto Deploy Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Host – AssociateMachine	Allows users to associate a host with a machine.	vCenter Server	AutoDeploy.Host.AssociateMachine
• Image Profile – Create – Edit	Create allows creation of image profiles. Edit allows editing of image profiles.	vCenter Server	AutoDeploy.Profile.Create AutoDeploy.Profile.Edit
• Rule – Create – Edit – Delete	Create allows creation of Auto Deploy rules. Edit allows editing of Auto Deploy rules. Delete allows deletion of Auto Deploy rules.	vCenter Server	AutoDeploy.Rule.Create AutoDeploy.Rule.Edit AutoDeploy.Rule.Delete
• RuleSet – Activate – Edit	Activate allows activation of Auto Deploy rule sets. Edit allows editing of Auto Deploy rule sets.	vCenter Server	AutoDeploy.RuleSet.Activate AutoDeploy.RuleSet.Edit

Certificates Privileges

Certificates privileges control which users can manage ESXi certificates.

This privilege determines who can perform certificate management for ESXi hosts. See Required Privileges for Certificate Management Operations in the *vSphere Authentication* documentation for information on vCenter Server certificate management.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 277: Host Certificates Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Manage Certificates	Allows certificate management for ESXi hosts.	vCenter Server	Certificate.Manage

Certificate Authority Privileges

Certificate authority privileges control aspects of VMware Certificate Authority (VMCA) certificates.

Table 278: Certificate Authority Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create/Delete (Admins priv).	Allows full administrative-level access for managing vCenter Server certificates.	vCenter Server	CertificateAuthority.Administer

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create/Delete (below Admins priv).	Allows viewing the VMCA root certificate in the Certificate Management page in the vSphere Client.	vCenter Server	CertificateAuthority.Manage

Certificate Management Privileges

Certificate management privileges control which users can manage vCenter Server certificates.

Table 279: Certificate Management Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create/Delete (Admins priv).	Allows full administrative-level access to various internal APIs and functionality for vCenter Server certificate-related operations.	vCenter Server	CertificateManagement.Administer
Create/Delete (below Admins priv).	Allows reduced administrative access to various internal APIs and functionality. This privilege restricts certificate related operations so that the user cannot escalate non-administrator privileges. Allowed operations are: • Generating certificate signing requests • Creating and retrieving Trusted Root chains • Deleting Trusted Root chains created by a user with the privilege Certificate Management > Create/Delete (below Admins priv). • Retrieving Machine SSL certificates • Retrieving the signing certificate chains for validating tokens issued by vCenter Server	vCenter Server	CertificateManagement.Manage

Cns Privileges

Cloud Native Store (Cns) privileges control which users can access the Cloud Native Storage UI.

Table 280: Cns Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Searchable	Allows storage administrator to view the Cloud Native Storage UI.	Root vCenter Server	Cns.Searchable

Compute Policy Privileges

Compute policy privileges control the ability to manage compute policies.

Table 281: Compute Policy Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create and Delete Compute Policy	Allows creating and deleting compute policies.	Root vCenter Server	ComputePolicy.Manage

Content Library Privileges

Content Libraries provide simple and effective management for virtual machine templates and vApps. Content library privileges control who can view or manage different aspects of content libraries.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Note: Inheritance of permissions for content libraries works in the context of a single vCenter Server instance. However, content libraries are not direct children of vCenter Server system from an inventory perspective. The direct parent for content libraries is the global root. This relationship means that if you set a permission at a vCenter Server level and propagate it to the children objects, the permission applies to data centers, folders, clusters, hosts, virtual machines, and so on, but does not apply to the content libraries that you see and operate with in this vCenter Server instance. To assign a permission on a content library, an Administrator must grant the permission to the user as a global permission. Global permissions support assigning privileges across solutions from a global root object.

Table 282: Content Library Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Add library item	Allows addition of items in a library.	Library	ContentLibrary.AddLibraryItem
Add root certificate to trust store	Allows addition of root certificates to the Trusted Root Certificates Store.	vCenter Server	ContentLibrary.AddCertToTrustStore
Check in a template	Allows checking in of templates.	Library	ContentLibrary.CheckInTemplate
Check out a template	Allows checking out of templates.	Library	ContentLibrary.CheckOutTemplate
Create a subscription for a published library	Allows creation of a library subscription.	Library	ContentLibrary.AddSubscription

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create local library	Allows creation of local libraries on the specified vCenter Server system.	vCenter Server	ContentLibrary.CreateLocalLibrary
Create or delete a Harbor registry	Allows creation or deletion of the VMware Tanzu Harbor Registry service.	vCenter Server for creation. Registry for deletion.	ContentLibrary.ManageRegistry
Create subscribed library	Allows creation of subscribed libraries.	vCenter Server	ContentLibrary.CreateSubscribedLibrary
Create, delete or purge a Harbor registry project	Allows creation, deletion, or purging of VMware Tanzu Harbor Registry projects.	Registry	ContentLibrary.ManageRegistryProject
Delete library item	Allows deletion of library items.	Library. Set this permission to propagate to all library items.	ContentLibrary.DeleteLibraryItem
Delete local library	Allows deletion of a local library.	Library	ContentLibrary.DeleteLocalLibrary
Delete root certificate from trust store	Allows deletion of root certificates from the Trusted Root Certificates Store.	vCenter Server	ContentLibrary.DeleteCertFromTrustStore
Delete subscribed library	Allows deletion of a subscribed library.	Library	ContentLibrary.DeleteSubscribedLibrary
Delete subscription of a published library	Allows deletion of a subscription to a library.	Library	ContentLibrary.DeleteSubscription
Download files	Allows download of files from the content library.	Library	ContentLibrary.DownloadSession
Evict library item	Allows eviction of items. The content of a subscribed library can be cached or not cached. If the content is cached, you can release a library item by evicting it if you have this privilege.	Library. Set this permission to propagate to all library items.	ContentLibrary.EvictLibraryItem
Evict subscribed library	Allows eviction of a subscribed library. The content of a subscribed library can be cached or not cached. If the content is cached, you can release a library by evicting it if you have this privilege.	Library	ContentLibrary.EvictSubscribedLibrary

Privilege Name in thevSphere Client	Description	Required On	Privilege Name in the API
Import Storage	Allows a user to import a library item if the source file URL starts with <code>ds://</code> or <code>file://</code> . This privilege is deactivated for content library administrator by default. Because an import from a storage URL implies import of content, activate this privilege only if necessary and if no security concern exists for the user who performs the import.	Library	<code>ContentLibrary.ImportStorage</code>
Manage Harbor registry resources on specified compute resource	Allows management of VMware Tanzu Harbor Registry resources.	Compute cluster	<code>ContentLibrary.ManageClusterRegistryResource</code>
Probe subscription information	This privilege allows solution users and APIs to probe a remote library's subscription info including URL, SSL certificate, and password. The resulting structure describes whether the subscription configuration is successful or whether there are problems such as SSL errors.	Library	<code>ContentLibrary.ProbeSubscription</code>
Publish a library item to its subscribers	Allows publication of library items to subscribers.	Library. Set this permission to propagate to all library items.	<code>ContentLibrary.PublishLibraryItem</code>
Publish a library to its subscribers	Allows publication of libraries to subscribers.	Library	<code>ContentLibrary.PublishLibrary</code>
Read storage	Allows reading of content library storage.	Library	<code>ContentLibrary.ReadStorage</code>
Sync library item	Allows synchronization of library items.	Library. Set this permission to propagate to all library items.	<code>ContentLibrary.SyncLibraryItem</code>
Sync subscribed library	Allows synchronization of subscribed libraries.	Library	<code>ContentLibrary.SyncLibrary</code>
Type introspection	Allows a solution user or API to introspect the type support plug-ins for the content library service.	Library	<code>ContentLibrary.TypeIntrospection</code>

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Update configuration settings	Allows you to update the configuration settings. No vSphere Client user interface elements are associated with this privilege.	Library	ContentLibrary.UpdateConfiguration
Update files	Allows you to upload content into the content library. Also allows you to remove files from a library item.	Library	ContentLibrary.UpdateSession
Update library	Allows updates to the content library.	Library	ContentLibrary.UpdateLibrary
Update library item	Allows updates to library items.	Library. Set this permission to propagate to all library items.	ContentLibrary.UpdateLibraryItem
Update local library	Allows updates of local libraries.	Library	ContentLibrary.UpdateLocalLibrary
Update subscribed library	Allows you to update the properties of a subscribed library.	Library	ContentLibrary.UpdateSubscribedLibrary
Update subscription of a published library	Allows updates of subscription parameters. Users can update parameters such as the subscribed library's vCenter Server instance specification and placement of its virtual machine template items.	Library	ContentLibrary.UpdateSubscription
View configuration settings	Allows you to view the configuration settings. No vSphere Client user interface elements are associated with this privilege.	Library	ContentLibrary.GetConfiguration

Cryptographic Operations Privileges

Cryptographic operations privileges control who can perform which type of cryptographic operation on which type of object.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 283: Cryptographic Operations Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Direct Access	Allows users access to encrypted resources. Users can export virtual machines, have NFC access to virtual machines, and open a console session to an encrypted virtual machine.	Virtual machine, host, or datastore	Cryptographer.Access
Add disk	Allows users to add a disk to an encrypted virtual machine.	Virtual machine	Cryptographer.AddDisk
Clone	Allows users to clone an encrypted virtual machine.	Virtual machine	Cryptographer.Clone
Decrypt	Allows users to decrypt a virtual machine or disk.	Virtual machine	Cryptographer.Decrypt
Encrypt	Allows users to encrypt a virtual machine or a virtual machine disk.	Virtual machine	Cryptographer.Encrypt
Encrypt new	Allows users to encrypt a virtual machine during virtual machine creation or a disk during disk creation.	Virtual machine folder	Cryptographer.EncryptNew
Manage encryption policies	Allows users to manage virtual machine storage policies with encryption IO filters. By default, virtual machines that use the Encryption storage policy do not use other storage policies.	vCenter Server root folder	Cryptographer.ManageEncryptionPolicy
Manage KMS	Allows users to manage the Key Management Server for the vCenter Server system.	vCenter Server system	Cryptographer.ManageKeyServers

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
	Management tasks include adding and removing KMS instances, and establishing a trust relationship with the KMS.		
Manage keys	Allows users to perform key management operations. These operations are not supported from the vSphere Client but can be performed by using <code>crypto-util</code> or the API.	vCenter Server root folder	Cryptographer.ManageKeys
Migrate	Allows users to migrate an encrypted virtual machine to a different ESXi host. Supports migration with or without vMotion and storage vMotion. Supports migration to a different vCenter Server instance.	Virtual machine	Cryptographer.Migrate
Recrypt	Allows users to recrypt virtual machines or disks with a different key. This privilege is required for both deep and shallow recrypt operations.	Virtual machine	Cryptographer.Recrypt
Register VM	Allows users to register an encrypted virtual machine with an ESXi host.	Virtual machine folder	Cryptographer.RegisterVM
Register host	Allows users to enable encryption on a host. You can enable encryption on a host explicitly, or the virtual machine creation process can enable it.	Host folder for standalone hosts, cluster for hosts in cluster	Cryptographer.RegisterHost

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Read KMS information	Allows users to list vSphere Native Key Providers on the vCenter Server and on hosts. Also allows users to get vSphere Native Key Provider information.	vCenter Server or host	Cryptographer.ReadKeyServersInfo

dvPort Group Privileges

Distributed virtual port group privileges control the ability to create, delete, and modify distributed virtual port groups.

The table describes the privileges required to create and configure distributed virtual port groups.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 284: Distributed Virtual Port Group Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create	Allows creation of a distributed virtual port group.	Virtual port groups	DVPortgroup.Create
Delete	Allows deletion of distributed virtual port group. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Virtual port groups	DVPortgroup.Delete
Modify	Allows modification of a distributed virtual port group configuration.	Virtual port groups	DVPortgroup.Modify
Policy operation	Allows setting the policy of a distributed virtual port group.	Virtual port groups	DVPortgroup.PolicyOp
Scope operation	Allows setting the scope of a distributed virtual port group.	Virtual port groups	DVPortgroup.ScopeOp

Distributed Switch Privileges

Distributed Switch privileges control the ability to perform tasks related to the management of Distributed Switch instances.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 285: vSphere Distributed Switch Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create	Allows creation of a distributed switch.	Data centers, Network folders	DVSwitch.Create
Delete	Allows removal of a distributed switch. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Distributed switches	DVSwitch.Delete
Host operation	Allows changing the host members of a distributed switch.	Distributed switches	DVSwitch.HostOp
Modify	Allows changing the configuration of a distributed switch.	Distributed switches	DVSwitch.Modify
Move	Allows moving a vSphere Distributed Switch to another folder.	Distributed switches	DVSwitch.Move
Network I/O control operation	Allow changing the resource settings for a vSphere Distributed Switch.	Distributed switches	DVSwitch.ResourceManagement
Policy operation	Allows changing the policy of a vSphere Distributed Switch.	Distributed switches	DVSwitch.PolicyOp
Port configuration operation	Allow changing the configuration of a port in a vSphere Distributed Switch.	Distributed switches	DVSwitch.PortConfig
Port setting operation	Allows changing the setting of a port in a vSphere Distributed Switch.	Distributed switches	DVSwitch.PortSetting
VSPAN operation	Allows changing the VSPAN configuration of a vSphere Distributed Switch.	Distributed switches	DVSwitch.Vspan

Datacenter Privileges

Datacenter privileges control the ability to create and edit data centers in the vSphere Client inventory.

All data center privileges are used in vCenter Server only. The **Create datacenter** privilege is defined on data center folders or the root object. All other data center privileges are pair with data centers, data center folders, or the root object.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 286: Datacenter Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create datacenter	Allows creation of new data center.	Data center folder or root object	Datacenter.Create
Move datacenter	Allows moving a data center. Privilege must be present at both the source and destination.	Data center, source and destination	Datacenter.Move

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Network protocol profile configuration	Allows configuration of the network profile for a data center.	Data center	Datacenter.IpPoolConfig
Query IP pool allocation	Allows configuration of a pool of IP addresses.	Data center	Datacenter.IpPoolQueryAllocations
Reconfigure datacenter	Allows reconfiguration of a data center.	Data center	Datacenter.Reconfigure
Release IP allocation	Allows releasing the assigned IP allocation for a data center.	Data center	Datacenter.IpPoolReleaseIp
Remove datacenter	Allows removal of a data center. To have permission to perform this operation, you must have this privilege assigned to both the object and its parent object.	Data center plus parent object	Datacenter.Delete
Rename datacenter	Allows changing the name of a data center.	Data center	Datacenter.Rename
Update datacenter carbon info	Allows gathering metrics related to energy and carbon measurement.	Data center	Datacenter.UpdateCarbonInfo

Datastore Privileges

Datastore privileges control the ability to browse, manage, and allocate space on datastores.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 287: Datastore Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Allocate space	Allows allocating space on a datastore for a virtual machine, snapshot, clone, or virtual disk.	Data stores	Datastore.AllocateSpace
Browse datastore	Allows browsing files on a datastore.	Data stores	Datastore.Browse
Configure datastore IO management	Allows configuring Storage I/O Control.	Data stores	Datastore.ConfigIOManagement
Configure datastore	Allows configuration of a datastore.	Data stores	Datastore.Config

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Low level file operations	Allows performing read, write, delete, and rename operations in the datastore browser.	Data stores	Datastore.FileManagement
Move datastore	Allows moving a datastore between folders. Privileges must be present at both the source and destination.	Datastore, source and destination	Datastore.Move
Remove datastore	Allows removal of a datastore. This privilege is deprecated. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Data stores	Datastore.Delete
Remove file	Allows deletion of files in the datastore. This privilege is deprecated. Assign the Low level file operations privilege.	Data stores	Datastore.DeleteFile
Rename datastore	Allows renaming a datastore.	Data stores	Datastore.Rename
Update virtual machine files	Allows updating file paths to virtual machine files on a datastore after the datastore has been resignatured.	Data stores	Datastore.UpdateVirtualMachineFiles
Update virtual machine metadata	Allows updating virtual machine metadata associated with a datastore.	Data stores	Datastore.UpdateVirtualMachineMetadata

Datastore Cluster Privileges

Datastore cluster privileges control the configuration of datastore clusters for Storage DRS.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 288: Datastore Cluster Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Configure a datastore cluster	Allows creation of and configuration of settings for datastore clusters for Storage DRS.	Datastore clusters	StoragePod.Config

ESX Agent Manager Privileges

ESX Agent Manager privileges control operations related to ESX Agent Manager and agent virtual machines. The ESX Agent Manager is a service that lets you install management virtual machines, which are tied to a host and not affected by VMware DRS or other services that migrate virtual machines.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 289: ESX Agent Manager

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Config	Allows deployment of an agent virtual machine on a host or cluster.	Virtual machines	EAM.Config
Modify	Allows modifications to an agent virtual machine such as powering off or deleting the virtual machine.	Virtual machines	EAM.Modify
View	Allows viewing of an agent virtual machine.	Virtual machines	EAM.View

Extension Privileges

Extension privileges control the ability to install and manage extensions.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 290: Extension Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Register extension	Allows registration of an extension (plug-in).	Root vCenter Server	Extension.Register
Unregister extension	Allows unregistering an extension (plug-in).	Root vCenter Server	Extension.Unregister
Update extension	Allows updates to an extension (plug-in).	Root vCenter Server	Extension.Update

External Stats Provider Privileges

External stats provider privileges control the ability to notify vCenter Server of Proactive Distributed Resource Scheduler (DRS) statistics.

These privileges apply to an API that is VMware-internal only.

Folder Privileges

Folder privileges control the ability to create and manage folders.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 291: Folder Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create folder	Allows creation of a new folder.	Folders	Folder.Create
Delete folder	Allows deletion of a folder. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Folders	Folder.Delete
Move folder	Allows moving a folder. Privilege must be present at both the source and destination.	Folders	Folder.Move
Rename folder	Allows changing the name of a folder.	Folders	Folder.Rename

Global Privileges

Global privileges control global tasks related to tasks, scripts, and extensions.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 292: Global Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Act as vCenter Server	Allows preparation or initiation of a vMotion send operation or a vMotion receive operation.	Root vCenter Server	Global.VCServer
Cancel task	Allows cancelation of a running or queued task.	Inventory object related to the task	Global.CancelTask
Capacity planning	Allows activating the use of capacity planning for planning consolidation of physical machines to virtual machines.	Root vCenter Server	Global.CapacityPlanning
Diagnostics	Allows retrieval of a list of diagnostic files, log header, binary files, or diagnostic bundle. To avoid potential security breaches, limit this privilege to the vCenter Server Administrator role.	Root vCenter Server	Global.Diagnostics
Disable methods	Allows servers for vCenter Server extensions to deactivate certain operations on objects managed by vCenter Server.	Root vCenter Server	Global.DisableMethods
Enable methods	Allows servers for vCenter Server extensions to activate certain operations on objects managed by vCenter Server.	Root vCenter Server	Global.EnableMethods
Global tag	Allows adding or removing global tags.	Root host or vCenter Server	Global.GlobalTag
Health	Allows viewing the health of vCenter Server components.	Root vCenter Server	Global.Health
Licenses	Allows viewing installed licenses and adding or removing licenses.	Root host or vCenter Server	Global.Licenses
Log event	Allows logging a user-defined event against a particular managed entity.	Any object	Global.LogEvent
Manage custom attributes	Allows adding, removing, or renaming custom field definitions.	Root vCenter Server	Global.ManageCustomFields
Proxy	Allows access to an internal interface for adding or removing endpoints to or from the proxy.	Root vCenter Server	Global.Proxy
Script action	Allows scheduling a scripted action along with an alarm.	Any object	Global.ScriptAction
Service managers	Allows use of the <code>resxtop</code> command in ESXCLI.	Root host or vCenter Server	Global.ServiceManagers
Set custom attribute	Allows viewing, creating, or removing custom attributes for a managed object.	Any object	Global.SetCustomField
Settings	Allows reading and modifying runtime vCenter Server configuration settings.	Root vCenter Server	Global.Settings
System tag	Allows adding or removing system tags.	Root vCenter Server	Global.SystemTag

Interact with the Guest Data Publisher Privileges

Interact with the Guest Data Publisher privileges control access to the published guest data on the host GDP service.

Table 293: Interact with the Guest Data Publisher Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Subscribe to the guest data publisher service on ESX hosts	Allows access to the published guest data on the host GDP service.	Hosts	GuestDataPublisher.GetData

Hybrid Linked Mode Privileges

Hybrid Linked Mode privileges control aspects of linking your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain. (Applies to VMware Cloud on AWS.)

Table 294: Hybrid Linked Mode Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create	Allows full administrative-level access for creating and deleting communities.	SDDC	HLM.Create
Manage	Allows creating trust for sources and accessing communities (read-level).	SDDC	HLM.Manage

Health Update Provider Privileges

Health update provider privileges control the ability for hardware vendors to notify vCenter Server of Proactive HA events.

These privileges apply to an API that is VMware-internal only.

Host CIM Privileges

Host CIM privileges control the use of CIM for host health monitoring.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 295: Host CIM Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
CIM — CIM Interaction	Allows a client to obtain a ticket to use for CIM services.	Hosts	Host.Cim.CimInteraction

Host Configuration Privileges

Host configuration privileges control the ability to configure hosts.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 296: Host Configuration Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Configuration — Advanced Settings	Allows setting advanced host configuration options.	Hosts	Host.Config.AdvancedConfig
• Configuration — Authentication Store	Allows configuring Active Directory authentication stores.	Hosts	Host.Config.AuthenticationStore
• Configuration — Change PciPassthru settings	Allows changes to PciPassthru settings for a host.	Hosts	Host.Config.PciPassthru
• Configuration — Change SNMP settings	Allows changes to SNMP settings for a host.	Hosts	Host.Config.Snmp
• Configuration — Change date and time settings	Allows changes to date and time settings on the host.	Hosts	Host.Config.DateTime
• Configuration — Change settings	Allows setting of lockdown mode on ESXi hosts.	Hosts	Host.Config.Settings
• Configuration — Connection	Allows changes to the connection status of a host (connected or disconnected).	Hosts	Host.Config.Connection
• Configuration — Firmware	Allows updates to the ESXi host's firmware.	Hosts	Host.Config.Firmware
• Configuration — GuestStore settings	Allows changes to the GuestStore.	GuestStore repository	Host.Config.GuestStore
• Configuration — Hyperthreading	Allows activating and deactivating hyperthreading in a host CPU scheduler.	Hosts	Host.Config.HyperThreading
• Configuration — Image configuration	Allows changes to the image associated with a host.		Host.Config.Image
• Configuration — Maintenance	Allows putting the host in and out of maintenance mode and shutting down and restarting the host.	Hosts	Host.Config.Maintenance
• Configuration — Memory configuration	Allows modifications to the host configuration.	Hosts	Host.Config.Memory
• Configuration — NVDIMM	Allows reading and configuring Non-Volatile DIMMs.	Hosts	Host.Config.Nvdim

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Configuration – Network configuration	Allows configuration of network, firewall, and vMotion network.	Hosts	Host.Config.Network
• Configuration – Power	Allows configuration of host power management settings.	Hosts	Host.Config.Power
• Configuration – ProductLocker settings	Allows configuration of the ESXi productlocker folder.	Hosts	Host.Config.ProductLocker
• Configuration – Quarantine	Allows putting a host into Quarantine mode.	Hosts	Host.Config.Quarantine
• Configuration – Query patch	Allows querying for installable patches and installing patches on the host.	Hosts	Host.Config.Patch
• Configuration – Security profile and firewall	Allows configuration of Internet services, such as SSH, Telnet, SNMP, and of the host firewall.	Hosts	Host.Config.NetService
• Configuration – Storage partition configuration	Allows VMFS datastore and diagnostic partition management. Users with this privilege can scan for new storage devices and manage iSCSI.	Hosts	Host.Config.Storage
• Configuration – System Management	Allows extensions to manipulate the file system on the host.	Hosts	Host.Config.SystemManagement
• Configuration – System resources	Allows updates to the configuration of the system resource hierarchy.	Hosts	Host.Config.Resources
• Configuration – Virtual machine autostart configuration	Allows changes to the auto-start and auto-stop order of virtual machines on a single host.	Hosts	Host.Config.AutoStart

Host Entropy Pool Privileges

Host Entropy Pool privileges control the ability to view and add ESXi host entropy.

Table 297: Host Entropy Pool Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Entropy Pool – Read	Allows reading the host entropy pool information.	Hosts	Host.Entropy.Read
• Entropy Pool – Write	Allows adding entropy to the host entropy pool.	Hosts	Host.Entropy.Write

Host Intel Software Guard Extensions Privileges

Host Intel Software Guard Extensions privileges control aspects of remote attestation on multi-socket ESXi hosts.

Table 298: Host Intel Software Guard Extensions (SGX) Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Intel Software Guard Extensions (SGX) Intel Software Guard Extensions (SGX) Register host	Allows registering of hosts with Intel SGX registration service (for SGX workloads to be able to perform remote attestation of SGX when running on multi-socket SGX capable hosts).	Hosts	Host.Sgx.Register

Host Inventory Privileges

Host inventory privileges control adding hosts to the inventory, adding hosts to clusters, and moving hosts in the inventory.

The table describes the privileges required to add and move hosts and clusters in the inventory.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 299: Host Inventory Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Inventory Add host to cluster	Allows addition of a host to an existing cluster.	Clusters	Host.Inventory.AddHostToCluster
Inventory Add standalone host	Allows addition of a standalone host.	Host folders	Host.Inventory.AddStandaloneHost
Inventory Create cluster	Allows creation of a new cluster.	Host folders	Host.Inventory.CreateCluster
Inventory Manage Cluster Lifecycle	Allows managing the cluster.	Clusters	Host.Inventory.ManageClusterLifecycle
Inventory Modify cluster	Allows changing the properties of a cluster.	Clusters	Host.Inventory.EditCluster
Inventory Move cluster or standalone host	Allows moving a cluster or standalone host between folders. Privilege must be present at both the source and destination.	Clusters	Host.Inventory.MoveCluster
Inventory Move host	Allows moving a set of existing hosts into or out of a cluster. Privilege must be present at both the source and destination.	Clusters	Host.Inventory.MoveHost

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Inventory Remove cluster	Allows deletion of a cluster or standalone host. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Clusters, Hosts	Host.Inventory.DeleteCluster
Inventory Remove host	Allows removal of a host. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Hosts plus parent object	Host.Inventory.RemoveHostFromCluster
Inventory Rename cluster	Allows renaming a a cluster.	Clusters	Host.Inventory.RenameCluster

Host Local Operations Privileges

Host local operations privileges control actions performed when the VMware Host Client is connected directly to a host.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 300: Host Local Operations Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Local operations Add host to vCenter	Allows installation and removal of vCenter agents, such as vpxa and aam, on a host.	Root host	Host.Local.InstallAgent
Local operations Create virtual machine	Allows creation of a new virtual machine from scratch on a disk without registering it on the host.	Root host	Host.Local.CreateVM
Local operations Delete virtual machine	Allows deletion of a virtual machine on disk. Supported for registered and unregistered virtual machines.	Root host	Host.Local.DeleteVM
Local operations Manage user groups	Allows management of local accounts on a host.	Root host	Host.Local.ManageUserGroups
Local operations Reconfigure virtual machine	Allows reconfiguring a virtual machine.	Root host	Host.Local.ReconfigVM

Host Statistics Privileges

Hosts Statistics privileges control the ability to access statistical information from a data processing unit (DPU).

These privileges apply to an API that is VMware-internal only.

Host Trusted Platform Module Privileges

Host Trusted Platform Module privileges control operations related to managing Trusted Platform Module (TPM) chips.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 301: Host Trusted Platform Module Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Trusted Platform Module – Read – Unseal	Read allows reading detailed information about the state of the TPM installed in the ESXi host. Unseal allows requesting an ESXi host to decrypt a challenge to prove its state.	Hosts	Host.Tpm.Read Host.Tpm.Unseal

Host vSphere Replication Privileges

Host vSphere replication privileges control the use of virtual machine replication by VMware vCenter Site Recovery Manager™ for a host.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 302: Host vSphere Replication Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• vSphere Replication – Manage Replication	Allows management of virtual machine replication on this host.	Hosts	Host.Hbr.HbrManagement

Host Profile Privileges

Host Profile privileges control operations related to creating and modifying host profiles.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 303: Host Profile Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Clear	Allows clearing of profile related information.	Root vCenter Server	Profile.Clear
Create	Allows creation of a host profile.	Root vCenter Server	Profile.Create

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Delete	Allows deletion of a host profile.	Root vCenter Server	Profile.Delete
Edit	Allows editing a host profile.	Root vCenter Server	Profile.Edit
Export	Allows exporting a host profile	Root vCenter Server	Profile.Export
View	Allows viewing a host profile.	Root vCenter Server	Profile.View

vCenter Server Profiles Privileges

vCenter Server Profiles privileges control aspects of listing profiles, and exporting and importing configurations from one vCenter Server to another.

Table 304: vCenter Server Profiles Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
vCenter Server Profiles Read Privileges	Allows listing and exporting of vCenter Server profiles.	vCenter Server	Infraprofile.Read
vCenter Server Profiles Write Privileges	Allows importing a profile to another vCenter Server and validating it.	vCenter Server	Infraprofile.Write

vSphere Namespaces Privileges

Namespaces privileges control who can create and manage VMware vSphere® with VMware Tanzu™ namespaces.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 305: Namespaces Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Allows disk decommission operations	Allows for decommissioning operations of data stores.	Data stores	Namespaces.ManageDisks
Backup Workloads component files	Allows for backing up the contents of the etcd cluster (used only in VMware Cloud on AWS).	Clusters	Namespaces.Backup
List accessible namespaces	Allows listing the accessible namespaces.	Clusters	Namespaces.ListAccess
Modify cluster-wide configuration	Allows modifying the cluster-wide configuration, and activating and deactivating cluster namespaces.	Clusters	Namespaces.Manage
Modify cluster-wide namespace self-service configuration	Allows modifying the namespace self-service configuration.	Clusters (for activating and deactivating)	Namespaces.SelfServiceManage

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
		Templates (for modifying the configuration) vCenter Server (for creating a template)	
Modify namespace configuration	Allows modifying namespace configuration options such as resource allocation and user permissions.	Clusters	Namespaces.Configure
Toggle cluster capabilities	Allows manipulating the state of cluster capabilities (used internally only for VMware Cloud on AWS).	Clusters	Namespaces.ManageCapabilities
Upgrade clusters to newer versions	Allows initiation of the cluster upgrade.	Clusters	Namespaces.Upgrade

Network Privileges

Network privileges control tasks related to network management.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 306: Network Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Assign network	Allows assigning a network to a virtual machine.	Networks, Virtual Machines	Network.Assign
Configure	Allows configuring a network.	Networks, Virtual Machines	Network.Config
Move network	Allows moving a network between folders. Privilege must be present at both the source and destination.	Networks	Network.Move
Remove	Allows removal of a network. This privilege is deprecated. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Networks	Network.Delete

NSX Privileges

NSX privileges control tasks related to NSX management.

Table 307: NSX Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Read NSX configuration	Allows reading of NSX objects.	NSX	Nsx.Read
Manage NSX configuration	Allows managing of NSX objects from a vSphere administrator's perspective.	NSX	Nsx.Manage
Modify NSX configuration	Allows managing of NSX objects from an enterprise administrator's perspective.	NSX	Nsx.ModifyAll

VMware Observability Privileges

VMware Observability privileges control the ability for an agent to access the observability APIs on vCenter Server.

These privileges apply to an API that is VMware-internal only.

OvfManager Privileges

OvfManager privileges control the ability to access vService Manager.

These privileges apply to an API that is VMware-internal only.

Interact with Partner Rest Daemons Privileges

Interact with Partner Rest Daemons privileges control access to read and write operations.

Table 308: Interact with Partner Rest Daemons Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Perform a GET operation from a Partner's Rest Daemon	Allows the partner-provisioned REST client to perform GET operations.	The partner's user that performs GET operations.	PartnerRestDaemon.Read
Perform a modifying operation to a Partner's Rest Daemon	Allows the partner-provisioned REST client to perform POST, PUT, and DELETE operations.	The partner's user that performs POST, PUT, or DELETE operations.	PartnerRestDaemon.Write

Performance Privileges

Performance privileges control modifying performance statistics settings.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 309: Performance Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Modify intervals	Allows creating, removing, and updating performance data collection intervals.	Root vCenter Server	Performance.ModifyIntervals

Plug-in Privileges

Plug-in privileges control the managing of vSphere Client plug-ins.

Table 310: Plug-in Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Manage Plug-ins	Allows for managing vSphere Client plug-ins.	vCenter Server	Plugin.Management

Replication as a Service Privileges

Replication as a Service privileges control access to various internal APIs and functionality related to vCenter Server linking.

These privileges apply to an API that is VMware-internal only.

Permissions Privileges

Permissions privileges control the assigning of roles and permissions.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 311: Permissions Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Modify permission	Allows defining one or more permission rules on an entity, or updating rules if rules are already present for the given user or group on the entity. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Any object plus parent object	Authorization.ModifyPermissions
Modify privilege	Allows modifying a privilege's group or description. No vSphere Client user interface elements are associated with this privilege.	Any object	Authorization.ModifyPrivileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Modify role	Allows updating a role's name and the privileges that are associated with the role.	Any object	Authorization.ModifyRoles
Modify vTContainer	Allows creating, updating, and deleting vTContainer instances.	vTContainer objects	Authorization.ModifyVTContainers
Modify vTContainer mappings	Allows creating and deleting a vTContainer mapping.	vTContainer mapping objects	Authorization.ModifyVTContainerMappings
Reassign role permissions	Allows reassigning all permissions of a role to another role.	Any object	Authorization.ReassignRolePermissions

VM Storage Policies Privileges

VM storage privileges control operations related to storage profiles.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 312: VM Storage Privileges

Privilege Name	Description	Required On	Privilege Name in the API
Update VM storage policies	Allows changes to be made to storage profiles, such as creating and updating storage capabilities and virtual machine storage profiles.	Root vCenter Server	StorageProfile.Update
View VM storage policies	Allows using and viewing defined storage capabilities and storage profiles.	Root vCenter Server	StorageProfile.View

Resource Privileges

Resource privileges control the creation and management of resource pools, as well as the migration of virtual machines.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 313: Resource Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Apply recommendation	Allows accepting a suggestion by the server to perform a migration with vMotion.	Clusters	Resource.ApplyRecommendation
Assign vApp to resource pool	Allows assignment of a vApp to a resource pool.	Resource pools	Resource.AssignVAppToPool
Assign virtual machine to resource pool	Allows assignment of a virtual machine to a resource pool.	Resource pools	Resource.AssignVMToPool

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create resource pool	Allows creation of resource pools.	Resource pools, clusters	Resource.CreatePool
Migrate powered off virtual machine	Allows migration of a powered off virtual machine to a different resource pool or host.	Virtual machines	Resource.ColdMigrate
Migrate powered on virtual machine	Allows migration with vMotion of a powered on virtual machine to a different resource pool or host.	Virtual machines	Resource.HotMigrate
Modify resource pool	Allows changes to the allocations of a resource pool.	Resource pools	Resource.EditPool
Move resource pool	Allows moving a resource pool. Privilege must be present at both the source and destination.	Resource pools	Resource.MovePool
Query vMotion	Allows querying the general vMotion compatibility of a virtual machine with a set of hosts.	Root vCenter Server	Resource.QueryVMotion
Remove resource pool	Allows deletion of a resource pool. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Resource pools	Resource.DeletePool
Rename resource pool	Allows renaming of a resource pool.	Resource pools	Resource.RenamePool

Scheduled Task Privileges

Scheduled task privileges control creation, editing, and removal of scheduled tasks.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 314: Scheduled Task Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create tasks	Allows scheduling of a task. Required in addition to the privileges to perform the scheduled action at the time of scheduling.	Any object	ScheduledTask.Create
Modify task	Allows reconfiguration of the scheduled task properties.	Any object	ScheduledTask.Edit
Remove task	Allows removal of a scheduled task from the queue.	Any object	ScheduledTask.Delete

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Run task	Allows running the scheduled task immediately. Creating and running a scheduled task also requires permission to perform the associated action.	Any object	ScheduledTask.Run

Sessions Privileges

Sessions privileges control the ability of extensions to open sessions on the vCenter Server system.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Note: Assign Sessions privileges only to administrators or trusted users.

Table 315: Session Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Impersonate user	Allows impersonation of another user. This capability is used by extensions.	Root vCenter Server	Sessions.ImpersonateUser
Message	Allows setting of the global login message.	Root vCenter Server	Sessions.GlobalMessage
Validate session	Allows verification of session validity.	Root vCenter Server	Sessions.ValidateSession
View and stop sessions	Allows viewing sessions and forcing log out of one or more logged-on users.	Root vCenter Server	Sessions.TerminateSession
privilege.StorageProfile.ViewPermissions	Allows viewing of sessions.	Root vCenter Server	Sessions.CollectPrivilegeChecks

Storage Views Privileges

Storage Views privileges control privileges for Storage Monitoring Service APIs.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 316: Storage Views Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Configure service	Allows privileged users to use all Storage Monitoring Service APIs. Use Storage views > View for privileges to read-only Storage Monitoring Service APIs.	Root vCenter Server	StorageViews.ConfigureService
View	Allows privileged users to use read-only Storage Monitoring Service APIs.	Root vCenter Server	StorageViews.View

Supervisor Services Privileges

Supervisor Services privileges control who can create and manage Supervisor Services on the vSphere with Tanzu environment.

Table 317: Supervisor Services Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Manage Supervisor Services	Allows for creating, updating, or deleting a Supervisor Service. Also allows for installing a Supervisor Service on a cluster, and creating or deleting a Supervisor Service version.	Clusters	SupervisorServices.Manage

Tasks Privileges

Tasks privileges control the ability of extensions to create and update tasks on the vCenter Server.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 318: Tasks Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create task	Allows an extension to create a user-defined task. No vSphere Client user interface elements are associated with this privilege.	Root vCenter Server	Task.Create
Update task	Allows an extension to update a user-defined task.	Root vCenter Server	Task.Update

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
	No vSphere Client user interface elements are associated with this privilege.		

Tenant Management Privileges

Tenant Management privileges control aspects of defining and retrieving tenant management entities. (Applies to VMware Cloud on AWS.)

Table 319: Tenant Management Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Tenant provisioning operations	Allows defining a set of resources to use for tenant management.	Root folder and each entity currently marked as a service provider.	TenantManager.Update
Tenant query operations	Allows retrieving the list of tenant management resources.	Root folder and each entity currently marked as a service provider.	TenantManager.Query

Transfer Service Privileges

Transfer service privileges are VMware internal. Do not use these privileges.

VcTrusts/VcIdentity Privileges

VcTrusts/VcIdentity privileges control access to various internal APIs and functionality related to trust between vCenter Server systems.

Table 320: VcTrusts/VcIdentity Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create/Update/Delete (Admin privs)	Allows full administrative-level access to various internal APIs and functionality related to trust between vCenter Server systems.	N/A	Trust.Administer
Create/Update/Delete (below Admin privs)	Allows reduced administrative access to various internal APIs and functionality related to trust between vCenter Server systems. This privilege restricts creating/updating/deleting VcTrusts/VcIdentity so that the user cannot escalate non-administrator privileges.	N/A	Trust.Manage

Trusted Infrastructure Administrator Privileges

Trusted Infrastructure administrator privileges configure and manage a vSphere Trust Authority deployment.

These privileges determine who can perform configuration and management tasks for a vSphere Trust Authority deployment. See [Prerequisites and Required Privileges for vSphere Trust Authority](#) for more information about the Trust Authority roles and the TrustedAdmins group.

Table 321: Trusted Infrastructure Administrator Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Configure Key Server Trust	Allows managing the Key Providers of the Key Provider Service.	Root vCenter Server	TrustedAdmin.ManageKMSTrust
Configure Trust Authority Host TPM certificates	Allows creation and modification of the Attestation Service settings.	Root vCenter Server	TrustedAdmin.ConfigureHostCertificates
Configure Trust Authority Host metadata	Allows editing the base images to be attested by the Attestation Service.	Root vCenter Server	TrustedAdmin.ConfigureHostMetadata
Configure attesting SSO	Allows editing which hosts can be trusted by the Trust Authority Hosts.	Root vCenter Server	TrustedAdmin.ManageAttestingSSO
Configure token conversion policy	Allows configuring the token conversion policy.	Root vCenter Server	TrustedAdmin.ConfigureTokenConversionPolicy
List Trusted Infrastructure Hosts	Allows reading information regarding the Trusted Hosts and the Trust Authority Hosts.	Root vCenter Server	TrustedAdmin.ReadTrustedHosts
List information about the STS	Allows exporting the Trusted Host details, so that they can be imported to the Trust Authority Cluster.	Root vCenter Server	TrustedAdmin.ReadStsInfo
Manage Trusted Infrastructure Hosts	Allows editing the information regarding the Trusted Hosts and the Trust Authority Hosts.	Root vCenter Server	TrustedAdmin.ManageTrustedHosts
Read Key Server Trust	Allows reading the Key Providers of the Key Provider Service.	Root vCenter Server	TrustedAdmin.ReadKMSTrust
Read attesting SSO	Allows reading which hosts can be trusted by the Trust Authority Hosts.	Root vCenter Server	TrustedAdmin.ReadAttestingSSO

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Retrieve TPM Trust Authority Host certificates	Allows reading the settings of the Attestation Service.	Root vCenter Server	TrustedAdmin.RetrieveTPMHostCertificates
Retrieve Trust Authority Host metadata	Allows reading which base images can be attested by the Attestation Service.	Root vCenter Server	TrustedAdmin.RetrieveHostMetadata

vApp Privileges

vApp privileges control operations related to deploying and configuring a vApp.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 322: vApp Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Add virtual machine	Allows adding a virtual machine to a vApp.	vApps	VApp.AssignVM
Assign resource pool	Allows assigning a resource pool to a vApp.	vApps	VApp.AssignResourcePool
Assign vApp	Allows assigning a vApp to another vApp	vApps	VApp.AssignVApp
Clone	Allows cloning of a vApp.	vApps	VApp.Clone
Create	Allows creation of a vApp.	vApps	VApp.Create
Delete	Allows deletion a vApp. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	vApps	VApp.Delete
Export	Allows export of a vApp from vSphere.	vApps	VApp.Export
Import	Allows import of a vApp into vSphere.	vApps	VApp.Import
Move	Allows moving a vApp to a new inventory location.	vApps	VApp.Move
Power Off	Allows power off operations on a vApp.	vApps	VApp.PowerOff
Power On	Allows power on operations on a vApp.	vApps	VApp.PowerOn
Pull from URL	Allows listing of remote source file descriptors.	vApps	VApp.PullFromUrls
Rename	Allows renaming a vApp.	vApps	VApp.Rename
Suspend	Allows suspension of a vApp.	vApps	VApp.Suspend

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Unregister	Allows unregistering a vApp. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	vApps	VApp.Unregister
View OVF Environment	Allows viewing the OVF environment of a powered-on virtual machine within a vApp.	vApps	VApp.ExtractOvfEnvironment
vApp application configuration	Allows modification of a vApp's internal structure, such as product information and properties.	vApps	VApp.ApplicationConfig
vApp instance configuration	Allows modification of a vApp's instance configuration, such as policies.	vApps	VApp.InstanceConfig
vApp managedBy configuration	Allows an extension or solution to mark a vApp as being managed by that extension or solution. No vSphere Client user interface elements are associated with this privilege.	vApps	VApp.ManagedByConfig
vApp resource configuration	Allows modification of a vApp's resource configuration. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	vApps	VApp.ResourceConfig

VcIdentityProviders Privileges

VcIdentityProviders privileges control access to the VcIdentityProviders API.

Table 323: VcIdentityProviders Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create	Allows create-only access to the VcIdentityProviders API (vCenter Server identity providers).	N/A	VcIdentityProviders.Create
Manage	Allows administrative-level write access (create, read, update, delete) to the VcIdentityProviders API (vCenter Server identity providers).	N/A	VcIdentityProviders.Manage

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Read	Allows read access to the VcIdentityProviders API (vCenter Server identity providers).	N/A	VcIdentityProviders.Read

VMware vSphere Lifecycle Manager Configuration Privileges

VMware vSphere Lifecycle Manager configuration privileges control the ability to configure the vSphere Lifecycle Manager service.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Note: Assign privileges that authorize users to invoke VMware vSphere Lifecycle Manager APIs that accept URLs only to administrators or trusted users.

Table 324: VMware vSphere Lifecycle Manager Configuration Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Configure Configure Service	Allows configuring the vSphere Lifecycle Manager service and the scheduled patch download task.	Root vCenter Server	VcIntegrity.General.com.vmware.vcIntegrity.Configure

VMware vSphere Lifecycle Manager Manager Desired Configuration Management Privileges

VMware vSphere Lifecycle Manager Desired Configuration Management privileges control the ability to manage the vSphere Lifecycle Manager configuration.

Table 325: VMware vSphere Lifecycle Manager Desired Configuration Management Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Desired Configuration Management Privileges – Export desired cluster configuration – Modify desired cluster configuration – Read-only access to desired configuration management platform – Remediate cluster to the desired configuration	Export desired cluster configuration allows exporting the configuration or the configuration schema. Modify desired cluster configuration allows importing a configuration or extracting the configuration from a reference host. Read-only access to desired configuration management platform allows checking compliance, running the remediation pre-check, viewing compliance, and viewing pre-check results. Remediate cluster to the desired configuration allows remediating a cluster and transitioning to vSphere Configuration Profiles.	Root vCenter Server	VcIntegrity.ClusterConfiguration.Export VcIntegrity.ClusterConfiguration.Modify VcIntegrity.ClusterConfiguration.View VcIntegrity.ClusterConfiguration.Remediate

VMware vSphere Lifecycle Manager ESXi Health Perspectives Privileges

VMware vSphere Lifecycle Manager ESXi health perspective privileges control the ability to check the health of ESXi hosts and clusters.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 326: VMware vSphere Lifecycle Manager ESXi Health Perspectives Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• ESXi Health Perspectives – Read – Write	Read allows querying the health of ESXi hosts and clusters. Write is currently not used.	Hosts Clusters	VcIntegrity.lifecycleHealth.Read VcIntegrity.lifecycleHealth.Write

VMware vSphere Lifecycle Manager Depots Privileges

VMware vSphere Lifecycle Manager Depots Privileges control the ability to manage depots.

Table 327: VMware vSphere Lifecycle Manager Depots Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Lifecycle Manager: Depots Privileges Delete	Allows deleting a vSphere Lifecycle Manager depot.	Root vCenter Server	VcIntegrity.lifecycleDepots.Delete

VMware vSphere Lifecycle Manager General Privileges

VMware vSphere Lifecycle Manager General privileges control the ability to read and write Lifecycle Manager resources.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 328: VMware vSphere Lifecycle Manager General Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Lifecycle Manager: General Privileges Read Write	Read allows reading of the vSphere Lifecycle Manager resources. This privilege is required to get task information. Write allows writing of the vSphere Lifecycle Manager resources. This privilege is required to cancel a vSphere Lifecycle Manager task.	Root vCenter Server	VcIntegrity.lifecycleGeneral.Read VcIntegrity.lifecycleGeneral.Write

VMware vSphere Lifecycle Manager Hardware Compatibility Privileges

VMware vSphere Lifecycle Manager Hardware Compatibility privileges control the ability to discover and resolve potential hardware compatibility issues.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 329: VMware vSphere Lifecycle Manager Hardware Compatibility Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Lifecycle Manager: Hardware Compatibility Privileges Access Hardware Compatibility Write	Access Hardware Compatibility and Write allow accessing the hardware compatibility data and resolving potential hardware compatibility issues.	Hosts	VcIntegrity.HardwareCompatibility.Read VcIntegrity.HardwareCompatibility.Write

VMware vSphere Lifecycle Manager Image Privileges

VMware vSphere Lifecycle Manager Image privileges control the ability to manage images.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Note: Assign privileges that authorize users to invoke VMware vSphere Lifecycle Manager APIs that accept URLs only to administrators or trusted users.

Table 330: VMware vSphere Lifecycle Manager Image Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Lifecycle Manager: Image Privileges – Read – Write	Read allows reading of vSphere Lifecycle Manager images. This privilege is required to: • List all the drafts for a cluster • Get more information on a draft • Perform a scan on a draft • Validate a draft • Retrieve the contents of a draft • Compute the effective component list • Get the contents of the current desired state document • Start a scan on a cluster • Get the compliance result • Get a recommendation • Export the current desired state as a depot, JSON file, or ISO Write allows managing of vSphere Lifecycle Manager images. This privilege is required to: • Create, delete, or commit a draft • Import the desired state • Generate recommendations • Set or delete different portions of a draft	Root vCenter Server	VcIntegrity.lifecycleSettings.Read VcIntegrity.lifecycleSettings.Write

VMware vSphere Lifecycle Manager Image Remediation Privileges

VMware vSphere Lifecycle Manager Image privileges control the ability to remediate images.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 331: VMware vSphere Lifecycle Manager Image Remediation Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Lifecycle Manager: Image Remediation Privileges – Read – Write	Read allows performing the remediation pre-check. Write allows performing the remediation.	Clusters	VcIntegrity.lifecycleSoftwareRemediation.Read VcIntegrity.lifecycleSoftwareRemediation.Write

VMware vSphere Lifecycle Manager Settings Privileges

VMware vSphere Lifecycle Manager Settings privileges control the ability to manage depots and remediation policies.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Note: Assign privileges that authorize users to invoke VMware vSphere Lifecycle Manager APIs that accept URLs only to administrators or trusted users.

Table 332: VMware vSphere Lifecycle Manager Settings Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Lifecycle Manager: Settings Privileges – Read – Write	Read allows reading of vSphere Lifecycle Manager depots and remediation policies. Write allows writing of vSphere Lifecycle Manager depots and remediation policies.	Root vCenter Server	VcIntegrity.lifecycleSoftwareSpecification.Read VcIntegrity.lifecycleSoftwareSpecification.Write

VMware vSphere Lifecycle Manager Manage Baseline Privileges

VMware vSphere Lifecycle Manager Manage Baseline privileges control the ability to manage baselines and baseline groups.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 333: VMware vSphere Lifecycle Manager Manage Baseline Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Manage Baseline – Attach Baseline – Manage Baseline	Attach Baseline allows attaching baselines and baseline groups to objects in the vSphere inventory. Manage Baseline allows creating, editing, or deleting baselines and baseline groups.	Root vCenter Server	VcIntegrity.Baseline.com.vmware.vcIntegrity.AssignBaselines VcIntegrity.Baseline.com.vmware.vcIntegrity.ManageBaselines

VMware vSphere Lifecycle Manager Manage Patches and Upgrades Privileges

VMware vSphere Lifecycle Manager Manage Patches and Upgrades privileges control the ability to view, scan, and remediate applicable patches, extensions, or upgrades.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 334: VMware vSphere Lifecycle Manager Manage Patches and Upgrades Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Manage Patches and Upgrades – Remediate to Apply Patches, Extensions, and Upgrades – Scan for Applicable Patches, Extensions, and Upgrades – Stage Patches and Extensions – View Compliance Status	Remediate to Apply Patches, Extensions, and Upgrades allows remediation of virtual machines and hosts to apply patches, extensions, or upgrades when you are using baselines. In addition, this privilege allows viewing the compliance status. Scan for Applicable Patches, Extensions, and Upgrades allows scanning virtual machines and hosts to search for applicable patches, extensions, or upgrades when you are using baselines. Stage Patches and Extensions allows staging patches or extensions to ESXi hosts when you are using baselines. In addition, this privilege allows viewing the compliance status of ESXi hosts. View Compliance Status allows viewing the baseline compliance information for an object in the vSphere inventory.	Root vCenter Server	VcIntegrity.Updates.com.vmware.vcIntegrity.Remediate VcIntegrity.Updates.com.vmware.vcIntegrity.Scan VcIntegrity.Updates.com.vmware.vcIntegrity.Stage VcIntegrity.Updates.com.vmware.vcIntegrity.ViewStatus

VMware vSphere Lifecycle Manager Upload File Privileges

VMware vSphere Lifecycle Manager Upload File privileges control the ability to import updates to the vSphere Lifecycle Manager depot.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Note: Assign privileges that authorize users to invoke VMware vSphere Lifecycle Manager APIs that accept URLs only to administrators or trusted users.

Table 335: VMware vSphere Lifecycle Manager Upload File Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Upload file – Upload upgrade images and offline bundles	Allows uploading upgrade ISO and offline patch bundles.	Root vCenter Server	VcLifecycle.Upgrade

Virtual Machine Change Configuration Privileges

Virtual Machine Change Configuration privileges control the ability to configure virtual machine options and devices.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 336: Virtual Machine Change Configuration Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Change Configuration – Acquire disk lease	Allows disk lease operations for a virtual machine.	Virtual machines	VirtualMachine.Config.DiskLease
• Change Configuration – Add existing disk	Allows adding an existing virtual disk to a virtual machine.	Virtual machines	VirtualMachine.Config.AddExistingDisk
• Change Configuration – Add new disk	Allows creation of a new virtual disk to add to a virtual machine.	Virtual machines	VirtualMachine.Config.AddNewDisk
• Change Configuration – Add or remove device	Allows addition or removal of any non-disk device.	Virtual machines	VirtualMachine.Config.AddRemoveDevice
• Change Configuration – Advanced configuration	Allows addition or modification of advanced parameters in the virtual machine's configuration file.	Virtual machines	VirtualMachine.Config.AdvancedConfig
• Change Configuration – Change CPU count	Allows changing the number of virtual CPUs.	Virtual machines	VirtualMachine.Config.CPUCount
• Change Configuration – Change Memory	Allows changing the amount of memory allocated to the virtual machine.	Virtual machines	VirtualMachine.Config.Memory

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Change Configuration — Change Settings	Allows changing general virtual machine settings.	Virtual machines	VirtualMachine.Config.Settings
• Change Configuration — Change Swapfile placement	Allows changing the swapfile placement policy for a virtual machine.	Virtual machines	VirtualMachine.Config.SwapPlacement
• Change Configuration — Change resource	Allows changing the resource configuration of a set of virtual machine nodes in a given resource pool.	Virtual machines	VirtualMachine.Config.Resource
• Change Configuration — Configure Host USB device	Allows attaching a host-based USB device to a virtual machine.	Virtual machines	VirtualMachine.Config.HostUSBDevice
• Change Configuration — Configure Raw device	Allows adding or removing a raw disk mapping or SCSI pass through device. Setting this parameter overrides any other privilege for modifying raw devices, including connection states.	Virtual machines	VirtualMachine.Config.RawDevice
• Change Configuration — Configure managedBy	Allows an extension or solution to mark a virtual machine as being managed by that extension or solution.	Virtual machines	VirtualMachine.Config.ManagedBy
• Change Configuration — Display connection settings	Allows configuration of virtual machine remote console options.	Virtual machines	VirtualMachine.Config.MksControl
• Change Configuration — Extend virtual disk	Allows expansion of the size of a virtual disk.	Virtual machines	VirtualMachine.Config.DiskExtend
• Change Configuration — Modify device settings	Allows changing the properties of an existing device.	Virtual machines	VirtualMachine.Config.EditDevice
• Change Configuration — Query Fault Tolerance compatibility	Allows checking if a virtual machine is compatible for Fault Tolerance.	Virtual machines	VirtualMachine.Config.QueryFTCompatibility
• Change Configuration — Query unowned files	Allows querying of unowned files.	Virtual machines	VirtualMachine.Config.QueryUnownedFiles

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Change Configuration — Reload from path	Allows changing a virtual machine configuration path while preserving the identity of the virtual machine. Solutions such as VMware vCenter Site Recovery Manager use this operation to maintain virtual machine identity during failover and failback.	Virtual machines	VirtualMachine.Config.ReloadFromPath
• Change Configuration — Remove disk	Allows removal of a virtual disk device.	Virtual machines	VirtualMachine.Config.RemoveDisk
• Change Configuration — Rename	Allows renaming a virtual machine or modifying the associated notes of a virtual machine.	Virtual machines	VirtualMachine.Config.Rename
• Change Configuration — Reset guest information	Allows editing the guest operating system information for a virtual machine.	Virtual machines	VirtualMachine.Config.ResetGuestInfo
• Change Configuration — Set annotation	Allows adding or editing a virtual machine annotation.	Virtual machines	VirtualMachine.Config.Annotation
• Change Configuration — Toggle disk change tracking	Allows activating or deactivating of change tracking for the virtual machine's disks.	Virtual machines	VirtualMachine.Config.ChangeTracking
• Change Configuration — Toggle fork parent	Allows activating or deactivating a vmfork parent.	Virtual machines	VirtualMachine.Config.ToggleForkParent
• Change Configuration — Upgrade virtual machine compatibility	Allows upgrade of the virtual machine's virtual machine compatibility version.	Virtual machines	VirtualMachine.Config.UpgradeVirtualHardware

Virtual Machine Guest Operations Privileges

Virtual Machine Guest Operations privileges control the ability to interact with files and applications inside a virtual machine's guest operating system with the API.

See the *vSphere Web Services API Reference* documentation for more information on these operations.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 337: Virtual Machine Guest Operations

Privilege Name in the vSphere Client	Description	Effective on Object	Privilege Name in the API
Guest operations Guest operation alias modification	Allows virtual machine guest operations that involve modifying the alias for the virtual machine.	Virtual machines	VirtualMachine.GuestOperations.ModifyAliases
Guest operations Guest operation alias query	Allows virtual machine guest operations that involve querying the alias for the virtual machine.	Virtual machines	VirtualMachine.GuestOperations.QueryAliases
Guest operations Guest operation modifications	Allows virtual machine guest operations that involve modifications to a guest operating system in a virtual machine, such as transferring a file to the virtual machine. No vSphere Client user interface elements are associated with this privilege.	Virtual machines	VirtualMachine.GuestOperations.Modify
Guest operations Guest operation program execution	Allows virtual machine guest operations that involve running an application in the virtual machine. No vSphere Client user interface elements are associated with this privilege.	Virtual machines	VirtualMachine.GuestOperations.Execute
Guest operations Guest Operation Queries	Allows virtual machine guest operations that involve querying the guest operating system, such as listing files in the guest operating system. No vSphere Client user interface elements are associated with this privilege.	Virtual machines	VirtualMachine.GuestOperations.Query

Virtual Machine Interaction Privileges

Virtual Machine Interaction privileges control the ability to interact with a virtual machine console, configure media, perform power operations, and install VMware Tools.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 338: Virtual Machine Interaction

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
- Interaction - Answer question	Allows resolution of issues with virtual machine state transitions or runtime errors.	Virtual machines	VirtualMachine.Interact.AnswerQuestion
- Interaction - Backup operation on virtual machine	Allows performance of backup operations on virtual machines.	Virtual machines	VirtualMachine.Interact.Backup
- Interaction - Configure CD media	Allows configuration of a virtual DVD or CD-ROM device.	Virtual machines	VirtualMachine.Interact.SetCDMedia
- Interaction - Configure floppy media	Allows configuration of a virtual floppy device.	Virtual machines	VirtualMachine.Interact.SetFloppyMedia
- Interaction - Console interaction	Allows interaction with the virtual machine's virtual mouse, keyboard, and screen.	Virtual machines	VirtualMachine.Interact.ConsoleInteract
- Interaction - Create screenshot	Allows creation of a virtual machine screen shot.	Virtual machines	VirtualMachine.Interact.CreateScreenshot
- Interaction - Defragment all disks	Allows defragment operations on all disks of the virtual machine.	Virtual machines	VirtualMachine.Interact.DefragmentAllDisks
- Interaction - Device connection	Allows changing the connected state of a virtual machine's disconnectable virtual devices.	Virtual machines	VirtualMachine.Interact.DeviceConnection
- Interaction - Drag and Drop	Allows drag and drop of files between a virtual machine and a remote client.	Virtual machines	VirtualMachine.Interact.DnD
- Interaction - Guest operating system management by VIX API	Allows management of the virtual machine's operating system through the VIX API.	Virtual machines	VirtualMachine.Interact.GuestControl
- Interaction - Inject USB HID scan codes	Allows injection of USB HID scan codes.	Virtual machines	VirtualMachine.Interact.PutUsbScanCodes
- Interaction - Pause or Unpause	Allows pausing or unpausing of the virtual machine.	Virtual machines	VirtualMachine.Interact.Pause
- Interaction - Perform wipe or shrink operations	Allows performing wipe or shrink operations on the virtual machine.	Virtual machines	VirtualMachine.Interact.SESparseMaintenance

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Interaction — Power Off	Allows powering off a powered-on virtual machine. This operation powers down the guest operating system.	Virtual machines	VirtualMachine.Interact.PowerOff
• Interaction — Power On	Allows powering on a powered-off virtual machine, and resuming a suspended virtual machine.	Virtual machines	VirtualMachine.Interact.PowerOn
• Interaction — Record session on Virtual Machine	Allows recording a session on a virtual machine.	Virtual machines	VirtualMachine.Interact.Record
• Interaction — Replay session on Virtual Machine	Allows replaying of a recorded session on a virtual machine.	Virtual machines	VirtualMachine.Interact.Replay
• Interaction — Reset	Allows resetting of a virtual machine and reboots the guest operating system.	Virtual machines	VirtualMachine.Interact.Reset
• Interaction — Resume Fault Tolerance	Allows resuming of fault tolerance for a virtual machine.	Virtual machines	VirtualMachine.Interact.EnableSecondary
• Interaction — Suspend	Allows suspending a powered-on virtual machine. This operation puts the guest in standby mode.	Virtual machines	VirtualMachine.Interact.Suspend
• Interaction — Suspend Fault Tolerance	Allows suspension of fault tolerance for a virtual machine.	Virtual machines	VirtualMachine.Interact.DisableSecondary
• Interaction — Suspend to memory	Allows suspension of memory for a virtual machine.	Virtual machines	VirtualMachine.Interact.SuspendToMemory
• Interaction — Test failover	Allows testing of Fault Tolerance failover by making the Secondary virtual machine the Primary virtual machine.	Virtual machines	VirtualMachine.Interact.MakePrimary
• Interaction — Test restart Secondary VM	Allows termination of a Secondary virtual machine for a virtual machine using Fault Tolerance.	Virtual machines	VirtualMachine.Interact.DisableSecondary
• Interaction — Turn Off Fault Tolerance	Allows turning off Fault Tolerance for a virtual machine.	Virtual machines	VirtualMachine.Interact.TurnOffFaultTolerance
• Interaction — Turn On Fault Tolerance	Allows turning on Fault Tolerance for a virtual machine.	Virtual machines	VirtualMachine.Interact.CreateSecondary

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Interaction VMware Tools install	Allows mounting and unmounting the VMware Tools CD installer as a CD-ROM for the guest operating system.	Virtual machines	VirtualMachine.Interact.ToolsInstall

Virtual Machine Edit Inventory Privileges

Virtual Machine Edit Inventory privileges control adding, moving, and removing virtual machines.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 339: Virtual Machine Edit Inventory Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Edit inventory Create from existing	Allows creation of a virtual machine based on an existing virtual machine or template, by cloning or deploying from a template.	Clusters, Hosts, Virtual machine folders	VirtualMachine.Inventory.CreateFromExisting
Edit inventory Create new	Allows creation of a virtual machine and allocation of resources for its execution.	Clusters, Hosts, Virtual machine folders	VirtualMachine.Inventory.Create
Edit inventory Move	Allows relocating a virtual machine in the hierarchy. The privilege must be present at both the source and destination.	Virtual machines	VirtualMachine.Inventory.Move
Edit inventory Register	Allows adding an existing virtual machine to a vCenter Server or host inventory.	Clusters, Hosts, Virtual machine folders	VirtualMachine.Inventory.Register
Edit inventory Remove	Allows deletion of a virtual machine. Deletion removes the virtual machine's underlying files from disk. To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.	Virtual machines	VirtualMachine.Inventory.Delete
Edit inventory Unregister	Allows unregistering a virtual machine from a vCenter Server or host inventory.	Virtual machines	VirtualMachine.Inventory.Unregister

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
	To have permission to perform this operation, a user or group must have this privilege assigned in both the object and its parent object.		

Virtual Machine Provisioning Privileges

Virtual Machine Provisioning privileges control activities related to deploying and customizing virtual machines.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 340: Virtual Machine Provisioning Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Provisioning Allow disk access	Allows opening a disk on a virtual machine for random read and write access. Used mostly for remote disk mounting.	Virtual machines	VirtualMachine.Provisioning.DiskRandomAccess
Provisioning Allow file access	Allows operations on files associated with a virtual machine, including vmx, disks, logs, and nvram.	Virtual machines	VirtualMachine.Provisioning.FileRandomAccess
Provisioning Allow read-only disk access	Allows opening a disk on a virtual machine for random read access. Used mostly for remote disk mounting.	Virtual machines	VirtualMachine.Provisioning.DiskRandomRead
Provisioning Allow virtual machine download	Allows read operations on files associated with a virtual machine, including vmx, disks, logs, and nvram.	Root host or vCenter Server	VirtualMachine.Provisioning.GetVmFiles
Provisioning Allow virtual machine files upload	Allows write operations on files associated with a virtual machine, including vmx, disks, logs, and nvram.	Root host or vCenter Server	VirtualMachine.Provisioning.PutVmFiles
Provisioning Clone template	Allows cloning of a template.	Templates	VirtualMachine.Provisioning.CloneTemplate

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Provisioning Clone virtual machine	Allows cloning of an existing virtual machine and allocation of resources.	Virtual machines	VirtualMachine.Provisioning.Clone
Provisioning Create template from virtual machine	Allows creation of a new template from a virtual machine.	Virtual machines	VirtualMachine.Provisioning.CreateTemplateFromVM
Provisioning Customize guest	Allows customization of a virtual machine's guest operating system without moving the virtual machine.	Virtual machines	VirtualMachine.Provisioning.Customize
Provisioning Deploy template	Allows deployment of a virtual machine from a template.	Templates	VirtualMachine.Provisioning.DeployTemplate
Provisioning Mark as template	Allows marking an existing powered off virtual machine as a template.	Virtual machines	VirtualMachine.Provisioning.MarkAsTemplate
Provisioning Mark as virtual machine	Allows marking an existing template as a virtual machine.	Templates	VirtualMachine.Provisioning.MarkAsVM
Provisioning Modify customization specification	Allows creation, modification, or deletion of customization specifications.	Root vCenter Server	VirtualMachine.Provisioning.ModifyCustSpecs
Provisioning Promote disks	Allows promote operations on a virtual machine's disks.	Virtual machines	VirtualMachine.Provisioning.PromoteDisks
Provisioning Read customization specifications	Allows reading a customization specification.	Virtual machines	VirtualMachine.Provisioning.ReadCustSpecs

Virtual Machine Service Configuration Privileges

Virtual machine service configuration privileges control who can perform monitoring and management tasks on the service configuration.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 341: Virtual Machine Service Configuration Privileges

Privilege Name in the vSphere Client	Description	Privilege Name in the API
• Service configuration – Allow notifications	Allows generating and consuming notification about service status.	VirtualMachine.Namespace.Event
• Service configuration – Allow polling of global event notifications	Allows querying whether any notifications are present.	VirtualMachine.Namespace.EventNotify
• Service configuration – Manage service configurations	Allows creating, modifying, and deleting virtual machine services.	VirtualMachine.Namespace.Management
• Service configuration – Modify service configuration	Allows modification of existing virtual machine service configuration.	VirtualMachine.Namespace.ModifyContent
• Service configuration – Query service configurations	Allows retrieval of list of virtual machine services.	VirtualMachine.Namespace.Query
• Service configuration – Read service configuration	Allows retrieval of existing virtual machine service configuration.	VirtualMachine.Namespace.ReadContent

Virtual Machine Snapshot Management Privileges

Virtual machine snapshot management privileges control the ability to take, delete, rename, and restore snapshots.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 342: Virtual Machine Snapshot Management Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
• Snapshot management – Create snapshot	Allows creation of a snapshot from the virtual machine's current state.	Virtual machines	VirtualMachine.State.CreateSnapshot
• Snapshot management – Remove Snapshot	Allows removal of a snapshot from the snapshot history.	Virtual machines	VirtualMachine.State.RemoveSnapshot
• Snapshot management – Rename Snapshot	Allows renaming a snapshot with a new name, a new description, or both.	Virtual machines	VirtualMachine.State.RenameSnapshot
• Snapshot management – Revert to snapshot	Allows setting the virtual machine to the state it was in at a given snapshot.	Virtual machines	VirtualMachine.State.RevertToSnapshot

Virtual Machine vSphere Replication Privileges

Virtual Machine vSphere replication privileges control the use of replication by VMware vCenter Site Recovery Manager™ for virtual machines.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 343: Virtual Machine vSphere Replication Privileges

Privilege Name in the	Description	Required On	Privilege Name in the API
vSphere Replication Configure Replication	Allows configuration of replication for the virtual machine.	Virtual machines	VirtualMachine.Hbr.ConfigureReplication
vSphere Replication Manage Replication	Allows triggering of full sync, online sync or offline sync on a replication.	Virtual machines	VirtualMachine.Hbr.ReplicaManagement
vSphere Replication Monitor Replication	Allows monitoring of replication.	Virtual machines	VirtualMachine.Hbr.MonitorReplication

Virtual Machine Classes Privileges

Virtual Machine Classes privileges control who can add and remove virtual machine classes on a Kubernetes namespace.

Table 344: Virtual Machine Classes Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Manage Virtual Machine Classes	Allows managing virtual machine classes on Kubernetes namespaces on a Supervisor Cluster.	Clusters	VirtualMachineClasses.Manage

vSAN Privileges

vSAN privileges control who can perform shallow rekey operations and update client information.

Table 345: vSAN Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Cluster ShallowRekey	Allows performing a shallow rekey for a cluster.	Clusters	Vsan.Cluster.ShallowRekey
Xvc UpdateClientInfo	Used internally.	Service user	Vsan.Xvc.UpdateClientInfo

vSAN Stats Privileges

vSphere Stats privileges control the ability to access vSAN metrics.

Table 346: vSAN Stats Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Access service discovery endpoint of vSAN stats	Allows accessing the service discovery endpoint <code>https://vCenterServer-IP/vsan/metrics/serviceDiscovery</code> .	Service account role.	vSANStats.Access

vSphere Zones Privileges

vSphere Zones privileges control who can create and manage vSphere Zones on vSphere with Tanzu.

Table 347: vSphere Zones Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Attach and Detach vSphere objects for vSphere Zones	Allows for associating objects with a vSphere Zone.	Clusters	Zone.ObjectAttachable
Create, Update and Delete vSphere Zones and their associations	Allows for creating and deleting a vSphere Zone.	Clusters	Zone.Manage

vService Privileges

vService privileges control the ability to create, configure, and update vService dependencies for virtual machines and vApps.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 348: vService Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Create dependency	Allows creation of a vService dependency for a virtual machine or vApp.	vApps and virtual machines	vService.CreateDependency
Destroy dependency	Allows removal of a vService dependency for a virtual machine or vApp.	vApps and virtual machines	vService.DestroyDependency
Reconfigure dependency configuration	Allows reconfiguration of a dependency to update the provider or binding.	vApps and virtual machines	vService.ReconfigureDependency
Update dependency	Allows updates of a dependency to configure the name or description.	vApps and virtual machines	vService.UpdateDependency

vSphere Tagging Privileges

vSphere Tagging privileges control the ability to create and delete tags and tag categories, and assign and remove tags on vCenter Server inventory objects.

You can set this privilege at different levels in the hierarchy. For example, if you set a privilege at the folder level, you can propagate the privilege to one or more objects within the folder. The object listed in the Required On column must have the privilege set, either directly or inherited.

Table 349: vSphere Tagging Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Assign or Unassign vSphere Tag	Allows assignment or unassignment of a tag for an object in the vCenter Server inventory.	Any object	InventoryService.Tagging.AttachTag
Assign or Unassign vSphere Tag on Object	Allows objects to have tags assigned or unassigned. Use this privilege to limit which objects users are able to assign or unassign tags to.	Any object	InventoryService.Tagging.ObjectAttachable
Create vSphere Tag	Allows creation of a tag.	Any object	InventoryService.Tagging.CreateTag
Create vSphere Tag Category	Allows creation of a tag category.	Any object	InventoryService.Tagging.CreateCategory
Delete vSphere Tag	Allows deletion of a tag.	Any object	InventoryService.Tagging.DeleteTag
Delete vSphere Tag Category	Allows deletion of a tag category.	Any object	InventoryService.Tagging.DeleteCategory
Edit vSphere Tag	Allows editing of a tag.	Any object	InventoryService.Tagging.EditTag
Edit vSphere Tag Category	Allows editing of a tag category.	Any object	InventoryService.Tagging.EditCategory
Modify UsedBy Field for Category	Allows changing the UsedBy field for a tag category.	Any object	InventoryService.Tagging.ModifyUsedByForCategory
Modify UsedBy Field for Tag	Allows changing the UsedBy field for a tag.	Any object	InventoryService.Tagging.ModifyUsedByForTag

vSphere Client Privileges

vSphere Client privileges control offline access to vCenter Server.

These privileges apply to VMware Cloud only.

vSphere Data Protection Privileges

vSphere Data Protection privileges control the ability to manage the VMware vSphere® Data Protection™ backup and recovery solution.

Table 350: vSphere Data Protection Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Protection	Allows performing data protection operations such as creating and managing backups.	vCenter Server	vSphereDataProtection.Protection
Recovery	Allows performing data protection operations such as restoring backups.	vCenter Server	vSphereDataProtection.Recovery

vSphere Stats Privileges

vSphere Stats privileges control the ability to access vStats state and stats data on objects such as virtual machines and hosts.

Table 351: vSphere Stats Privileges

Privilege Name in the vSphere Client	Description	Required On	Privilege Name in the API
Collect Stats Data	Allows creating and updating acquisition specs in vStats.	The object for which stats data is being collected.	vStats.CollectAny
Modify Stats Configuration	Allows managing the vStats service configuration settings.	vCenter Server	vStats.Settings
Query Stats Data	Allows enumerating stats providers, and metrics and counters the providers expose, for which providers can collect stats data.	The object for which stats data is being queried.	vStats.QueryAny

vSphere Hardening and Compliance

Organizations expect to keep their data secure by reducing the risk of data theft, cyberattack, or unauthorized access. Organizations also must often comply with one or more regulations from government standards to private standards, such as the National Institute of Standards and Technology (NIST) and Defense Information Systems Agency Security Technical Implementation Guides (DISA STIG). Ensuring that your vSphere environment is in compliance with such standards involves understanding a broader set of considerations including people, processes, and technology.

A high-level overview of security and compliance topics that require attention helps you plan your compliance strategy. You also benefit from other compliance-related resources on the VMware Web site.

Security Versus Compliance in the vSphere Environment

The terms security and compliance are often used interchangeably. However, they are unique and distinct concepts.

Security, often thought of as information security, is commonly defined as a set of technical, physical, and administrative controls that you implement to provide confidentiality, integrity, and availability. For example, you secure a host by locking down which accounts can log into it, and by what means (SSH, direct console, and so on). Compliance, by contrast, is a set of requirements necessary to meet the minimum controls established by different regulatory frameworks that provide limited guidance on any specific type of technology, vendor, or configuration. For example, the Payment Card Industry (PCI) has established security guidelines to help organizations proactively protect customer account data.

Security reduces the risk of data theft, cyberattack, or unauthorized access, while compliance is the proof that a security control is in place, typically within a defined time line. Security is primarily outlined in the design decisions and highlighted within the technology configurations. Compliance is focused on mapping the correlation between security controls and specific requirements. A compliance mapping provides a centralized view to list out many of the required security controls. Those controls are further detailed by including compliance citations for each respective security control as dictated by a domain such as NIST, PCI, FedRAMP, HIPAA, and so forth.

Effective cybersecurity and compliance programs are built on three pillars: people, process, and technology. A general misconception is that technology alone can solve all your cybersecurity needs. Technology does play a large and important role in the development and execution of an information security program. However, technology without process and procedures, and awareness and training, creates a vulnerability within your organization.

When defining your security and compliance strategies, keep the following in mind:

- People need general awareness and training, whereas IT staff need specific training.
- Process defines how activities, roles, and documentation within an organization are used to mitigate risk. Processes are only effective if people follow them correctly.
- Technology can be used to prevent or reduce the impact of cybersecurity risk to your organization. Which technology to use depends on the risk acceptance level within an organization.

VMware provides Compliance Kits that contain both an Audit Guide and a Product Applicability Guide, helping to bridge the gap between compliance and regulatory requirements and implementation guides. For more information, see <https://core.vmware.com/compliance>.

Glossary of Compliance Terms

Compliance introduces specific terms and definitions that are important to understand.

Table 352: Compliance Terms

Term	Definition
CJIS	Criminal Justice Information Services. In the context of compliance, the CJIS produces a Security Policy for how local, state, and federal criminal justice and law enforcement agencies must take security precautions to protect sensitive information such as fingerprints and criminal backgrounds.
DISA STIG	Defense Information Systems Agency Security Technical Implementation Guide. The Defense Information Systems Agency (DISA) is the entity responsible for maintaining the security posture of the Department of Defense (DoD) IT infrastructure. DISA accomplishes this task by developing and using Security Technical Implementation Guides, or "STIGs."
FedRAMP	Federal Risk and Authorization Management Program. FedRAMP is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services.
HIPAA	Health Insurance Portability and Accountability Act. Passed by Congress in 1996, HIPAA does the following: • Gives millions of American workers and their families the ability to transfer and continue health insurance coverage for when they change or lose jobs • Reduces health care fraud and abuse • Mandates industry-wide standards for health care information on electronic billing and other processes

Term	Definition
	Requires the protection and confidential handling of protected health information The latter bullet is of most importance to <i>vSphere Security</i> documentation.
NCCoE	National Cybersecurity Center of Excellence. NCCoE is a U.S. government organization that produces and publicly shares solutions to cybersecurity problems that U.S. businesses encounter. The center forms a team of people from cybersecurity technology companies, other federal agencies, and academia to address each problem.
NIST	National Institute of Standards and Technology. Founded in 1901, NIST is a non-regulatory federal agency within the U.S. Department of Commerce. The mission of NIST is to advocate for U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that increase economic security and improve our quality of life.
PAG	Product Applicability Guide. A document that provides general guidance for organizations that are considering a company's solutions to help them address compliance requirements.
PCI DSS	Payment Card Industry Data Security Standard. A set of security standards designed to ensure that all companies that accept, process, store, or transmit credit card information maintain a secure environment.
VVD/VCF Compliance Solutions	VMware Validated Design/VMware Cloud Foundation. The VMware Validated Designs provide comprehensive and extensively tested blueprints to build and operate a Software-Defined Data Center. VVD/VCF compliance solutions enable customers to meet compliance requirements for multiple government and industry regulations.

vSphere Security Controls Reference

VMware Security Hardening Guides provide prescriptive guidance about deploying and operating VMware products in a secure manner. For vSphere, this guide is called the *vSphere Security Configuration Guide* (formerly known as the *Hardening Guide*). Starting in vSphere 8.0 Update 3, the information from the *vSphere Security Configuration Guide* known as security controls is now included in this guide.

Security controls provide security best practices for vSphere. The security controls do not map directly to regulatory guidelines or frameworks. Thus, do not use them as a way towards achieving compliance. Also, the security controls are not intended for use as a security checklist.

Security is always a tradeoff. When you implement security controls, you might affect usability, performance, or other operational tasks negatively. Consider your workloads, usage patterns, organizational structure, and so on carefully before making security changes, whether the advice is from VMware or from other industry sources.

If your organization is subject to regulatory compliance needs, see <https://core.vmware.com/compliance>. This site features compliance kits and product audit guides to help vSphere administrators and regulatory auditors secure and attest virtual infrastructure for regulatory frameworks, such as NIST 800-53v4, NIST 800-171, PCI DSS, HIPAA, CJIS, ISO 27001, and more.

These vSphere security controls do not discuss securing the following items:

- Software running inside the virtual machine, such as the Guest OS and applications

- Traffic running through the virtual machine networks
- Security of add-on products

These vSphere security controls are not meant to be used as a "compliance" tool. These security controls do enable you to take initial steps towards compliance, but used by themselves, they do not ensure that your deployment is compliant. For more information about compliance, see [Security Versus Compliance in the vSphere Environment](#).

Do not blindly apply security controls to your environment. Rather, take time to evaluate each setting and make an informed decision whether you want to apply it. At a minimum, you can use the instructions in the Assessment sections to verify the security of your deployment.

These security controls are an aid to begin implementing compliance in your deployment. When used with the Defense Information Systems Agency (DISA) and other compliance guidelines, you can map vSphere security controls to the compliance flavor per each guideline.

Definitions for Security Controls Terms

In the security control sections that follow, the following terms and definitions are used.

Table 353: Security Controls Definitions

Control Term	Definition
Installation Default Value	What the control defaults to in this version of vSphere when you first install the product.
Baseline Suggested Value	A reasonable recommendation for how you should configure this control, if no other guidance is present. Regulatory compliance guidance might supersede these recommendations, for example.
Action Needed	The suggested action for a particular control. Modify: Make the change. For controls that are outside vSphere, such as hardware settings, this documentation always assumes that the control is set insecurely by default and recommends modifying the configuration. Audit: Ensure that the default is in use, the expected value is present, or exceptions to the control are documented. When auditing a control whose default is the suggested value, two lines of thought are possible. First: Only by setting the parameters explicitly can they be audited and known. Second: All configuration changes require "care and feeding" over time, so where there is a secure default, you can use it to help simplify an environment. This documentation takes the latter approach, but you can choose your own course. Controls that are unimplemented have zero effect on security. They are listed as "audit" in this documentation, but can be removed.
Potential Functional Impact if Default Value Is Changed	Does this change potentially cause problems? Most security controls present tradeoffs in some way. What might changing this control require in exchange?
PowerCLI Command Assessment	An example PowerCLI command to determine how the control is set.
PowerCLI Command Remediation Example	An example PowerCLI command to set the control to the recommendation.

vSphere System Design Security Controls Reference

These security controls provide a baseline set of vSphere system design best practices.

Eliminate vCenter Server Third-Party Plug-ins

Reduce or eliminate third-party vCenter Server plug-ins.

Installation of plug-ins and other third-party cross-connections between systems can erode boundaries between different infrastructure systems, offering opportunities for attackers who have compromised one system to move laterally to another. Tight coupling of other systems to vSphere also often creates impediments to timely patching and upgrades. Ensure that any third-party plug-ins or add-ons to vSphere components create value. If you choose to use plug-ins rather than individual management consoles, be sure that their use offsets the risks that they create.

Use Caution with Infrastructure Management Interfaces

Use caution when connecting infrastructure management interfaces to general-purpose authentication and authorization sources.

Centralized enterprise directories are targets for attackers because of their role in authorization across an enterprise. An attacker can move freely inside an organization once that directory is compromised. Connecting IT infrastructure to centralized directories has proven to be a considerable risk for ransomware and other attacks. Isolate the authentication and authorization of all infrastructure systems.

For ESXi:

- Conduct all host management through vCenter Server
- Deactivate the ESXi Shell
- Place ESXi in normal lockdown mode
- Set the ESXi root password to a complex password

Activate vSphere Distributed Resource Scheduler

Activate vSphere Distributed Resource Scheduler (DRS) in Fully automated mode.

vSphere DRS uses vMotion to move workloads between physical hosts to ensure performance and availability. Fully automated mode ensures that the vSphere Lifecycle Manager can work with DRS to activate patching and update operations.

If specific VM-to-host mappings are needed, use DRS rules. Where possible, use "should" rules instead of "must" so that you can suspend the rule temporarily during patching and high availability recovery.

Activate vSphere High Availability

vSphere High Availability (HA) restarts workloads on other ESXi hosts in a cluster if an ESXi host fails suddenly. Ensure that the settings for HA are configured correctly for your environment.

Activate Enhanced vMotion Compatibility

vSphere Enhanced vMotion Compatibility (EVC) ensures that workloads can be live-migrated by using vMotion between ESXi hosts in a cluster that are running different CPU generations. EVC also assists in situations with CPU vulnerabilities, where new microcode instructions might be introduced to CPUs, which makes them temporarily incompatible with one another.

Protect Systems from Tampering

Ensure that ESXi hosts and related storage and networking components are protected from tampering, unauthorized access, and unauthorized removal. Also, protect hosts from damage from environmental factors such as flooding, extreme temperatures (low or high), and dust and debris.

Use of security features, such as vSphere Native Key Provider and ESXi Key Persistence, might cause secure material to be stored locally on ESXi hosts, enabling attackers to boot and unlock otherwise protected clusters. Consideration of physical security and appropriate threats, like theft, is important.

Beyond theft, being security-minded also means asking yourself and your organization questions such as the following:

- What could go wrong?
- How would I know if something went wrong?

These questions take on added importance when dealing with unstaffed data center locations and collocation facilities. With respect to data centers and rack configurations, ask the following questions:

- Do the doors to the data center automatically close and lock properly on their own?
- If the doors were left ajar, would there be a proactive alert?
- If your rack doors are locked, is it still possible to reach into the rack from the side or top and disconnect a cable? Can an unauthorized person connect a cable to a network switch?
- Is it possible to remove a device, like a storage device or even an entire server? What would happen in such a scenario?

Other questions to ask include:

- Could someone glean information about your environment or your business from information displays on the servers, such as LCD panels or consoles?
- If those information displays are inactive, could they be triggered from outside the rack, for example, with the use of a stiff metal wire?
- Are there other buttons, such as the power button, that could be pushed to create a service disruption to your company?

Finally, ask yourself, are there other physical threats, such as the possibility of flooding, freezing or high heat, or dust and debris from the environment, that would impact availability?

Name vSphere Objects Descriptively

Ensure that you name vSphere objects descriptively, changing the default names of objects to ensure accuracy and reduce confusion.

Use good naming practices for vSphere objects, changing default names such as "Datacenter," "vSAN Datastore," "DSwitch," "VM Network," and so on, to include additional information. This helps improve accuracy and reduce errors when developing, implementing, and auditing security policies and operational processes.

Port groups using 802.1Q VLAN tagging could include the VLAN number. Data centers and cluster names could reflect locations and purposes. Datastore and virtual distributed switch names could reflect the data center and cluster names to which they are attached. Key provider names are particularly important, especially when protecting encrypted virtual machines with replication to alternate sites. Work to avoid potential "name collisions" with objects present in other data centers and clusters.

Some organizations do not name systems with physical location identifiers such as street addresses, preferring to obscure the physical location of data centers through the use of terms like "Site A," "Site B," and so on. This also helps if sites are relocated, preventing the need to rename everything or endure inaccurate information.

When deciding on a naming scheme, keep in mind that many objects can have similar properties. For example, two port groups could both have the same VLAN assigned, but have different traffic filtering and marking rules. Incorporating a project name or short description in the name might be helpful for disambiguating objects of this type.

Lastly, consider automation when developing a naming scheme. Names that can be derived programmatically are often helpful when scripting and automating tasks.

Isolate Infrastructure Management Interfaces

Ensure that IT infrastructure management interfaces are isolated on their own network segment or as part of an isolated management network.

Ensure that all management interfaces configured for virtualization components are on a network segment (VLAN, and so on) that is dedicated only to virtualization management, free of workloads and unrelated systems. Ensure that management interfaces are controlled with perimeter security controls such that only authorized vSphere administrators can access those interfaces from authorized workstations.

Some system designs put vCenter Server and other management tools on their own network segments, isolated from ESXi, because it offers better monitoring of those systems. Other designs put vCenter Server in with ESXi management because of the relationship between the two products, and the possibility of firewall configuration errors or outages disrupting service. Whichever design you choose, do so thoughtfully.

Use vMotion Properly

Ensure that vMotion uses data-in-transit encryption (set to "Required" for virtual machines), or that VMkernel network interfaces used for vMotion are isolated on their own network segments that have perimeter controls.

vMotion and Storage vMotion copy virtual machine memory and storage data, respectively, across the network. Ensuring that the data is encrypted in transit ensures confidentiality. Isolation to a dedicated network segment with appropriate perimeter controls can add defense-in-depth and also allow for network traffic management.

Like all forms of encryption, vMotion encryption does introduce performance loss, but that performance change occurs on the background vMotion process and does not impact virtual machine operation.

Use vSAN Properly

Ensure that vSAN uses data-in-transit encryption or that VMkernel network interfaces used for vSAN are isolated on their own network segments that have perimeter controls.

vSAN features data-in-transit encryption that can help maintain confidentiality as vSAN nodes communicate. As with many security controls, there is a tradeoff with performance. Monitor storage latency and performance as data-in-transit encryption is activated. Organizations that do not or cannot activate vSAN data-in-transit encryption should isolate the network traffic to a dedicated network segment with appropriate perimeter controls.

Activate Network I/O Control

Ensure that you have resilience to network denial-of-service by activating Network I/O Control (NIOC).

vSphere Network I/O Control (NIOC) is a traffic management technology that offers quality of service at the hypervisor level, enhancing network performance by prioritizing resources in multi-tenant cloud and shared workload environments. Incorporated into the vSphere Distributed Switch (vDS), NIOC partitions network adapter bandwidth into "network resource pools" that correspond to different traffic types, such as vMotion and management traffic. Use of NIOC allows users to allocate shares, limits, and reservations to these pools.

NIOC preserves network availability for essential services and prevents congestion by limiting less critical traffic. This is achieved by enabling the creation of network control policies per business requirements, ensuring traffic type isolation, and allowing dynamic resource reallocation based on priority and usage.

Do Not Configure Vendor-Reserved VLANs

Ensure that the physical switch uplinks from ESXi hosts are not configured with vendor-reserved VLANs.

Some network vendors reserve particular VLAN IDs for internal or specific use. Ensure that your vSphere network configurations do not include these values.

Configure ESXi Uplinks as Access Ports

Ensure that the physical switch uplinks from ESXi hosts are configured as "access ports" assigned to a single VLAN, or as tagged 802.1Q VLAN trunks with no native VLAN. Ensure that vSphere port groups do not allow access to VLAN 1 or untagged native VLANs.

Network connections that have a "native" VLAN configured to accept untagged traffic, or that have access to VLAN 1, might offer opportunities for attackers to craft specialized packets that defeat network security controls. VLAN 1 is the default often used for network management and communications and should be isolated from workloads. Ensure that port groups are not configured for access to native VLANs. Ensure that VLAN trunk ports are configured with specific definitions of VLANs (not "all"). Finally, ensure that port groups are configured appropriately so that attackers cannot use a virtualized environment to circumvent network security controls.

Configure Storage Fabric Connections Properly

Ensure that the storage fabric connections use data-in-transit encryption or are isolated on their own network segments or SANs that have perimeter controls.

Protecting storage data while in transit helps ensure the confidentiality of the data. Encryption is not an option for many storage technologies, often because of availability or performance concerns. In those cases, isolation to a dedicated network segment with the appropriate perimeter controls can be an effective compensating control and can add defense-in-depth.

Use LUN Masking on Storage Systems

Ensure that the storage systems employ LUN masking, zoning, and other storage-side security techniques to ensure that storage allocations are only visible to the vSphere cluster in which it is to be used.

LUN masking on the storage controller and SAN zoning help to ensure that storage traffic is not visible to unauthorized hosts and that unauthorized hosts cannot mount the datastores, bypassing other security controls.

Limit Connections to Authorized Systems

Consider the use of the vCenter Server Appliance firewall to limit connections to authorized systems and administrators.

The vCenter Server Appliance contains a basic firewall that you can use to limit the incoming connections to vCenter Server. This can be an effective layer of defense-in-depth in conjunction with perimeter security controls.

As always, before adding rules to block connections, ensure that rules are in place to allow access from administrative workstations.

Do Not Store Encryption Keys on ESXi Hosts Without Securing Physical Access

The environment must not store encryption keys on ESXi hosts without also securing physical access to the hosts.

To prevent dependency loops, the vSphere Native Key Provider stores decryption keys directly on the ESXi hosts, either in a Trusted Platform Module (TPM) or as part of the encrypted ESXi configuration. However, if you do not physically secure a host, and an attacker steals the host, the attacker possesses the means to unlock and execute encrypted workloads. Therefore, it is crucial to ensure physical security (see [Protect Systems from Tampering](#)), or to opt for using a Standard Key Provider (see [What Is a Standard Key Provider](#)) that includes additional network security controls.

Use Adequately Sized Persistent, Non-SD, Non-USB Devices for ESXi Boot Volumes

The environment must use adequately sized persistent, non-SD, non-USB devices for ESXi boot volumes.

Flash memory is a component that wears out over time, with each data write shortening its lifespan. SSDs and NVMe devices have built-in features to reduce this wear, making them more reliable. However, SD cards and most USB flash drives do not have these features and can develop reliability issues, such as bad sectors, often without any obvious signs.

To lessen wear and make SD and USB devices last longer, when you install ESXi on these devices, you can save audit and system logs to a RAM disk instead of constantly writing to the device. This means that you must set up new, long-term storage locations for these logs and change the log output to go to these new locations.

Choosing a reliable boot device removes these extra steps and helps ESXi automatically pass security audits.

Properly Configure the vSAN iSCSI Target

Ensure that the vSAN iSCSI Target uses its own VMkernel network interfaces, isolated on its own network segment and employing separate perimeter controls using Distributed Port Group Traffic Filtering and Marking, NSX, or external network security controls.

Because the iSCSI Target clients are external to the cluster, isolate them on their own network interfaces. In this way, you can separately restrict other, internal-only network communications. Isolation of this type also helps diagnose and manage performance.

vSphere Hardware Security Controls Reference

These security controls provide a baseline set of vSphere hardware security best practices. They are structured in a way that explains the benefits and tradeoffs of implementing the control.

Variable Used

The PowerCLI commands in this section use the following variable:

- \$ESXi = "host_name"

Use Intel Trusted Execution Technology

Ensure that Intel Trusted Execution Technology (TXT) is activated, if available in the system firmware.

Intel Xeon Scalable Processor platforms have TXT, which provides authenticity of a platform and its operating system. When activated, ESXi takes advantage of the security benefits offered by this technology.

Suggested Value	Activated
Potential Impact if Default Value Is Changed	Early implementations of TXT occasionally caused sudden system shutdowns, triggering attestation alarms in vCenter Server, or even boot failures. A system restart resolves these problems, while an update to system firmware usually resolves it permanently. See the VMware knowledge base article at https://kb.vmware.com/s/article/78243 .

PowerCLI Command Assessment

```
(Get-VMHost -Name $ESXi | Get-View).Capability.TxtEnabled
```

Configure UEFI Secure Boot

Ensure that UEFI Secure Boot is activated.

Activating UEFI Secure Boot on the hardware of an ESXi host helps prevent malware and untrusted configurations.

Suggested Value	Activated
------------------------	-----------

Potential Impact if Default Value Is Changed

Activating UEFI Secure Boot after installation might prevent an ESXi host from starting. Run `/usr/lib/vmware/secureboot/bin/secureBoot.py -c` on an example host to determine if you can activate Secure Boot safely.

PowerCLI Command Assessment

N/A

Use TPM 2.0

Ensure that a Trusted Platform Module (TPM) 2.0 is installed and configured properly on your ESXi hosts.

ESXi can use a TPM to activate advanced security features that prevent malware, remove dependencies, and secure hardware lifecycle operations. When possible, configure your hosts to use TPM 2.0 and activate the TPM in the system firmware.

Suggested Value

TPM 2.0 installed and activated (SHA-256 hashing, TIS/FIFO interface)

Potential Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
(Get-VMHost -Name $ESXi | Get-View).Capability.TpmSupported
(Get-VMHost -Name $ESXi | Get-View).Capability.TpmVersion
```

Ensure that Hardware Firmware Is Up-to-date

Ensure that you apply the latest firmware updates to all components of your systems, and that the firmware is authentic and supplied by your hardware manufacturer.

Hardware firmware is not immune from serious issues affecting confidentiality, integrity, or availability. Attackers can use vulnerable system management controllers and management engines to establish persistence, and re-infect and re-compromise hosts after reboots and updates.

Suggested Value

N/A

Potential Impact if Default Value Is Changed

If you use vSAN, ensure that the storage device and the controller firmware versions are certified.

PowerCLI Command Assessment

N/A

Secure Integrated Hardware Management Controllers

Ensure that integrated hardware management controllers are fully secured.

Many servers have integrated hardware management controllers that can be extremely helpful when monitoring and updating hardware, settings, and firmware. For these controllers:

- Deactivate all unused functionality.
- Disable all unused access methods.
- Set passwords and password controls.
- Put firewalls and access control in place so that access only occurs from authorized access workstations for the virtualization administration team.

Deactivate all "first boot" configuration options, especially ones that reconfigure the system from an inserted USB device. Also, deactivate or protect USB ports attached to management controllers. Where possible, set USB ports to permit only keyboards.

Change default passwords for accounts.

Secure external information displays to prevent information from leaking. Secure power and information buttons against unauthorized use.

Many hardware management controllers provide alert mechanisms when hardware faults and configuration changes occur. Consider using these if you are not using another method for hardware monitoring.

Suggested Value	N/A
Potential Impact if Default Value Is Changed	Deactivating connection methods might cause future monitoring and management changes to the hardware management controller configurations across your deployed servers. When possible, use CLI and API management methods that you can script in lieu of using additional management software or applications. Learning these techniques saves time, avoids the additional effort of installing and maintaining additional tools, and allows for timely configuration changes.
PowerCLI Command Assessment	N/A

Synchronize Time on Integrated Hardware Management Controllers

Ensure that you synchronize time on integrated hardware management controllers.

Cryptography, audit logging, cluster operations, and incident responses depend on synchronized time. This recommendation extends to all devices in your infrastructure. Network Time Protocol (NTP) must have at least four sources. If you must choose between two sources and one source, one source is preferable.

Suggested Value	Site-Specific or: 0.vmware.pool.ntp.org, 1.vmware.pool.ntp.org, 2.vmware.pool.ntp.org, 3.vmware.pool.ntp.org
Potential Impact if Default Value Is Changed	None
PowerCLI Command Assessment	N/A

Secure How Integrated Hardware Management Controllers Use Active Directory

Ensure that you do not create a dependency loop or attack vector in how integrated hardware management controllers use Active Directory.

Either deactivate connections to Active Directory, or at a minimum, considered them as attack vectors and dependency loops (for authentication, authorization, DNS, DHCP, and time). Consider managing local accounts on these devices through APIs and CLIs. If you must use Active Directory for authentication, use local authorization so that attackers with access to Active Directory cannot promote themselves through group membership.

Suggested Value	N/A
Potential Impact if Default Value Is Changed	Not connecting hardware management controllers to centralized authentication and authorization sources entails additional management. Most hardware management controllers have CLI toolkits or APIs to automate the process.
PowerCLI Command Assessment	N/A

Deactivate Virtual Integrated Hardware Management Controllers

Ensure that integrated hardware management controllers with internal, emulated, or virtual network interfaces are deactivated.

Some hardware management controllers have the ability to present virtual network interfaces to ESXi as a management interface. These approaches create potential back doors for access that adversaries can use to circumvent network-based and perimeter firewalls, in either direction, and to avoid observation by IDS, IPS, and threat analysis tools. In many cases, this functionality is not strictly necessary to manage hosts.

Suggested Value	N/A
Potential Impact if Default Value Is Changed	Deactivating internal networking might limit vendor management tool effectiveness.
PowerCLI Command Assessment	N/A

Activate AMD Secure Encrypted Virtualization-Encrypted State

Ensure that AMD Secure Encrypted Virtualization-Encrypted State (SEV-ES) is activated, if available in the system firmware. Ensure that the value for Minimum SEV non-ES ASID is equal to the number of SEV-ES virtual machines plus one.

AMD EPYC platforms support SEV-ES, a technology to encrypt memory and CPU register state, and limit visibility to the hypervisor, to increase workload security and decrease exposure to certain types of attacks. When configured properly, SEV-ES provides enhanced security to the guest operating system on virtual machines and containers under vSphere and vSphere with Tanzu. Activating SEV-ES in system firmware eases future enablement inside virtual machines, containers, and guest operating systems.

Suggested Value	Activated (Minimum SEV non-ES ASID is equal to the number of SEV-ES virtual machines plus one)
Potential Impact if Default Value Is Changed	The guest operating system for a virtual machine must support SEV-ES, and so limits some features, such as vMotion, snapshots, and so on. For more information about these tradeoffs, see Unsupported VMware Features on SEV-ES .
PowerCLI Command Assessment	N/A

Activate Virtual Intel Software Guard Extensions (vSGX)

Ensure that Virtual Intel® Software Guard Extensions (vSGX) is activated, if available in the system firmware.

Intel Xeon Scalable Processor platforms have Software Guard Extensions, or SGX, a technology that helps applications protect data in system memory. When configured properly, vSphere supports the use of SGX inside virtual machines. Enabling SGX in system firmware eases future enablement inside virtual machines and guest operating systems.

Suggested Value	Suggested: Activated (software, unlocked)
Potential Impact if Default Value Is Changed	The guest operating system for a virtual machine must support vSGX, and so limits some features such as vMotion, snapshots, and so on. For more information about these tradeoffs, see Unsupported VMware Features on vSGX .
PowerCLI Command Assessment	<pre>(Get-VMHost -Name \$ESXi Get-View).Capability.SgxRegistrationSupported</pre>

Deactivate External Ports

Ensure that unused external ports are deactivated or protected against unauthorized use.

Unused ports, especially USB, can be used by attackers to attach storage, networking, and keyboards. Take reasonable steps to control access to these ports through disablement and access control. Where possible, use other means such as solid rack doors, rack side panels, and flooring to make the ports inaccessible from outside the rack when the rack door is closed. Be aware that cables fit easily through many gaps in and around racks and rack doors, and stiff wires can be used to push cables into sockets from outside the rack, as well as to dislodge cables to create a service disruption.

Where possible, set USB ports to permit only keyboards.

When deactivating this type of functionality, consider that you might need to access a server using a USB keyboard during an outage, or as part of lifecycle operations, and plan accordingly.

Suggested Value	N/A
Potential Impact if Default Value Is Changed	Security is always a tradeoff. When considering a security control such as deactivating external ports, make ease of recovery from an outage or an incident a part of the equation. In this case, deactivating external ports affects the ability to use the ESXi console in case of emergency. Some servers can dynamically deactivate and activate certain USB ports for management. Ensure that your choice for this security control meets the needs of your organization, and that you test these methods before implementing them.
PowerCLI Command Assessment	N/A

ESXi Security Controls Reference

These security controls provide a baseline set of ESXi security best practices. They are structured in a way that explains the benefits and tradeoffs of implementing the control. Most controls are in the form of advanced system settings. To change advanced system settings, you can use either the PowerCLI provided, or the vSphere Client (**Host > Configure > System > Advanced System Settings**).

Variable Used

The PowerCLI commands in this section use the following variables:

- `$ESXi = "host_name"`
- `$vmkernel_interface = "vmkernel_adapter"`

Ensure the DCUI Account Is Denied Access

The ESXi host must deny shell access for the dcui user account.

The dcui user account is used for process isolation for the DCUI itself. To reduce attack surface, deactivate shell access for the dcui user account.

Values	Installation Default Value: True Baseline Suggested Value: False
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.account.list.Invoke() Where-Object { \$_.UserID -eq 'dcui' } Select-Object -Expand- Property Shellaccess</pre>

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.account.set.CreateArgs()
$arguments.id = "dcui"
$arguments.shellaccess = "false"
$ESXcli.system.account.set.Invoke($arguments)
```

Annotations.WelcomeMessage

Configures the text of the login message that is displayed on the VMware Host Client and the DCUI.

ESXi provides the ability to display a login message. Usages of the login message include informing intruders that their activities are illegal, and conveying to authorized users the expectations and obligations they must meet and agree to while using the system.

Values

Installation Default Value: Undefined

Baseline Suggested Value: Consult the legal advisers for your organization for text that is applicable to your environment.

Example message text: Authorized users only. Actual or attempted unauthorized use of this system is prohibited and may result in criminal, civil, security, or administrative proceedings and/or penalties. Use of this information system indicates consent to monitoring and recording, without notice or permission. Users have no expectation of privacy in any use of this system. Any information stored on, or transiting this system, or obtained by monitoring and/or recording, may be disclosed to law enforcement and/or used in accordance with Federal law, State statute, and organization policy. If you are not an authorized user of this system, exit the system at this time.

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

Masks the "F2/F12" and IP address information on the DCUI. Might also require documentation and training for your environment.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Annotations.WelcomeMessage
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Annotations.WelcomeMessage | Set-AdvancedSetting -Value "your_message"
```

Config.HostAgent.vmacore.soap.sessionTimeout

Configures a session timeout for the vSphere API.

This practice helps mitigate potential security risks by ensuring that unattended sessions, which could be exploited by unauthorized users or malicious software, are not left open indefinitely.

Values

Installation Default Value: 30

Baseline Suggested Value: 30

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Config.HostAgent.vmacore.soap.sessionTimeout
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Con-
fig.HostAgent.vmacore.soap.sessionTimeout | Set-Ad-
vancedSetting -Value 30
```

Config.Etc.issue

Configures the text of the banner when a user connects to anESXi host by using SSH.

ESXi provides the ability to display banner for SSH connections. Usages of the banner include informing intruders that their activities are illegal, and conveying to authorized users the expectations and obligations they must meet and agree to while using the system. Keep the SSH service deactivated unless you are performing troubleshooting operations. An implementation inconsistency betweenESXi andvCenter Server requires that "issue" in Config.Etc.issue be lowercase, to work in both scenarios.

Values

Installation Default Value: Undefined

Baseline Suggested Value: Consult the legal advisers for your organization for text that is applicable to your environment.

Example message text: Authorized users only. Actual or attempted unauthorized use of this system is prohibited and may result in criminal, civil, security, or administrative proceedings and/or penalties. Use of this information system indicates consent to monitoring and recording, without notice or permission. Users have no expectation of privacy in any use of this system. Any information stored on, or transiting this system, or obtained by monitoring and/or recording, may be disclosed to law enforcement and/or used in accordance with Federal law, State statute, and organization policy. If you are not an authorized user of this system, exit the system at this time.

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Con-
fig.Etc.issue
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
Config.Etc.issue | Set-AdvancedSetting -Value
*****
Authorized users only. Actual or attempted unau-
thorized use of this      *`n* system is prohibit-
ed and may result in criminal, civil, security, or
      *`n* administrative proceedings and/or penal-
ties. Use of this information      *`n* system indi-
cates consent to monitoring and recording, without
notice      *`n* or permission. Users have no expec-
tation of privacy. Any information      *`n* stored
on or transiting this system, or obtained by mon-
itoring and/or      *`n* recording, may be disclosed
to law enforcement and/or used in accordance *`n*
with Federal law, State statute, and organization
policy. If you are not *`n* an authorized user of
this system, exit the system at this time.
*`n*****
```

Deactivate Shell Access for vpxuser

The ESXi host must deny shell access for the vpxuser account.

vCenter Server creates the vpxuser account when an ESXi host is first attached. The vpxuser account is subsequently used for privileged authentication to ESXi. While vCenter Server automatically rotates the password for the vpxuser account on an interval governed by the `VirtualCenter.VimPasswordExpirationInDays` option, the vpxuser account also has shell access. Deactivate the vpxuser account to reduce attack surface.

Values	Installation Default Value: True Baseline Suggested Value: False
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	User accounts that do not have shell access cannot reconfigure the shell access of other users, no matter their privilege levels. Because vCenter Server connects to an ESXi host as the vpxuser account, once you deactivate the shell access for vpxuser, it can no longer be used to change those account settings for other accounts. Further reconfiguration must occur on a host-by-host basis using an account that is authorized. ESXi 8.0 and later no longer support traditional password or account recovery operations, such as booting from media or changing init to a shell on boot. Ensure that the ESXi host retains at least one fully privileged user account, and that you protect this account accordingly.
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.account.list.Invoke() Where-Object { \$_.UserID -eq 'vpxuser' } Select-Object -Ex- pandProperty Shellaccess</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.account.set.CreateArgs() \$arguments.id = "vpxuser" \$arguments.shellaccess = "false" \$ESXcli.system.account.set.Invoke(\$arguments)</pre>

vCenter Server Must Use the vSphere Authentication Proxy to Avoid Storing Active Directory Credentials

The vSphere Authentication Proxy enables vCenter Server to connect to and manage Active Directory entities without the need to directly store Active Directory credentials, which reduces the risk of credential exposure or misuse.

Values	Installation Default Value: Not configured Baseline Suggested Value: Configure if using the feature
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-VMHostAuthentication Select-Object VMHost,Domain,DomainMembershipStatus</pre>
PowerCLI Command Remediation Example	N/A

DCUI.Access

The ESXi host must have an accurate DCUI.Access list.

Sets the Lockdown Mode Exception Users list to contain an accurate list of users, and ensures that only authorized users have direct console user interface (DCUI) access to the ESXi host when Lockdown Mode is activated.

You cannot remove the root user from the list.

To control ESXi Shell and SSH access, use the Lockdown Mode Exception Users list. See [Ensure the ESXi Host Has an Accurate Exception Users List](#).

Values	Installation Default Value: root Baseline Suggested Value: root
Action Needed	Audit the list.
Potential Functional Impact if Default Value Is Changed	Potential loss of administrative access to hosts. Ensure that you attach ESXi hosts to vCenter Server, and that you configure access lists and exception lists before configuring lockdown mode.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting DCUI.Access</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting DCUI.Access Set-AdvancedSetting -Value root</pre>

Ensure the ESXi Host Has an Accurate Exception Users List

The ESXi host must have an accurate Exception Users list.

Users on the Lockdown Mode Exception Users list do not lose their privileges when the host enters lockdown mode. This situation might possibly defeat the purpose of lockdown mode.

Values	Installation Default Value: Null Baseline Suggested Value: Null
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Potential loss of administrative access to ESXi hosts. Ensure that you attach ESXi hosts to vCenter Server, and that you configure access lists and exception lists before configuring lockdown mode.
PowerCLI Command Assessment	<pre>(Get-View (Get-VMHost -Name \$ESXi Get-View).Con- figManager.HostAccessManager).QueryLockdownExcep- tions()</pre>
PowerCLI Command Remediation Example	<pre>(Get-View (Get-VMHost -Name \$ESXi Get-View).Con- figManager.HostAccessManager).UpdateLockdownExcep- tions(\$NULL)</pre>

Activate Normal Lockdown Mode to Restrict Access to ESXi

Activating lockdown mode deactivates direct access to an ESXi host. Lockdown mode requires vCenter Server to manage the ESXi host directly.

Restricting access in this way ensures that vCenter Server enforces roles and permissions. Also, users cannot bypass these roles and permissions by logging into an ESXi host directly. Requiring all interaction to occur through vCenter Server reduces the risk of users inadvertently attaining elevated privileges or performing tasks for which they are not properly audited.

Users listed in the Exception Users list for each ESXi host are allowed to override lockdown mode and log in. By default, no users are present on the Exception Users list.

The lockdown mode settings are Disabled, Normal, and Strict. When lockdown mode is set to Strict, if the ESXi host loses contact with vCenter Server, you cannot manage it until that connection is restored. If you cannot restore the connection, you must rebuild the ESXi host. In general, Strict lockdown mode exceeds the needs of most deployments. As such, Normal lockdown mode typically suffices.

Values	Installation Default Value: lockdownDisabled Baseline Suggested Value: lockdownNormal
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Potential loss of administrative access to hosts. Ensure that you attach ESXi hosts to vCenter Server, and that you configure access lists and exception lists before configuring lockdown mode. Some operations, such as backup and troubleshooting, require direct access to the ESXi host. In these cases, you can deactivate lockdown mode temporarily for specific hosts, then re-activate lockdown mode when you are done.

PowerCLI Command Assessment

```
(Get-View (Get-VMHost -Name $ESXi | Get-View).ConfigManager.HostAccessManager).LockdownMode
```

PowerCLI Command Remediation Example

```
(Get-View (Get-VMHost -Name $ESXi | Get-View).ConfigManager.HostAccessManager).ChangeLockdownMode('lockdownNormal')
```

Syslog.global.auditRecord.storageEnable

Configures the ESXi host to store audit records locally.

You must activate audit record logging on ESXi hosts.

Values	Installation Default Value: False Baseline Suggested Value: True
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Additional storage space is consumed by logs.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Syslog.global.auditRecord.storageEnable</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Syslog.global.auditRecord.storageEnable Set-AdvancedSetting -Value TRUE</pre>

Syslog.global.auditRecord.storageCapacity

You must activate storage capacity for one week of audit records on ESXi hosts.

If a remote audit record storage facility is available, it is essential to ensure that the local storage capacity is sufficient to hold audit records that can accumulate during anticipated interruptions in the delivery of records to the facility. This ensures that audit records are not lost or overwritten during periods when the remote storage is unavailable, allowing for seamless continuity of the audit trail and compliance requirements.

Values	Installation Default Value: 4 Baseline Suggested Value: 100
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Additional storage space is consumed by logs.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.auditRecord.storageCapacity
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.auditRecord.storageCapacity | Set-Ad-
vancedSetting -Value 100
```

ScratchConfig.CurrentScratchLocation and Syslog.global.auditRecord.storageDirectory

Configures a persistent log location for all locally stored audit records on the ESXi host.

You can configure ESXi to store audit records on an in-memory file system. This occurs when the "/scratch" directory of the host is linked to "/tmp/scratch." When this is done, only a single day's worth of records are stored at any time. In addition, audit records are reinitialized upon each reboot. This presents a security risk as user activity logged on the host is only stored temporarily and does not persist across reboots. This can also complicate auditing and make it harder to monitor events and diagnose issues. Always configure ESXi host audit record logging to a persistent datastore.

You can detect if the scratch volume is temporary or persistent by querying the `ScratchConfig.CurrentScratchLocation` advanced setting. If, when queried, it returns "/tmp/scratch," then the volume is temporary and you should remap the audit record storage to a persistent device.

The storage cannot be a vSAN datastore. If your only local, non-vSAN storage is SD or USB media (which can become unreliable with repeated writes from logs), consider leaving the logs in the ramdisk and ensuring that a remote logging host is configured instead. Document the decision and rationale in preparation for future audits.

Values

Installation Default Value:

`ScratchConfig.CurrentScratchLocation`: Depends on boot device

`Syslog.global.auditRecord.storageDirectory`: /scratch/auditLog

Baseline Suggested Value: Persistent storage location

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

Additional storage space is consumed by logs.

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.syslog.config.get.Invoke() | Select
LocalLogOutput,LocalLogOutputIsPersistent

# If your LocalLogOutput is set to a directory in /
scratch, and LocalLogOutputIsPersistent is true,
that means your boot device is of a type and size
that makes /scratch persistent. Verify that your
audit storage is also on /scratch, and that /
scratch points to a VMFS datastore:
```

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
ScratchConfig.CurrentScratchLocation
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.auditRecord.storageDirectory
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.auditRecord.storageDirectory | Set-Ad-
vancedSetting -Value "/vmfs/volumes/$Datastore/au-
dit"
```

Syslog.global.auditRecord.remoteEnable

Configures theESXi host for transmission of audit logs to a remote host.

Values	Installation Default Value: False Baseline Suggested Value: True
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.auditRecord.remoteEnable</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.auditRecord.remoteEnable Set-Ad- vancedSetting -Value TRUE</pre>

Syslog.global.logFiltersEnable

Activates log filtering on theESXi host.

You can create log filters to reduce the number of repetitive entries, and to deny specific log events entirely.

Values	Installation Default Value: False Baseline Suggested Value: False
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.logFiltersEnable</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.logFiltersEnable Set-AdvancedSetting - Value FALSE</pre>

LocalLogOutputIsPersistent, ScratchConfig.CurrentScratchLocation, and Syslog.global.logDir

Configures persistent logging for all locally stored logs on theESXi host.

You can configureESXi to store log files on an in-memory file system. This occurs when the "/scratch" directory of the host is linked to "/tmp/scratch." When this is done, only a single day's worth of logs is stored at any time. Additionally, log files are reinitialized upon each reboot. This presents a security risk as user activity logged on the host is only stored temporarily and does not persist across reboots. It can also complicate auditing and make it harder to monitor events and diagnose issues. Always configureESXi host logging to a persistent datastore.

You can detect if the scratch volume is temporary or persistent by querying the `ScratchConfig.CurrentScratchLocation` advanced parameter. If, when queried, it returns "/tmp/scratch," then the volume is temporary and you should remap the audit record storage to a persistent device.

The storage cannot be avSAN datastore, unless you set `Syslog.global.vsanBacking`, which has caveats and dependencies. If your only local, non-vSAN storage is SD or USB media (which can become unreliable with repeated writes from logs), consider leaving the logs in the ramdisk and ensuring that a remote logging host is configured instead. Document the decision and rationale in preparation for future audits.

Values	Installation Default Value: <code>ScratchConfig.CurrentScratchLocation</code> : Depends on boot device
---------------	--

Action Needed**Potential Functional Impact if Default Value Is Changed****PowerCLI Command Assessment**

```
Syslog.global.logDir: /scratch/log
```

Baseline Suggested Value: Persistent storage location

Modify the installation default value.

None

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.syslog.config.get.Invoke() | Select
    LocalLogOutput,LocalLogOutputIsPersistent
```

```
# If your LocalLogOutput is set to a directory in /
scratch, and LocalLogOutputIsPersistent is true,
that means your boot device is of a type and size
that makes /scratch persistent. Verify that your
log storage is also on /scratch, , and that /
scratch points to a VMFS datastore:
```

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
    ScratchConfig.CurrentScratchLocation
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.logDir
```

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.logDir | Set-AdvancedSetting -Value "/
vmfs/volumes/$Datastore/logs"
```

PowerCLI Command Remediation Example**Syslog.global.logHost**

Configures remote logging.

When you configure remote logging to a central log host, you provide a secure, centralized store for ESXi logs. Gathering host log files onto a central host gives you the ability to monitor all hosts with a single tool. You can also perform aggregate analysis and search for items such as coordinated attacks on multiple hosts. Logging to a secure, central log server helps prevent log tampering and also provides a long-term audit record.

Values

Installation Default Value: Undefined

Baseline Suggested Value: Site-specific

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.logHost
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.logHost | Set-AdvancedSetting -Value
    "log_collector"
```

Syslog.global.certificate.checkSSLCerts

Verifies certificates for TLS.

The ESXi host must verify certificates for TLS remote logging endpoints. TLS certificates help ensure that the endpoint is authentic and trustworthy.

Values

Installation Default Value: True

Action Needed	Baseline Suggested Value: True Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.certificate.checkSSLCerts</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.certificate.checkSSLCerts Set-Ad- vancedSetting -Value TRUE</pre>

Syslog.global.certificate.strictX509Compliance

Performs strict x509 verification for TLS-enabled remote logging endpoints.

TheESXi host must use strict x509 verification for TLS-enabled remote logging endpoints. The `Syslog.global.certificate.strictX509Compliance` setting performs additional validity checks on CA root certificates during verification.

Values	Installation Default Value: False Baseline Suggested Value: True
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.certificate.strictX509Compliance</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.certificate.strictX509Compliance Set- AdvancedSetting -Value TRUE</pre>

Mem.MemEagerZero

Activates volatile key destruction.

By default,ESXi zeroes out pages allocated for virtual machines, userspace applications, and kernel threads at the time of allocation. This ensures that no non-zero pages are exposed to virtual machines or userspace applications. This measure is in place to prevent the exposure of cryptographic keys from virtual machines or userworlds to other clients.

However, if memory is not reused, these keys can remain present in host memory for an extended period. To address this, you can configure the `MemEagerZero` setting to enforce the zeroing out of userworld and guest memory pages when a userworld process or guest exits. For kernel threads, memory spaces holding keys are zeroed out as soon as the secret is no longer required.

Values	Installation Default Value: 0 Baseline Suggested Value: 1
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Additional shutdown time is required for virtual machines, corresponding to the amount of allocated memory.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Mem.MemEagerZero</pre>

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
Mem.MemEagerZero | Set-AdvancedSetting -Value 1
```

Check for Active Maintenance on the ESXi Version

Ensure that the ESXi version has not reached VMware End of General Support status.

Values

Installation Default Value: N/A
Baseline Suggested Value: N/A

Action Needed

Audit your ESXi version periodically.

Potential Functional Impact if Default Value Is Changed

Always read the Release Notes, and test and deploy new software versions by using staged roll-outs.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Select-Object Name,Version,Build
```

PowerCLI Command Remediation Example

N/A

Activate Time Synchronization Sources

The ESXi host must have time synchronization services activated and running.

Cryptography, audit logging, cluster operations, and incident response and forensics rely on synchronized time. To ensure time is synchronized across services and operations, activate NTP and/or PTP services to start with the host, and ensure that those services are running.

Values

Installation Default Value: Stopped, Start and stop manually
Baseline Suggested Value Value: Running, Stop and start with host

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHostService -VMHost $ESXi | Where-Object {$_.Key -eq "ntpd"}
```

PowerCLI Command Remediation Example

```
Get-VMHostService -VMHost $ESXi -ErrorAction:Stop |
Where-Object {$_.Key -eq "ntpd"} | Set-VMHostService
-policy "on" -Confirm:$false
Get-VMHostService -VMHost $ESXi -ErrorAction:Stop |
Where-Object {$_.Key -eq "ntpd"} | Restart-
VMHostService -Confirm:$false
```

Configure Reliable Time Synchronization Sources

The ESXi host must have reliable time synchronization sources configured.

Cryptography, audit logging, cluster operations, and incident response and forensics depend on synchronized time. Network Time Protocol (NTP) must have at least four sources. If you must choose between two sources and one source, one source is preferable.

Precision Time Protocol (PTP) is an alternative to NTP that provides sub-millisecond time accuracy. The architecture of PTP is different than NTP, and does not have the same resilience to primary server failure. Consider configuring NTP as a backup source to PTP so that a time source continues to be available, even if the accuracy is less.

Values

Installation Default Value: Undefined

Baseline Suggested Value Value:

Site-Specific or:

0.vmware.pool.ntp.org,
1.vmware.pool.ntp.org,
2.vmware.pool.ntp.org,
3.vmware.pool.ntp.org

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHostNtpServer -VMHost $ESXi
```

PowerCLI Command Remediation Example

```
$ntp0 = "0.vmware.pool.ntp.org"
$ntp1 = "1.vmware.pool.ntp.org"
$ntp2 = "2.vmware.pool.ntp.org"
$ntp3 = "3.vmware.pool.ntp.org"

Add-VMHostNTPServer -NtpServer $ntp0 , $ntp1 ,
    $ntp2 , $ntp3 -VMHost $ESXi -Confirm:$false
```

Use TLS Ciphers

The ESXi host must maintain confidentiality and integrity of transmissions by enabling modern TLS ciphers.

Starting in ESXi 8.0 Update 3, TLS profiles configure client and server TLS settings to use only strong ciphers. You can view the entire cipher list and suites by using the following commands:

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.tls.server.get.CreateArgs()
$arguments.showprofiledefaults = $true
$arguments.showcurrentbootprofile = $true
$ESXcli.system.tls.server.get.invoke($arguments)
```

You must reboot the ESXi host after making changes to the TLS profile. (In the vSphere Client, the host displays the suffix of "Reboot Required.")

Values

Installation Default Value: COMPATIBLE
Baseline Suggested Value: NIST_2024

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

Changes to cipher suites impact connectivity with external systems. You must reboot the host for this TLS profile change to take effect.

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.tls.server.get.invoke() | Select-Ob-
ject -ExpandProperty Profile
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.tls.server.set.Cre-
ateArgs()
$arguments.profile = "NIST_2024"
$ESXcli.system.tls.server.set.invoke($arguments)
```

UserVars.ESXiVPsDisabledProtocols

The ESXi host must enable the highest version of TLS supported.

ESXi 8.0 activates TLS 1.2 by default, but it is possible to activate other protocols if necessary. Starting with ESXi 8.0 Update 3, TLS 1.3 is activated by default.

Values	Installation Default Value: sslv3,tls1,tls1.1 Baseline Suggested Value Value: sslv3,tls1,tls1.1
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting User- Vars.ESXiVPsDisabledProtocols</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting User- Vars.ESXiVPsDisabledProtocols Set-AdvancedSetting -Value "sslv3,tls1,tls1.1"</pre>

Configure TPM-Based Encryption

The ESXi host must require TPM-based configuration encryption.

The configuration of an ESXi host consists of configuration files for each service that runs on the host. The configuration files typically reside in the `/etc` directory, but they can also reside in other namespaces. The configuration files contain run-time information about the state of the services. Over time, the default values in the configuration files might change, for example, when settings on the ESXi host are changed.

A cron job backs up the ESXi configuration files periodically, when ESXi shuts down gracefully or on demand, and creates an archived configuration file in the boot bank. When ESXi reboots, the system reads the archived configuration file and recreates the state that ESXi was in when the backup was taken.

Before vSphere 7.0 Update 2, the archived ESXi configuration file is not encrypted. In vSphere 7.0 Update 2 and later, the archived configuration file is encrypted. When the ESXi host is configured with a Trusted Platform Module (TPM), the TPM is used to "seal" the configuration to the host, providing a strong security guarantee and additional protection from offline attacks.

Configuration encryption uses the physical TPM when it is available and supported at install or upgrade time. If the TPM was added or enabled later, you must explicitly reconfigure the ESXi host to use the newly available TPM. Once the TPM configuration encryption is enabled, it cannot be deactivated.

Values	Installation Default Value: Site-specific Baseline Suggested Value: TPM
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Use of Secure Boot and TPM-enforced configuration encryption render traditional root password recovery efforts unusable. Ensure that you do not lose access to ESXi administrator accounts.
PowerCLI Command Assessment	<pre>\$ESXcli = Get-ESxCli -VMHost \$ESXi -V2 \$ESXcli.system.settings.encryption.get.Invoke() Select Mode</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-ESxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.settings.encryp- tion.set.CreateArgs() \$arguments.mode = "TPM"</pre>

```
$ESXcli.system.settings.encryption.set.Invoke($arguments)
```

Check thatESXi Software Is Up-to-date

By staying current onESXi patches, you mitigate vulnerabilities in the hypervisor.

An educated attacker can exploit known vulnerabilities when attempting to access or elevate privileges on anESXi host. Always updatevCenter Server first, if an update is available, then updateESXi.

Values	Installation Default Value: Downlevel Baseline Suggested Value: Current
Action Needed	AuditESXi patch level periodically. Always read the Release Notes, and test and deploy new software versions by using staged roll-outs.
Potential Functional Impact if Default Value Is Changed	vSphere Update releases add and change functionality. Patch releases only resolve issues.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Select-Object Name,Version,Build</pre>
PowerCLI Command Remediation Example	N/A

VMkernel.Boot.execInstalledOnly

Run binaries that are only delivered by a VIB.

ESXi conducts integrity checks of VIBs based upon the Acceptance Level. InstructingESXi to only execute binaries that originate from a valid VIB installed on the host makes it harder for attackers to use pre-built toolkits during a compromise, and increases chances of detection.

Values	Installation Default Value: False Baseline Suggested Value: True
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Third-party unsigned software might not install or execute.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting VMkernel.Boot.execInstalledOnly</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting VMkernel.Boot.execInstalledOnly Set-AdvancedSetting -Value True</pre>

Deactivate Management Services on VMkernel Adapters

Ensure thatvSAN, vMotion, and other dedicated VMkernel adapters do not have management services activated.

VMkernel network interfaces that are intended for specialized use can be configured with management capabilities, which might defeat network isolation and security efforts. Only enable management services on VMkernel interfaces that are intended for management.

Values	Installation Default Value: Site-specific Baseline Suggested Value: Site-specific
Action Needed	Audit your site-specific values.

Potential Functional Impact if Default Value Is Changed

Some third-party managed solutions might require that you activate management services on VMkernel adapters.

PowerCLI Command Assessment

```
Get-VMHostNetworkAdapter -VMHost $ESXi -VMKernel |
Select VMHost,Name,IP,ManagementTrafficEnabled
```

PowerCLI Command Remediation Example

```
Get-VMHostNetworkAdapter -VMHost $ESXi -Name $vmkernel_interface | Set-VMHostNetworkAdapter -ManagementTrafficEnabled $false
```

Configure theESXi Firewall to Block Traffic

You must configure theESXi host firewall to block network traffic by default.

Ensure that all incoming and outgoing network traffic is blocked unless explicitly allowed, reducing the attack surface and preventing unauthorized access to the host.

Values

Installation Default Value: Activated
Baseline Suggested Value: Activated

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

Firewall is simplistic, akin to router ACLs. You might need to reconfigure reflexive rules.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-VMHostFirewallDefaultPolicy
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.network.firewall.set.CreateArgs()
$arguments.defaultaction = $FALSE
$arguments.enabled = $true
$ESXcli.network.firewall.set.Invoke($arguments)
```

Configure theESXi Firewall for Authorized Networks

Configure theESXi firewall to allow traffic only from authorized networks.

Ensure that all incoming and outgoing network traffic is blocked unless explicitly allowed, reducing the attack surface and preventing unauthorized access to theESXi host.

Starting in vSphere 8.0 Update 2, firewall rules are categorized as 'user' or 'system' owned, where only 'user' owned rules are configurable. In vSphere 8 Update 2b and PowerCLI 13.2.1, there are additional queryable parameters to automate setting and checking for configurable rules.

Values

Installation Default Value: Connections allowed from any IP address
Baseline Suggested Value: Connections allowed only from authorized infrastructure and administration workstations

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

Firewall is simplistic, akin to router ACLs. You might need to reconfigure reflexive rules.

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
```

```
$list = $ESXcli.network.firewall.ruleset.list.Invoke() | Where {(($_.AllowedIPconfigurable -eq $true) -and ($_.EnableDisableconfigurable -eq $true))} |
Select -ExpandProperty Name
```

```
$arguments = $ESXcli.network.firewall.ruleset.allowedip.list.CreateArgs()
foreach ($rule in $list) {
    $arguments.rulesetid = $rule
    $ESXcli.network.firewall.ruleset.allowedip.list.Invoke($arguments)
}
```

PowerCLI Command Remediation Example

```
# Customize this example for your environment.
$ESXcli = Get-ESXcli -VMHost $ESXi -V2
# Deactivate firewall temporarily so we don't lose connectivity
$arguments = $ESXcli.network.firewall.set.CreateArgs()
$arguments.enabled = $false
$ESXcli.network.firewall.set.Invoke($arguments)

# Unset the "allow all" flag
$arguments = $ESXcli.network.firewall.ruleset.set.CreateArgs()
$arguments.allowedall = $false
$arguments.rulesetid = "sshServer"
$ESXcli.network.firewall.ruleset.set.Invoke($arguments)

# Add an IP range
$arguments = $ESXcli.network.firewall.ruleset.allowedip.add.CreateArgs()
$arguments.ipaddress = "192.168.0.0/16"
$arguments.rulesetid = "sshServer"
$ESXcli.network.firewall.ruleset.allowedip.add.Invoke($arguments)

# Enable the firewall
$arguments = $ESXcli.network.firewall.set.CreateArgs()
$arguments.enabled = $true
$ESXcli.network.firewall.set.Invoke($arguments)
```

Set the Forged Transmits Policy to Reject

Set the Forged Transmits policy to Reject on both the vSphere Standard Switch and its port groups.

If the virtual machine operating system changes the MAC address, the operating system can send frames with an impersonated source MAC address at any time. MAC address impersonation allows an operating system to stage malicious attacks on the devices in a network by impersonating a network adaptor authorized by the receiving network. When the Forged Transmits policy is set to Accept, ESXi does not compare source and effective MAC addresses. To protect against MAC impersonation, set the Forged Transmits policy to Reject. The host then compares the source MAC

address being transmitted by the guest operating system with the effective MAC address for its virtual machine adapter for a match. If the addresses do not match, the ESXi host drops the packet.

Values	Installation Default Value: Accept Baseline Suggested Value: Reject
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Some workloads, such as clustered applications and network devices and functions, rely on these techniques as a normal part of their operation. If needed, you can configure a separate port group that permits this behavior, and attach only authorized virtual machines to it.

PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-VirtualSwitch -Standard Get-SecurityPolicy select VirtualSwitch,Forged- Transmits Get-VMHost -Name \$ESXi Get-VirtualPortGroup - Standard Get-SecurityPolicy select VirtualPort- Group,ForgedTransmits</pre>
------------------------------------	---

PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-VirtualSwitch -Standard Get-SecurityPolicy Set-SecurityPolicy -Forged- Transmits \$false Get-VMHost -Name \$ESXi Get-VirtualPortGroup -Stan- dard Get-SecurityPolicy Set-SecurityPolicy - ForgedTransmitsInherited \$true</pre>
---	--

Set the MAC Address Changes Policy to Reject

Set the MAC Address Changes policy to Reject on both the vSphere Standard Switch and its port groups.

If the virtual machine operating system changes the MAC address, it can send frames with an impersonated source MAC address, enabling it to stage malicious attacks on devices within a network by impersonating a network adapter authorized by the receiving network. To prevent virtual machines from changing their effective MAC address, measures should be taken to enforce MAC address stability or restrict the ability to modify MAC addresses. This helps mitigate the risk of MAC impersonation and potential malicious activities.

Values	Installation Default Value: Accept Baseline Suggested Value: Reject
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Some workloads, such as clustered applications, network devices and functions, applications licensed by MAC address, and vCenter Server Reduced Downtime Upgrade, rely on these techniques as a normal part of their operation. If needed, you can configure a separate port group that permits this behavior, and attach only authorized virtual machines to it.

PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-VirtualSwitch -Standard Get-SecurityPolicy select VirtualSwitch,MacCha- nges Get-VMHost -Name \$ESXi Get-VirtualPortGroup - Standard Get-SecurityPolicy select VirtualPort- Group,MacChanges</pre>
------------------------------------	---

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-VirtualSwitch -Standard
| Get-SecurityPolicy | Set-SecurityPolicy -MacChanges $false
Get-VMHost -Name $ESXi | Get-VirtualPortGroup -Standard
| Get-SecurityPolicy | Set-SecurityPolicy -MacChangesInherited $true
```

Set the Promiscuous Mode Policy to Reject

Set the Promiscuous Mode policy to Reject on both the vSphere Standard Switch and its port groups.

When promiscuous mode is enabled for a port group, all virtual machines connected to that port group have the potential to read all packets transmitted across that port group, regardless of the intended recipient. Consider the potential impact and design considerations before changing the default value of promiscuous mode.

Values

Installation Default Value: Reject

Baseline Suggested Value: Reject

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

Certain workloads and types of work, such as DHCP servers, network devices, and security monitoring, incorporate these techniques as a regular part of their operation. If needed, you can configure a separate port group that permits this behavior, and attach only authorized virtual machines to it.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-VirtualSwitch -Standard
| Get-SecurityPolicy | select VirtualSwitch,AllowPromiscuous
Get-VMHost -Name $ESXi | Get-VirtualPortGroup -Standard
| Get-SecurityPolicy | select VirtualPortGroup,AllowPromiscuous
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-VirtualSwitch -Standard
| Get-SecurityPolicy | Set-SecurityPolicy -AllowPromiscuous $false
Get-VMHost -Name $ESXi | Get-VirtualPortGroup -Standard
| Get-SecurityPolicy | Set-SecurityPolicy -AllowPromiscuousInherited $true
```

Restrict Virtual Guest Tagging on Standard Switches

The ESXi host must restrict the use of Virtual Guest Tagging (VGT) on Standard Switches.

When a port group is set to VLAN 4095, the vSwitch passes all network frames to the attached virtual machines without modifying the VLAN tags. In vSphere, this is referred to as VGT. The virtual machine must process the VLAN information itself by using an 802.1Q driver in the operating system.

VLAN 4095 must only be implemented if the attached virtual machines have been specifically authorized and are capable of managing VLAN tags themselves. If VLAN 4095 is enabled inappropriately, it might cause denial of service or allow a virtual machine to interact with traffic on an unauthorized VLAN.

Values

Installation Default Value: Not VLAN 4095

Baseline Suggested Value: Not VLAN 4095

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-VirtualPortGroup -Standard | select Name,VlanID
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-VirtualPortGroup -Standard -Name $PG | Set-VirtualPortGroup -VlanID "new_VLAN"
```

Activate Secure Boot Enforcement

Secure Boot is part of the UEFI firmware standard. With UEFI Secure Boot activated, an ESXi host refuses to load any UEFI driver or app unless the operating system bootloader has a valid digital signature. Secure Boot for ESXi requires support from the firmware. Secure Boot for ESXi also requires that all ESXi kernel modules, drivers, and VIBs be signed by VMware or a partner subordinate.

Secure Boot is activated in the BIOS of the ESXi physical server and supported by the hypervisor boot loader. This control flips ESXi from merely supporting Secure Boot to requiring it. Without this setting activated, and configuration encryption, an ESXi host could be subject to offline attacks. An attacker could simply transfer the ESXi install drive to a non-Secure Boot host and boot.

Values

Installation Default Value: False
Baseline Suggested Value: True

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

Use of Secure Boot and TPM-enforced configuration encryption render traditional root password recovery efforts unusable. Ensure that you do not lose access to ESXi administrator accounts.

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.settings.encryption.get.Invoke() |
Select RequireSecureBoot
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.settings.encryption.set.CreateArgs()
$arguments.requiressecureboot = $true
$ESXcli.system.settings.encryption.set.Invoke($arguments)
```

Deactivate the ESXi Shell

The ESXi Shell should be deactivated.

Values

Installation Default Value: Stopped, Start and stop manually
Baseline Suggested Value: Stopped, Start and stop manually

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHostService -VMHost $ESXi | Where-Object {
    $_.Key -eq 'TSM' -and $_.Running -eq 'True'
}
Get-VMHostService -VMHost $ESXi | Where-Object {
    $_.Key -eq 'TSM' -and $_.Policy -eq 'On'
}
```

PowerCLI Command Remediation Example

```
Get-VMHostService -VMHost $ESXi | where { $_.Key -eq 'TSM' } | Set-VMHostService -Policy Off
```

```
Get-VMHostService -VMHost $ESXi | where {$_.Key -eq
'TSM'} | Stop-VMHostService
```

UserVars.ESXiShellInteractiveTimeOut

Sets a timeout to automatically terminate idleESXi Shell and SSH sessions.

If users forget to log out of their SSH session, the idle connection remains open indefinitely, increasing the potential for other users to gain privileged access to the host. You can configure idle shell sessions to terminate automatically.

Values	Installation Default Value: 0 Baseline Suggested Value: 900
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting User- Vars.ESXiShellInteractiveTimeOut</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting User- Vars.ESXiShellInteractiveTimeOut Set-AdvancedSet- ting -Value 900</pre>

Deactivate the SNMP Service

Deactivate the SNMP service if you are not using it.

Values	Installation Default Value: Stopped, Start and stop with host Baseline Suggested Value: Stopped, Start and stop manually
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHostService -VMHost \$ESXi Where-Object {\$_.Key -eq 'snmpd' -and \$_.Running -eq 'True'} Get-VMHostService -VMHost \$ESXi Where-Object {\$_.Key -eq 'snmpd' -and \$_.Policy -eq 'On'}</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHostService -VMHost \$ESXi where {\$_.Key -eq 'snmpd'} Set-VMHostService -Policy Off Get-VMHostService -VMHost \$ESXi where {\$_.Key -eq 'snmpd'} Stop-VMHostService</pre>

Deactivate the SSH Service

Deactivate SSH, and activate it only for troubleshooting.

ESXi is not a UNIX-like multi-user operating system.ESXi is a purpose-built hypervisor intended to be managed by theVMware Host Client, thevSphere Client, the CLIs, and the APIs. OnESXi, SSH is a troubleshooting and support interface, and is intentionally stopped and deactivated by default. Activation of the interface brings risk.

Values	Installation Default Value: Stopped, Start and stop manually Baseline Suggested Value: Stopped, Start and stop manually
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None

PowerCLI Command Assessment

```
Get-VMHostService -VMHost $ESXi | Where-Object
{$_ .Key -eq 'TSM-SSH' -and $_.Running -eq 'True'}
Get-VMHostService -VMHost $ESXi | Where-Object
{$_ .Key -eq 'TSM-SSH' -and $_.Policy -eq 'On'}
```

PowerCLI Command Remediation Example

```
Get-VMHostService -VMHost $ESXi | where {$_ .Key -eq
'TSM-SSH'} | Set-VMHostService -Policy Off
Get-VMHostService -VMHost $ESXi | where {$_ .Key -eq
'TSM-SSH'} | Stop-VMHostService
```

Use Entropy for Cryptographic Operations

The ESXi host must use sufficient entropy for cryptographic operations.

In vSphere 8.0 and later, the ESXi Entropy implementation supports the FIPS 140-3 and EAL4 certifications. Kernel boot options control which entropy sources to activate on an ESXi host.

Values

Installation Default Value:

disableHwrng = FALSE

entropySources = 0

Baseline Suggested Value:

disableHwrng = FALSE

entropySources = 0

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.settings.kernel.list.Invoke() | Where
{$_ .Name -eq "disableHwrng" -or $_.Name -eq "en-
tropySources"}
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.settings.kernel.set.Cre-
ateArgs()
$arguments.setting = "disableHwrng"
$arguments.value = "FALSE"
$ESXcli.system.settings.kernel.set.invoke($argu-
ments)
$arguments.setting = "entropySources"
$arguments.value = "0"
$ESXcli.system.settings.kernel.set.invoke($argu-
ments)
```

Verify the Image Profile and the VIB Acceptance Levels

The ESXi host image profile acceptance level must be PartnerSupported or higher.

The acceptance level controls what ESXi permits to be installed. See [Manage the Acceptance Levels of ESXi Hosts and vSphere Installation Bundles](#) for the VIB levels.

Neither VMware nor VMware partners test CommunitySupported VIBs, nor do CommunitySupported VIBs contain a digital signature. For these reasons, exercise caution when installing CommunitySupported VIBs.

Values	Installation Default Value: PartnerSupported Baseline Suggested Value: PartnerSupported or Higher
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	CommunitySupported packages are unsigned and are not able to be installed.
PowerCLI Command Assessment	<pre>(Get-EsxCli -VMHost \$ESXi -V2).software.acceptance.get.Invoke()</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.software.acceptance.set.CreateArgs() \$arguments.level = "PartnerSupported" # VMwareCertified, VMwareAccepted, PartnerSupported, CommunitySupported \$ESXcli.software.acceptance.set.Invoke(\$arguments)</pre>

Security.AccountUnlockTime

The ESXi host must unlock accounts after a specified timeout period.

`Security.AccountUnlockTime` ensures that user accounts on ESXi host are automatically unlocked after a defined period of inactivity. By enforcing automatic account unlocking, organizations can maintain a balance between security and usability, ensuring that idle accounts are reactivated promptly while minimizing the potential for unauthorized access.

Values	Installation Default Value: 900 seconds Baseline Suggested Value: 900 seconds
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Security.AccountUnlockTime</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Security.AccountUnlockTime Set-AdvancedSetting -Value 900</pre>

Security.AccountLockFailures

Sets the count of maximum failed login attempts before the account is locked out.

Protects against brute-force attacks and unauthorized access attempts by temporarily disabling the affected account, preventing further login attempts until the lockout period expires or is manually reset by an administrator. To unlock a locked account requires either administrative action or waiting for the account to automatically unlock if the `Security.AccountUnlockTime` setting is used.

Values	Installation Default Value: 5 Baseline Suggested Value: 5
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	A low threshold for login failures can potentially increase denial-of-service attacks, whether intentional or unintentional, such as with SSH connection retries.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Security.AccountLockFailures
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Security.AccountLockFailures | Set-AdvancedSetting -Value 5
```

Security.PasswordHistory

Does not permit password reuse.

This setting prevents the reuse of previous passwords, thus mitigating potential breaches from old, compromised credentials.

Values

Installation Default Value: 5

Baseline Suggested Value: 5

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Security.PasswordHistory
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Security.PasswordHistory | Set-AdvancedSetting -Value 5
```

Security.PasswordMaxDays

Sets the maximum number of days between password changes.

Modern best practices for passwords, as outlined in NIST 800-63B Section 5.1.1.2 and other relevant guidance, state that enforcing periodic password changes does not enhance security when passwords already possess adequate entropy.

Values

Installation Default Value: 99999

Baseline Suggested Value: 99999

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Security.PasswordMaxDays
```

PowerCLI Command

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Security.PasswordMaxDays | Set-AdvancedSetting -Value 99999
```

Security.PasswordQualityControl

Enforces password complexity.

Recommendations such as NIST 800-63B Section 5.1.1.2 suggest that composition rules, for example, mandating mixtures of character classes, should not be enforced on systems as they often fail to enhance password security and discourage the adoption of more secure passphrases.

Password strength and complexity rules are applicable to all ESXi users, including the root user. However, when the ESXi host is joined to a domain, these rules do not apply to Active Directory (AD) users because password policies for AD users are enforced by the AD system.

Values	Installation Default Value: retry=3 min=disabled,disabled,disabled,7,7 Baseline Suggested Value: retry=3 min=disabled,15,15,15,15 max=64 similar=deny passphrase=3
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Other products and services within the VMware ecosystem might not expect changes to password complexity requirements and might fail installation.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Security.PasswordQualityControl</pre>
PowerCLI Command	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Security.PasswordQualityControl Set-AdvancedSetting -Value "retry=3 min=disabled,15,15,15,15 max=64 similar=deny passphrase=3"</pre>

UserVars.SuppressHyperthreadWarning

Suppresses the warning for a potential hyperthreading security vulnerability.

Hyperthreading security warnings signify unaddressed CPU vulnerabilities in the system. Ignoring these warnings could mask potential risks. Ensure that hardware remediations align with accepted risk of your organization. When you do suppress a warning, document the decision and rationale.

Values	Installation Default Value: 0 Baseline Suggested Value: 0
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting UserVars.SuppressHyperthreadWarning</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting UserVars.SuppressHyperthreadWarning Set-AdvancedSetting -Value 0</pre>

UserVars.DcuiTimeOut

Sets a timeout to automatically terminate idle DCUI sessions.

DCUI allows direct login to the ESXi host for management tasks. To prevent unintended DCUI usage from leftover login sessions, terminate idle connections.

Values	Installation Default Value: 600 Baseline Suggested Value: 600
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting UserVars.DcuiTimeOut</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting UserVars.DcuiTimeOut Set-AdvancedSetting -Value 600</pre>

Deactivate CIM Service

TheESXi CIM service should be deactivated.

Services that are not in use and are non-essential for operations should be deactivated.

Values	Installation Default Value: Stopped, Start and stop with host Baseline Suggested Value: Stopped, Start and stop manually
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	

```
Get-VMHostService -VMHost $ESXi | Where-Object
{$_ .Key -eq 'sfcdb-watchdog' -and $_.Running -eq
'True'}
Get-VMHostService -VMHost $ESXi | Where-Object
{$_ .Key -eq 'sfcdb-watchdog' -and $_.Policy -eq
'On'}
```

PowerCLI Command Remediation Example

```
Get-VMHostService -VMHost $ESXi | where {$_ .Key -eq
'sfcdb-watchdog'} | Set-VMHostService -Policy Off
Get-VMHostService -VMHost $ESXi | where {$_ .Key -eq
'sfcdb-watchdog'} | Stop-VMHostService
```

Config.HostAgent.log.level

Sets the logging informational level.

When setting the log level, ensure that enough information is present in audit logs to conduct diagnostics and forensics.

Values	Installation Default Value: Info Baseline Suggested Value: Info
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Additional storage space is consumed by logs.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Con- fig.HostAgent.log.level</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Con- fig.HostAgent.log.level Set-AdvancedSetting -Value info</pre>

Syslog.global.logLevel

Logs sufficient information for events.

Without sufficient log data, critical indicators of compromise can go unnoticed, leading to increased vulnerability and potential failure to respond effectively to cybersecurity incidents.

Values	Installation Default Value: Error Baseline Suggested Value: Info
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Additional storage space is consumed by logs.
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Sys- log.global.logLevel</pre>

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Sys-
log.global.logLevel | Set-AdvancedSetting -Value in-
fo
```

Config.HostAgent.plugins.solo.enableMob

Deactivates the Managed Object Browser (MOB).

Services that are not in use and are non-essential for operations should be deactivated.

Values

Installation Default Value: False

Baseline Suggested Value: False

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Con-
fig.HostAgent.plugins.solo.enableMob
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Con-
fig.HostAgent.plugins.solo.enableMob | Set-Ad-
vancedSetting -Value False
```

Net.BlockGuestBPDU

Blocks guest operating system Bridge Protocol Data Unit (BPDU) transmissions.

BPDU s are used to transmit Spanning Tree Protocol (STP) information and detect network loops. BPDU Guard and Portfast are commonly activated on the physical switch directly connected to the ESXi host to reduce spanning tree convergence delay.

However, if a BPDU packet is sent from a virtual machine on the ESXi host to the configured physical switch, it can result in a cascading lockout of all uplink interfaces from the ESXi host. To prevent this type of lockout, you can activate the BPDU Filter on the ESXi host to drop any BPDU packets being sent to the physical switch.

Standard and Distributed Virtual Switches do not support STP and do not generate BPDU s.

Values

Installation Default Value: 1

Baseline Suggested Value: 1

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

Some network-oriented workloads can legitimately generate BPDU packets. Verify that there are no legitimate BPDU packets generated by virtual machines on the ESXi host prior to enabling BPDU Filter. If BPDU Filter is activated in this situation, enabling Reject Forged Transmits on the virtual switch port group adds protection against Spanning Tree loops.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
Net.BlockGuestBPDU
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
Net.BlockGuestBPDU | Set-AdvancedSetting -Value 1
```

Net.DVFilterBindIpAddress

Restricts the use of the dvFilter network APIs.

If you do not use a product such as VMware NSX that utilizes the dvFilter network API, do not configure the ESXi host to send network information to an IP address. Enabling the API and referencing an IP address that is compromised could potentially provide unauthorized access to the network of other virtual machines on the ESXi host.

If you are using a product that relies on this API, it is important to verify that the ESXi host has been configured correctly to ensure secure network communication.

Values	Installation Default Value: "" Baseline Suggested Value: ""
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Net.DV-FilterBindIpAddress</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting Net.DV-FilterBindIpAddress Set-AdvancedSetting -Value ""</pre>

UserVars.ESXiShellTimeout

Sets a timeout to limit how long the ESXi Shell and SSH services are allowed to run.

This advanced system setting defines a window of time after which the ESXi Shell and SSH services are automatically terminated.

Values	Installation Default Value: 0 Baseline Suggested Value: 600
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting User-Vars.ESXiShellTimeout</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting User-Vars.ESXiShellTimeout Set-AdvancedSetting -Value 600</pre>

UserVars.SuppressShellWarning

Suppresses the warning for support and troubleshooting interfaces.

The ESXi host must not suppress warnings that the ESXi Shell is activated.

Warnings indicating that SSH or the ESXi Shell is activated can be clues that an attack is in progress. It is important to ensure that SSH and the ESXi Shell are deactivated, and that this advanced system setting is not activated.

Values	Installation Default Value: 0 Baseline Suggested Value: 0
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-AdvancedSetting User-Vars.SuppressShellWarning</pre>

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting User-
Vars.SuppressShellWarning | Set-AdvancedSetting -
Value 0
```

Configure theESXi Secure Shell Daemon for FIPS

TheESXi host Secure Shell (SSH) daemon must be configured to only use FIPS 140-2/140-3 validated ciphers. You must harden and secure system services when they are activated.

Values

Installation Default Value: aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr
 Baseline Suggested Value: aes256-gcm@openssh.com,aes128-gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.ssh.server.config.list.invoke() |
Where-Object {$_.Key -eq 'ciphers'} | Select-Object
-ExpandProperty Value
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.ssh.server.con-
fig.set.CreateArgs()
$arguments.keyword = 'ciphers'
$arguments.value = 'aes256-gcm@openssh.com,aes128-
gcm@openssh.com,aes256-ctr,aes192-ctr,aes128-ctr'
$ESXcli.system.ssh.server.config.set.Invoke($argu-
ments)
```

Configure theESXi SSH Daemon for FIPS

TheESXi host SSH daemon must use FIPS 140-2/140-3 validated cryptographic modules.

OpenSSH on theESXi host ships with a FIPS 140-2/140-3 validated cryptographic module, activated by default. For backward compatibility reasons, you can deactivate this module. Audit and correct if necessary.

Values

Installation Default Value: True
 Baseline Suggested Value: True

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.security.fips140.ssh.get.Invoke()
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.security.fip-
s140.ssh.set.CreateArgs()
$arguments.enable = $true
$ESXcli.system.security.fips140.ssh.set.Invoke($ar-
guments)
```

Configure theESXi Secure Shell Daemon to Not Allow Gateway Ports

TheESXi host Secure Shell (SSH) daemon must be configured to not allow gateway ports.

You must harden and secure system services when activated.

Values	Installation Default Value: No Baseline Suggested Value: No
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.ssh.server.config.list.invoke() Where-Object {\$_.Key -eq 'gatewayports'} Select- Object -ExpandProperty Value</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.ssh.server.con- fig.set.CreateArgs() \$arguments.keyword = 'gatewayports' \$arguments.value = 'no' \$ESXcli.system.ssh.server.config.set.Invoke(\$argu- ments)</pre>

Configure theESXi Secure Shell Daemon to Not Use Host-Based Authentication

TheESXi host Secure Shell (SSH) daemon must not allow host-based authentication.

You must harden and secure system services when activated.

Values	Installation Default Value: No Baseline Suggested Value: No
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.ssh.server.config.list.invoke() Where-Object {\$_.Key -eq 'hostbasedauthentication'} Select-Object -ExpandProperty Value</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.ssh.server.con- fig.set.CreateArgs() \$arguments.keyword = 'hostbasedauthentication' \$arguments.value = 'no' \$ESXcli.system.ssh.server.config.set.Invoke(\$argu- ments)</pre>

Configure theESXi Secure Shell Daemon to Set a Timeout Count

TheESXi host Secure Shell (SSH) daemon must set a timeout count on idle sessions.

You must harden and secure system services when activated. The timeout count, multiplied by the idle timeout interval, is the total number of seconds the session can be idle until it is disconnected.

Values	Installation Default Value: 3
---------------	-------------------------------

Action Needed	Baseline Suggested Value: 3 Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.ssh.server.config.list.invoke() Where-Object {\$_.Key -eq 'clientalivecountmax'} Select-Object -ExpandProperty Value</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.ssh.server.config.set.CreateArgs() \$arguments.keyword = 'clientalivecountmax' \$arguments.value = '3' \$ESXcli.system.ssh.server.config.set.Invoke(\$arguments)</pre>

Configure theESXi Secure Shell Daemon to Set a Timeout Interval

TheESXi host Secure Shell (SSH) daemon must set a timeout count on idle sessions.

You must harden and secure system services when activated. The timeout count, multiplied by the idle timeout interval, is the total number of seconds the session may be idle until it is disconnected.

Values	Installation Default Value: 200 Baseline Suggested Value: 200
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.ssh.server.config.list.invoke() Where-Object {\$_.Key -eq 'clientaliveinterval'} Select-Object -ExpandProperty Value</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.ssh.server.config.set.CreateArgs() \$arguments.keyword = 'clientaliveinterval' \$arguments.value = '200' \$ESXcli.system.ssh.server.config.set.Invoke(\$arguments)</pre>

Configure theESXi Secure Shell Daemon to Display a Login Banner

TheESXi host Secure Shell (SSH) daemon must display the system login banner before granting access to the system.

You must harden and secure system services when activated. You must set the `Config.Etc.issue` setting as well, to provide text to this banner.

Values	Installation Default Value: <code>/etc/issue</code> Baseline Suggested Value: <code>/etc/issue</code>
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.ssh.server.config.list.invoke() |
  Where-Object {$_.Key -eq 'banner'} | Select-Object
  -ExpandProperty Value
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.ssh.server.config.set.CreateArgs()
$arguments.keyword = 'banner'
$arguments.value = '/etc/issue'
$ESXcli.system.ssh.server.config.set.Invoke($arguments)
```

Configure theESXi Secure Shell Daemon to Ignore .rhosts Files

TheESXi host Secure Shell (SSH) daemon must ignore `.rhosts` files.

You must harden and secure system services when activated.

Values

Installation Default Value: Yes

Baseline Suggested Value: Yes

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.ssh.server.config.list.invoke() |
  Where-Object {$_.Key -eq 'ignorerhosts'} | Select-Object
  -ExpandProperty Value
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$arguments = $ESXcli.system.ssh.server.config.set.CreateArgs()
$arguments.keyword = 'ignorerhosts'
$arguments.value = 'yes'
$ESXcli.system.ssh.server.config.set.Invoke($arguments)
```

Configure theESXi Secure Shell Daemon to Deactivate Stream Local Forwarding

TheESXi host Secure Shell (SSH) daemon must deactivate stream local forwarding.

You must harden and secure system services when activated.

Values

Installation Default Value: No

Baseline Suggested Value: No

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
$ESXcli.system.ssh.server.config.list.invoke() |
  Where-Object {$_.Key -eq 'allowstreamlocalforwarding'} | Select-Object
  -ExpandProperty Value
```

PowerCLI Command Remediation Example

```
$ESXcli = Get-EsxCli -VMHost $ESXi -V2
```

```
$arguments = $ESXcli.system.ssh.server.config.set.CreateArgs()
$arguments.keyword = 'allowstreamlocalforwarding'
$arguments.value = 'no'
$ESXcli.system.ssh.server.config.set.Invoke($arguments)
```

Configure the ESXi Secure Shell Daemon to Deactivate TCP Forwarding

The ESXi host Secure Shell (SSH) daemon must deactivate TCP forwarding.

You must harden and secure system services when activated.

Values	Installation Default Value: No Baseline Suggested Value: No
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.ssh.server.config.list.invoke() Where-Object {\$_.Key -eq 'allowtcpforwarding'} Select-Object -ExpandProperty Value</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.ssh.server.config.set.CreateArgs() \$arguments.keyword = 'allowtcpforwarding' \$arguments.value = 'no' \$ESXcli.system.ssh.server.config.set.Invoke(\$arguments)</pre>

Configure the ESXi Secure Shell Daemon to Not Permit Tunnels

The ESXi host Secure Shell (SSH) daemon must not permit tunnels.

You must harden and secure system services when activated.

Values	Installation Default Value: No Baseline Suggested Value: No
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.ssh.server.config.list.invoke() Where-Object {\$_.Key -eq 'permtunnel'} Select-Object -ExpandProperty Value</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.ssh.server.config.set.CreateArgs() \$arguments.keyword = 'permtunnel' \$arguments.value = 'no' \$ESXcli.system.ssh.server.config.set.Invoke(\$arguments)</pre>

Configure the ESXi Secure Shell Daemon to Not Permit User Environment Settings

The ESXi host Secure Shell (SSH) daemon must not permit user environment settings.

You must harden and secure system services when activated.

Values	Installation Default Value: No Baseline Suggested Value: No
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.ssh.server.config.list.invoke() Where-Object {\$_.Key -eq 'permituserenvironment'} Select-Object -ExpandProperty Value</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.ssh.server.config.set.CreateArgs() \$arguments.keyword = 'permituserenvironment' \$arguments.value = 'no' \$ESXcli.system.ssh.server.config.set.Invoke(\$arguments)</pre>

Deactivate the Service Location Protocol Service

Deactivate the Service Location Protocol (SLP) service if you are not using it.

Values	Installation Default Value: Stopped, Start and stop manually Baseline Suggested Value: Stopped, Start and stop manually
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHostService -VMHost \$ESXi Where-Object {\$_.Key -eq 'slpd' -and \$_.Running -eq 'True'} Get-VMHostService -VMHost \$ESXi Where-Object {\$_.Key -eq 'slpd' -and \$_.Policy -eq 'On'}</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHostService -VMHost \$ESXi where {\$_.Key -eq 'slpd'} Set-VMHostService -Policy Off Get-VMHostService -VMHost \$ESXi where {\$_.Key -eq 'slpd'} Stop-VMHostService</pre>

Mem.ShareForceSalting

Restricts transparent page sharing to virtual machines that are configured with `sched.mem.pshare.salt`.

Transparent Page Sharing (TPS) is a method to reduce the memory footprint of virtual machines. Under highly controlled conditions, attackers can use TPS to gain unauthorized access to data on neighboring virtual machines. Virtual machines that do not have the `sched.mem.pshare.salt` setting configured cannot share memory with other virtual machines. Large page sizes, a performance optimization in the hypervisor on many modern CPUs, is incompatible with TPS.

Values	Installation Default Value: 2 Baseline Suggested Value: 2
Action Needed	Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
Mem.ShareForceSalting
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting
Mem.ShareForceSalting | Set-AdvancedSetting -Value
2
```

UserVars.HostClientSessionTimeout

Sets a timeout to automatically terminate idleESXi host client sessions.

TheESXi host must automatically terminate idle host client sessions. This practice helps mitigate potential security risks by ensuring that unattended sessions, which could be exploited by unauthorized users or malicious software, are not left open indefinitely.

Values

Installation Default Value: 900
Baseline Suggested Value: 900

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting User-
Vars.HostClientSessionTimeout
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting User-
Vars.HostClientSessionTimeout | Set-AdvancedSetting
-Value 900
```

Net.BMCNetworkEnable

Deactivates virtual hardware management network interfaces.

Hardware management controllers often present virtual or USB NICs to theESXi host. These can be used as backdoors and should be deactivated both in the hardware configuration and in theESXi configuration.

Values

Installation Default Value: 1
Baseline Suggested Value: 0

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

This functionality might be required by some third-party managed solutions.

PowerCLI Command Assessment

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Net.BM-
CNetworkEnable
```

PowerCLI Command Remediation Example

```
Get-VMHost -Name $ESXi | Get-AdvancedSetting Net.BM-
CNetworkEnable | Set-AdvancedSetting -Value 0
```

Activate Bidirectional/Mutual CHAP Authentication for iSCSI Traffic

Set the iSCSI storage adapter authentication to "Use bidirectional CHAP" and supply credentials.

Mutual CHAP provides an additional layer of protection by requiring both the initiator (client) and the target (server) to verify their identities to each other, thereby ensuring data transmitted between the two is not intercepted or altered by unauthorized entities.

Values	Installation Default Value: Not configured Baseline Suggested Value: Activated
Action Needed	Modify the installation default value if you use this feature.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VMHost -Name \$ESXi Get-VMHostHba Where {\$_ .Type -eq "Iscsi"} Select VMHost, Device, ChapType, @{N="CHAPName";E={\$_.AuthenticationProp- erties.ChapName}}</pre>
PowerCLI Command Remediation Example	<pre>Get-VMHost -Name \$ESXi Get-VMHostHba Where {\$_ .Type -eq "Iscsi"} Set-VMHostHbaparameters</pre> For information about parameters, see <i>ESXCLI Reference</i> .

Do Not Store Encryption Keys on ESXi Hosts Without Securing Physical Access

The ESXi host must not store encryption keys on the ESXi host itself without securing physical access to the host.

Key Persistence is a mechanism that uses a local Trusted Platform Module (TPM) to store Standard Key Provider keys, usually found only in an external Key Management System (KMS). While this setup can improve the management of dependencies, using Key Persistence changes the encryption risks. If an attacker steals the host, they have access to the encryption keys for data on that host, bypassing the access controls of the external KMS. Therefore, only use Key Persistence when you can assure the physical security of your hosts. If the physical hosts are not secure and an attacker can steal the host, the attacker also has the means to access and use encrypted workloads.

Key Persistence and vSphere Native Key Provider are often conflated because both store encryption data on hosts. However, the vSphere Native Key Provider does not use Key Persistence, so turning off Key Persistence does not affect it. Like Key Persistence, the vSphere Native Key Provider also requires careful consideration of physical security. See [vSphere System Design Security Controls Reference](#).

Values	Installation Default Value: False Baseline Suggested Value: False
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Default is the desired behavior. Change away from the default might negatively impact confidentiality in environments where physical access by attackers is possible.
PowerCLI Command Assessment	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$ESXcli.system.security.keypersistance.get.invoke() Select-Object -ExpandProperty Enabled</pre>
PowerCLI Command Remediation Example	<pre>\$ESXcli = Get-EsxCli -VMHost \$ESXi -V2 \$arguments = \$ESXcli.system.security.keypersis- tence.disable.CreateArgs() \$arguments.removeallstoredkeys = \$true \$ESXcli.system.security.keypersistance.disable.In- voke(\$arguments)</pre>

vCenter Server Security Controls Reference

These security controls provide a baseline set of vCenter Server security best practices. They are structured in a way that explains the benefits and tradeoffs of implementing the control. To make changes, you can use the vSphere Client, PowerCLI, or the vCenter Server Management Interface, depending on the control.

PowerCLI and Variables Used

Some of the PowerCLI examples used here requires that the VMware.vSphere.SsoAdmin module be installed.

The PowerCLI commands in this section use the following variables:

- \$VC="vcenter_server_name"
- \$VDS="vsphere_distributed_switch_name"
- \$VDPG="vsphere_distributed_port_group"

Set vSphere Client Inactivity Timeout

The vCenter Server must terminate vSphere Client sessions after 15 minutes of inactivity.

Idle vSphere Client sessions can be left open indefinitely if a user forgets to log out, thereby increasing the risk of unauthorized privileged access.

Values	Installation Default Value: 120 minutes Baseline Suggested Value: 15 minutes
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	N/A (No public API available)
PowerCLI Command Remediation Example	N/A (No public API available)
Setting Location in the vSphere Client	Administration > Client Configuration > Session timeout

Set Failed Login Attempts Interval

The vCenter Server must set the interval for counting failed login attempts to at least 15 minutes.

By limiting the number of failed login attempts, the risk of unauthorized access using user password guessing, otherwise known as brute-forcing, is reduced.

Values	Installation Default Value: 180 Baseline Suggested Value: 900
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<code>Get-SsoLockoutPolicy Select FailedAttemptIntervalSec</code>
PowerCLI Command Remediation Example	<code>Get-SsoLockoutPolicy Set-SsoLockoutPolicy -FailedAttemptIntervalSec 900</code>
Setting Location in the vSphere Client	Administration > Single Sign On > Configuration > Local Accounts > Lockout Policy

Configure the vSphere SSO Lockout Policy Maximum Attempts

The vCenter Server must lock an account after a specified number of failed login attempts.

Repeated failed logins for an account can signal security issues. To limit brute force attempts, lock the account after a certain threshold, balancing between avoiding automatic connection retries and potential denial-of-service attacks.

Values	Installation Default Value: 5 Baseline Suggested Value: 5
---------------	--

Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<code>Get-SsoLockoutPolicy Select MaxFailedAttempts</code>
PowerCLI Command Remediation Example	<code>Get-SsoLockoutPolicy Set-SsoLockoutPolicy -Max-FailedAttempts 5</code>
Setting Location in the vSphere Client	Administration > Single Sign On > Configuration > Local Accounts > Lockout Policy

Configure the vSphere SSO Lockout Policy Unlock Time

The vCenter Server must unlock accounts after a specified timeout period.

Repeated failed logins can suggest security threats. vCenter Server accounts should not automatically unlock when they have been locked due to multiple login failures. Ensure that you have your administrator@vsphere.local information and that it is valid.

Values	Installation Default Value: 300 Baseline Suggested Value: 0
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	There is the opportunity for denial-of-service when accounts do not automatically unlock.
PowerCLI Command Assessment	<code>Get-SsoLockoutPolicy Select AutoUnlockIntervalSec</code>
PowerCLI Command Remediation Example	<code>Get-SsoLockoutPolicy Set-SsoLockoutPolicy -AutoUnlockIntervalSec 0</code>
Setting Location in the vSphere Client	Administration > Single Sign On > Configuration > Local Accounts > Lockout Policy

Enforce Password Complexity

The vCenter Server must enforce password complexity.

Modern best practices for passwords (see NIST 800-63B Section 5.1.1.2, among other guidance) indicate that with adequate password entropy, security is not improved by arbitrarily requiring users to change their passwords at certain intervals. Many automated security tools and regulatory compliance frameworks do not reflect this guidance, and might override this recommendation.

Password strength and complexity rules apply to accounts created in vSphere SSO, including administrator@vsphere.local (or, if you specified a different domain during installation, administrator@mydomain). These rules do not apply to Active Directory users when vCenter Server is joined to a domain, because AD enforces those password policies.

Values	Installation Default Value: Maximum Length: 20 Minimum Length: 8 At least 1 special At least 2 alphabetic At least 1 uppercase At least 1 lowercase
---------------	---

At least 1 numeric
 3 identical adjacent
 Baseline Suggested Value:
 Maximum Length: 64
 Minimum Length: 15
 At least 1 special
 At least 2 alphabetic
 At least 1 uppercase
 At least 1 lowercase
 At least 1 numeric
 3 identical adjacent

Recommended Action

Modify the installation default values.

Potential Functional Impact if Default Value Is Changed

Other products and services within the VMware ecosystem might not expect changes to password complexity requirements and could fail installation.

PowerCLI Command Assessment

```
Get-SsoPasswordPolicy
```

PowerCLI Command Remediation Example

```
Get-SsoPasswordPolicy | Set-SsoPasswordPolicy -Min-
Length 15 -MaxLength 64 -MinNumericCount 1 -MinSpe-
cialCharCount 1 -MinAlphabeticCount 2 -MinUpper-
caseCount 1 -MinLowercaseCount 1 -MaxIdenticalAdja-
centCharacters 3
```

Setting Location in the vSphere Client

**Administration > Single Sign On > Configuration > Local
 Accounts > Password Policy**

Configure the Maximum Number of Days Between Password Changes

The vCenter Server must be configured with an appropriate maximum password age.

Modern best practices for passwords (see NIST 800-63B Section 5.1.1.2, among other guidance) indicate that with adequate password entropy, security is not improved by arbitrarily requiring users to change their passwords at certain intervals. Many automated security tools and regulatory compliance frameworks do not reflect this guidance, and might override this recommendation.

Values

Installation Default Value: 90
 Baseline Suggested Value: 99999

Recommended Action

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-SsoPasswordPolicy | Select PasswordLifetimeDays
```

PowerCLI Command Remediation Example

```
Get-SsoPasswordPolicy | Set-SsoPasswordPolicy -Pass-
wordLifetimeDays 9999
```

Setting Location in the vSphere Client

**Administration > Single Sign On > Configuration > Local
 Accounts > Password Policy**

Restrict Password Reuse

Configure the password history setting to restrict the reuse of passwords on the vCenter Server.

Password complexity guidelines sometimes cause users to reuse older passwords. Configuring the password history setting on the vCenter Server can help prevent this situation.

Values	Installation Default Value: 5 Baseline Suggested Value: 5
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<code>Get-SsoPasswordPolicy Select ProhibitedPreviousPasswordsCount</code>
PowerCLI Command Remediation Example	<code>Get-SsoPasswordPolicy Set-SsoPasswordPolicy -ProhibitedPreviousPasswordsCount 5</code>
Setting Location in the vSphere Client	Administration > Single Sign On > Configuration > Local Accounts > Password Policy

Configure the Login Banner Text for SSH Access

Configure the vCenter Server login banner text for access using SSH.

vCenter Server allows a login message, which deters intruders and communicates obligations to authorized users. This configuration establishes the text displayed when a client connects using SSH. The default text leaks information to attackers about system configuration and should be changed.

Values	Installation Default Value: VMware vCenter Server <i>version</i> Type: vCenter Server with an embedded Platform Services Controller Baseline Suggested Value: Consult your organization's legal advisers for text that is applicable to your environment. Example text: Authorized users only. Actual or attempted unauthorized use of this system is prohibited and may result in criminal, civil, security, or administrative proceedings and/or penalties. Use of this information system indicates consent to monitoring and recording, without notice or permission. Users have no expectation of privacy in any use of this system. Any information stored on, or transiting this system, or obtained by monitoring and/or recording, may be disclosed to law enforcement and/or used in accordance with Federal law, State statute, and organization policy. If you are not an authorized user of this system, exit the system at this time.
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<code>Get-AdvancedSetting -Entity \$VC -Name etc.issue</code>
PowerCLI Command Remediation Example	<code>Get-AdvancedSetting -Entity \$VC -Name etc.issue Set-AdvancedSetting -Value "Authorized users only. Actual or attempted unauthorized use of this system is prohibited and may result in criminal, civil, security, or administrative proceedings and/or</code>

penalties. Use of this information system indicates consent to monitoring and recording, without notice or permission. Users have no expectation of privacy in any use of this system. Any information stored on, or transiting this system, or obtained by monitoring and/or recording, may be disclosed to law enforcement and/or used in accordance with Federal law, State statute, and organization policy. If you are not an authorized user of this system, exit the system at this time."

Setting Location in the vSphere Client

Administration > Single Sign On > Configuration > Login Message

Set Task and Retention Interval

The vCenter Server must have task and event retention set to an appropriate interval.

vCenter Server retains task and event data, which ages out to save storage space. The age is configurable. This only impacts local storage of event data on the vCenter Server Appliance.

Values	Installation Default Value: 30 Baseline Suggested Value: 30
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	N/A (No public API available)
PowerCLI Command Remediation Example	N/A (No public API available)
Setting Location in the vSphere Client	N/A

Activate Remote Logging

Activate remote logging of vCenter Server events.

Remote logging to a central host enhances the security of vCenter Server by storing logs securely. Remote logging simplifies monitoring across hosts and supports aggregate analysis for detecting coordinated attacks. Centralized logging prevents tampering and serves as a reliable long-term audit record. The `vpdx.event.syslog.enabled` setting activates remote logging.

Values	Installation Default Value: True Baseline Suggested Value: True
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<code>Get-AdvancedSetting -Entity \$VC -Name vpdx.event-syslog.enabled</code>
PowerCLI Command Remediation Example	<code>Get-AdvancedSetting -Entity \$VC -Name vpdx.event-syslog.enabled Set-AdvancedSetting -Value true</code>
Setting Location in the vSphere Client	Select vCenter Server > Configure > Advanced Settings

Activate FIPS

The vCenter Server must activate FIPS-validated cryptography.

FIPS cryptography makes a number of changes to the system to remove weaker ciphers. Activating FIPS causes vCenter Server to restart.

Values	Installation Default Value: False Baseline Suggested Value: True
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	FIPS cryptography makes a number of changes to the system to remove weaker ciphers. Enabling FIPS causes vCenter Server to restart.
PowerCLI Command Assessment	Invoke-GetSystemGlobalFips
PowerCLI Command Remediation Example	<pre>\$spec = Initialize-SystemSecurityGlobalFipsUpdateSpec -Enabled \$true Invoke-SetSystemGlobalFips -SystemSecurityGlobalFipsUpdateSpec \$spec</pre>
Setting Location in the vSphere Client	See Activate and Deactivate FIPS on the vCenter Server Appliance .

Configure Audit Records

The vCenter Server must produce audit records that contain information to establish what type of events occurred.

It is important to ensure that sufficient information is present in audit logs for diagnostics and forensics purposes. The `config.log.level` setting configures audit records.

Values	Installation Default Value: Info Baseline Suggested Value: Info
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-AdvancedSetting -Entity \$VC -Name config.log.level</pre>
PowerCLI Command Remediation Example	<pre>Get-AdvancedSetting -Entity \$VC -Name config.log.level Set-AdvancedSetting -Value info</pre>
Setting Location in the vSphere Client	vCenter Server Host > Configure > Advanced Settings

Deactivate MAC Learning

All distributed switch port groups must deactivate MAC Learning unless used intentionally.

MAC Learning enables a distributed switch to provide network connectivity to systems where more than one MAC address is used on a vNIC. This can be useful in special cases like nested virtualization (running ESXi inside ESXi, for example). MAC learning also supports unknown unicast flooding. Normally, when a packet that is received by a port has an unknown destination MAC address, the packet is dropped. With unknown unicast flooding enabled, the port floods unknown unicast traffic to every port on the switch that has MAC learning and unknown unicast flooding enabled. This property is activated by default, but only if MAC learning is enabled. Deactivate MAC Learning unless it is in use intentionally for a known workload that requires it.

Values	Installation Default Value: Disabled
---------------	--------------------------------------

Recommended Action**Potential Functional Impact if Default Value Is Changed****PowerCLI Command Assessment****PowerCLI Command Remediation Example**

Baseline Suggested Value: Disabled

Audit the installation default value.

Some workloads do legitimately use these network tactics and are negatively affected by the defaults and desired state.

```
(Get-VDPortgroup -Name $VDPG).ExtensionData.Config.DefaultPortConfig.MacManagementPolicy.MacLearningPolicy | Select-Object -ExpandProperty Enabled
```

```
$VDPGview = Get-VDPortgroup -Name $VDPG | Get-View
$ConfigSpec = New-Object VMware.Vim.DVPortgroupConfigSpec
$ConfigSpec.DefaultPortConfig = New-Object VMware.Vim.VMwareDVSPortSetting
$ConfigSpec.DefaultPortConfig.MacManagementPolicy = New-Object VMware.Vim.DVSMacManagementPolicy
$ConfigSpec.DefaultPortConfig.MacManagementPolicy.MacLearningPolicy = New-Object VMware.Vim.DVSMacLearningPolicy
$ConfigSpec.DefaultPortConfig.MacManagementPolicy.MacLearningPolicy.Enabled = $false
$ConfigSpec.ConfigVersion = $VDPGview.Config.ConfigVersion
$VDPGview.ReconfigureDVPortgroup_Task($ConfigSpec)
```

Setting Location in the vSphere Client

N/A. MAC Learning can be enabled on a Distributed Virtual Port group using the vSphere API. See *vSphere Web Services API Reference* for more information.

Configure the Login Message Banner Details

Configure the vCenter Server login banner details for the vSphere Client.

vCenter Server provides the ability to display a login message. Usages of the login message include informing intruders that their activities are illegal, and conveying to authorized users the expectations and obligations they must meet and agree to while using the system. This configuration sets the detailed text from the vSphere Client login page message.

Values

Installation Default Value: Not configured

Baseline Suggested Value: Consult your organization's legal advisers for text that is applicable to your environment.

An example text: Authorized users only. Actual or attempted unauthorized use of this system is prohibited and may result in criminal, civil, security, or administrative proceedings and/or penalties. Use of this information system indicates consent to monitoring and recording, without notice or permission. Users have no expectation of privacy in any use of this system. Any information stored on, or transiting this system, or obtained by monitoring and/or recording, may be disclosed to law enforcement and/or used in accordance with Federal law, State statute, and organization policy. If you are not an authorized user of this system, exit the system at this time.

Recommended Action

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment	N/A (No public API available) You can configure the login message by running the following command in an appliance shell: <pre>/opt/vmware/bin/sso-config.sh -set_login_banner -title login_banner_titlelogonBannerFile</pre> Remember to deactivate the shell again when you are done.
PowerCLI Command Remediation Example	N/A (No public API available) You can configure the login message by running the following command in an appliance shell: <pre>/opt/vmware/bin/sso-config.sh -set_login_banner -title login_banner_titlelogonBannerFile</pre> Remember to deactivate the shell again when you are done.
Setting Location in the vSphere Client	Administration > Single Sign On > Configuration > Login Message > Edit
<u>Activate the Login Banner</u>	
Activate the vCenter Server login banner for the vSphere Client.	
vCenter Server provides the ability to display a login message. Usages of the login message include informing intruders that their activities are illegal, and conveying to authorized users the expectations and obligations they must meet and agree to while using the system. This configuration activates the display of the message on the vSphere Client login page.	
Values	Installation Default Value: False Baseline Suggested Value: True
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	N/A (No public API available) You can configure the login message by running the following command in an appliance shell: <pre>/opt/vmware/bin/sso-config.sh -set_logon_banner -title logon_banner_titlelogonBannerFile</pre> Remember to deactivate the shell again when you are done.
PowerCLI Command Remediation Example	N/A (No public API available) You can configure the login message by running the following command in an appliance shell: <pre>/opt/vmware/bin/sso-config.sh -set_logon_banner -title logon_banner_titlelogonBannerFile</pre> Remember to deactivate the shell again when you are done.
Setting Location in the vSphere Client	Administration > Single Sign On > Configuration > Login Message > Edit
<u>Configure the Login Banner Text</u>	
Configure the vCenter Server login banner text for the vSphere Client.	

vCenter Server provides the ability to display a login message. Usages of the login message include informing intruders that their activities are illegal, and conveying to authorized users the expectations and obligations they must meet and agree to while using the system. This configuration establishes the text displayed on the vSphere Client login page.

Values	Installation Default Value: Not configured Baseline Suggested Value: Consult your organization's legal advisers for specific text. Example text: Use of this system indicates awareness of, and consent to, the organizational policies governing this system.
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	N/A (No public API available) You can configure the login message by running the following command in an appliance shell: <pre>/opt/vmware/bin/sso-config.sh -set_logon_banner -title logon_banner_titlelogonBannerFile</pre> Remember to deactivate the shell again when you are done.
PowerCLI Command Remediation Example	N/A (No public API available) You can configure the login message by running the following command in an appliance shell: <pre>/opt/vmware/bin/sso-config.sh -set_logon_banner -title logon_banner_titlelogonBannerFile</pre> Remember to deactivate the shell again when you are done.
Setting Location in the vSphere Client	Administration > Single Sign On > Configuration > Login Message > Edit

Separate Authentication and Authorization for Administrators

The vCenter Server must separate authentication and authorization for administrators.

Combining authentication and authorization, as services such as Active Directory do, risks infrastructure breaches if compromised. Thus, for vCenter Server, ensure that you segregate authentication and authorization for administrators. Consider using local SSO groups for authorization to better manage risk where feasible.

Values	Installation Default Value: Not configured Baseline Suggested Value: Not configured
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Provisioning vCenter Server access required interaction with vCenter Server SSO. Automation is possible with PowerCLI.
PowerCLI Command Assessment	N/A (No public API available)
PowerCLI Command Remediation Example	N/A (No public API available)
Setting Location in the vSphere Client	N/A

Set the Forged Transmits Policy to Reject

Set all distributed switches and their port groups to reject forged transmits.

A virtual machine can impersonate network adapters by changing MAC addresses, posing security threats. By setting the Forged Transmits option to Reject on all distributed switches and port groups, ESXi verifies MAC addresses and prevents such impersonation.

Values	Installation Default Value: Reject Baseline Suggested Value: Reject
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Some workloads do legitimately use these network tactics and are affected negatively by default setting.
PowerCLI Command Assessment	<pre>Get-VDSwitch -Name \$VDS Get-VDSecurityPolicy Get-VDPortgroup -Name \$VDPG Get-VDSecurityPolicy</pre>
PowerCLI Command Remediation Example	<pre>Get-VDSwitch -Name \$VDS Get-VDSecurityPolicy Set-VDSecurityPolicy -ForgedTransmits \$false Get-VDPortgroup -Name \$VDPG Get-VDSecurityPolicy Set-VDSecurityPolicy -ForgedTransmits \$false</pre>
Setting Location in the vSphere Client	See the <i>vSphere Networking</i> documentation.

Set the MAC Address Changes Policy to Reject

Set the MAC Address Changes policy to Reject on both the vSphere Standard Switch and its port groups.

Allowing virtual machines to change MAC addresses poses security risks, enabling potential network adapter impersonation. Rejecting MAC changes on all distributed switches and port groups prevents this, but might impact certain applications like Microsoft Clustering or MAC address-dependent licensing. Make exceptions to this security guidance as necessary.

Values	Installation Default Value: Reject Baseline Suggested Value: Reject
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Some workloads do legitimately use these network tactics and are affected negatively by the a setting of Reject.
PowerCLI Command Assessment	<pre>Get-VDSwitch -Name \$VDS Get-VDSecurityPolicy Get-VDPortgroup -Name \$VDPG Get-VDSecurityPolicy</pre>
PowerCLI Command Assessment	<pre>Get-VDSwitch -Name \$VDS Get-VDSecurityPolicy Set-VDSecurityPolicy -MacChanges \$false Get-VDPortgroup -Name \$VDPG Get-VDSecurityPolicy Set-VDSecurityPolicy -MacChanges \$false</pre>
Setting Location in the vSphere Client	See the <i>vSphere Networking</i> documentation.

Set the Promiscuous Mode Policy to Reject

Set the Promiscuous Mode policy to Reject on both the vSphere Standard Switch and its port groups.

Activating promiscuous mode on a port group allows all connected virtual machines to read all network packets, posing a potential security risk. Although allowing promiscuous mode is sometimes necessary for debugging or monitoring, the default setting of Reject is recommended. Make exceptions for specific port groups as required.

Values	Installation Default Value: Reject Baseline Suggested Value: Reject
---------------	--

Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Some workloads do legitimately use these network tactics and are affected negatively by the setting of Reject.
PowerCLI Command Assessment	<pre>Get-VDSwitch -Name \$VDS Get-VDSecurityPolicy Get-VDPortgroup -Name \$VDPG Get-VDSecurityPolicy</pre>
PowerCLI Command Remediation Example	<pre>Get-VDSwitch -Name \$VDS Get-VDSecurityPolicy Set-VDSecurityPolicy -AllowPromiscuous \$false Get-VDPortgroup -Name \$VDPG Get-VDSecurityPolicy Set-VDSecurityPolicy -AllowPromiscuous \$false</pre>
Setting Location in the vSphere Client	See the <i>vSphere Networking</i> documentation.

Reset Port Configuration When Virtual Machines Are Disconnected

The vCenter Server must reset port configuration when virtual machines are disconnected.

When a virtual machine is disconnected from the virtual switch port, it is desirable to reset the port configuration, so that another virtual machine that attaches has a port in a known state.

Values	Installation Default Value: Activated Baseline Suggested Value: Activated
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>(Get-VDPortgroup -Name \$VDPG).ExtensionData.Con- fig.Policy Select-Object -ExpandProperty PortCon- figResetAtDisconnect</pre>
PowerCLI Command Remediation Example	<pre>\$VDPGview = Get-VDPortgroup -Name \$VDPG Get-View \$ConfigSpec = New-Object VMware.Vim.DVPortgroupCon- figSpec \$ConfigSpec.DefaultPortConfig = New-Object VMware.Vim.VMwareDVSPortSetting \$ConfigSpec.Policy = New-Object VMware.Vim.VMwareD- VSPortgroupPolicy \$ConfigSpec.Policy.PortConfigResetAtDisconnect = \$true \$ConfigSpec.ConfigVersion = \$VDPGview.Config.Con- figVersion \$VDPGview.ReconfigureDVPortgroup_Task(\$ConfigSpec)</pre>
Setting Location in the vSphere Client	See the <i>vSphere Networking</i> documentation.

Deactivate Cisco Discovery Protocol or Link Layer Discovery Protocol

Deactivate participation in Cisco Discovery Protocol (CDP) or Link Layer Discovery Protocol (LLDP) on distributed switches unless used intentionally.

The vSphere Distributed Virtual Switch can engage in CDP or LLDP, potentially sharing sensitive unencrypted information, like IP addresses and system names, on the network. Thus, CDP and LLDP can aid adversaries in understanding or impersonating your environment. However, CDP and LLDP are also extremely helpful for legitimate use cases. Deactivate CDP and LLDP unless necessary for troubleshooting or configuration validation.

Values	Installation Default Value: Listen
---------------	------------------------------------

Recommended Action**Potential Functional Impact if Default Value Is Changed****PowerCLI Command Assessment****PowerCLI Command Remediation Example**

Baseline Suggested Value: None

Audit the installation default value.

None

```
(Get-VDSwitch -Name $VDS).ExtensionData.config.LinkDiscoveryProtocolConfig | Select-Object -ExpandProperty Operation
```

```
$VDview = Get-VDSwitch -Name $VDS | Get-View
$ConfigSpec = New-Object VMware.Vim.VMwareDVSSConfigSpec
$ConfigSpec.LinkDiscoveryProtocolConfig = New-Object VMware.Vim.LinkDiscoveryProtocolConfig
$ConfigSpec.LinkDiscoveryProtocolConfig.Protocol = 'cdp'
$ConfigSpec.LinkDiscoveryProtocolConfig.Operation = 'none'
$ConfigSpec.ConfigVersion = $VDview.Config.ConfigVersion
$VDview.ReconfigureDvs_Task($ConfigSpec)
```

Setting Location in the vSphere ClientSee the *vSphere Networking* documentation.**Ensure That Authorized Collectors Receive NetFlow Traffic**

The vCenter Server must ensure that NetFlow traffic is being sent to authorized collectors.

The vSphere Distributed Switch can export unencrypted NetFlow data, revealing details about virtual network and traffic patterns. Verify that NetFlow usage is authorized and configured correctly to prevent information leaks.

Values

Installation Default Value: Listen

Baseline Suggested Value: None

Recommended Action

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
(Get-VDSwitch -Name $VDS).ExtensionData.config.IpfixConfig.CollectorIpAddress | Select-Object -ExpandProperty CollectorIpAddress
(Get-VDPortgroup -Name $VDPG).ExtensionData.Config.DefaultPortConfig.IpfixEnabled | Select-Object -ExpandProperty Value
```

PowerCLI Command Remediation Example

```
$VDPGview = Get-VDPortgroup -Name $VDPG | Get-View
$ConfigSpec = New-Object VMware.Vim.DVPortgroupConfigSpec
$ConfigSpec.DefaultPortConfig = New-Object VMware.Vim.VMwareDVSPortSetting
$ConfigSpec.DefaultPortConfig.IpfixEnabled = New-Object VMware.Vim.BoolPolicy
$ConfigSpec.DefaultPortConfig.IpfixEnabled.Inherited = $false
$ConfigSpec.DefaultPortConfig.IpfixEnabled.Value = $false
```

```
$ConfigSpec.ConfigVersion = $VDPGview.Config.ConfigVersion
$VDPGview.ReconfigureDVPortgroup_Task($ConfigSpec)
```

Setting Location in the vSphere Client

See the *vSphere Networking* documentation.

Configure Virtual Machine Port Security

The vCenter Server must not override port group settings at the port level on distributed switches, except to block ports.

While port-level configuration overrides might be needed for unique virtual machine setups, be sure to monitor them to prevent unauthorized use. Unmonitored overrides could allow broader access if a less secure Distributed Switch configuration is exploited.

Values

Installation Default Value:

Block Ports Override: TRUE

All Other Overrides: FALSE

Baseline Suggested Value:

Block Ports Override: TRUE

All Other Overrides: FALSE

Recommended Action

Audit the installation default setting.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
(Get-VDPortgroup -Name $VDPG).ExtensionData.Config.Policy
```

PowerCLI Command Remediation Example

```
$VDPGview = Get-VDPortgroup -Name $VDPG | Get-View
$ConfigSpec = New-Object VMware.Vim.DVPortgroupConfigSpec
$ConfigSpec.DefaultPortConfig = New-Object VMware.Vim.VMwareDVSPortSetting
$ConfigSpec.Policy = New-Object VMware.Vim.VMwareDVSPortgroupPolicy
$ConfigSpec.Policy.UplinkTeamingOverrideAllowed = $false
$ConfigSpec.Policy.BlockOverrideAllowed = $true
$ConfigSpec.Policy.LivePortMovingAllowed = $false
$ConfigSpec.Policy.VlanOverrideAllowed = $false
$ConfigSpec.Policy.SecurityPolicyOverrideAllowed = $false
$ConfigSpec.Policy.VendorConfigOverrideAllowed = $false
$ConfigSpec.Policy.ShapingOverrideAllowed = $false
$ConfigSpec.Policy.IpfixOverrideAllowed = $false
$ConfigSpec.Policy.TrafficFilterOverrideAllowed = $false
$ConfigSpec.ConfigVersion = $VDPGview.Config.ConfigVersion
$VDPGview.ReconfigureDVPortgroup_Task($ConfigSpec)
```

Setting Location in the vSphere Client

See the *vSphere Networking* documentation.

Remove Port Mirroring

The vCenter Server must remove unauthorized port mirroring sessions on distributed switches.

The vSphere Distributed Switch can mirror traffic between ports, enabling traffic observation. To maintain security, any unauthorized port mirroring sessions on distributed switches must be removed.

Values	Installation Default Value: Not configured Baseline Suggested Value: Not configured
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>(Get-VDSwitch -Name \$VDS).ExtensionData.config.Vs-panSession</pre>
PowerCLI Command Remediation Example	N/A
Setting Location in the vSphere Client	See the <i>vSphere Networking</i> documentation.

Restrict Virtual Guest Tagging

The vCenter Server must restrict the use of Virtual Guest Tagging (VGT) on Distributed Switches.

Setting a port group to VLAN 4095 allows Virtual Guest Tagging (VGT), requiring the virtual machine to process VLAN tags. Activate VGT only for those virtual machines that are authorized and equipped to manage VLAN tags. Inappropriate use can cause service denial or unauthorized VLAN traffic interaction.

Values	Installation Default Value: Not configured Baseline Suggested Value: Not configured
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VDPortgroup -Name \$VDPG Where {\$_.ExtensionData.Config.Uplink -ne "True"} Select Name,VlanConfiguration</pre>
PowerCLI Command Remediation Example	<pre>Get-VDPortgroup \$VDPG Set-VDVlanConfiguration -VlanId "New_VLAN#"</pre>
Setting Location in the vSphere Client	See the <i>vSphere Networking</i> documentation.

Check for VMware Maintenance on the vCenter Server Version

Ensure that the vCenter Server version has not reached VMware End of General Support status.

Values	Installation Default Value: N/A Baseline Suggested Value: N/A
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	N/A
PowerCLI Command Remediation Example	N/A
Setting Location in the vCenter Server Management Interface	Update

Restrict Access to SSH

The vCenter Server SSH service must be deactivated.

vCenter Server Appliance is delivered as an appliance, and intended to be managed through the vCenter Server Management Interface, the vSphere Client, and the APIs. SSH is a troubleshooting and support tool to be activated only when necessary. vCenter Server High Availability uses SSH to coordinate the replication and failover between the nodes. Use of this feature requires SSH to remain activated.

Values	Installation Default Value: Deactivated Baseline Suggested Value: Deactivated
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	Note: You must first connect to the vCenter Server host by using the <code>Connect-CIServer</code> cmdlet. <pre>(Get-CisService -Name "com.vmware.appliance.access.ssh").get()</pre>
PowerCLI Command Remediation Example	<pre>(Get-CisService -Name "com.vmware.appliance.access.ssh").set(\$false)</pre>

Setting Location in the vCenter Server Management Interface **Access**

Verify the root User Password Expiration

The vCenter Server root account password expiration must be configured appropriately.

Modern best practices for passwords (NIST 800-63B Section 5.1.1.2, among other guidance) indicate that with adequate password entropy, security is not improved by arbitrarily requiring users to change their passwords at certain intervals. Many automated security tools and regulatory compliance frameworks do not reflect this guidance, and might override this recommendation.

Values	Installation Default Value: Yes Baseline Suggested Value: No
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Failure to reset the password before the expiration means that recovery procedures are necessary.
PowerCLI Command Assessment	N/A
PowerCLI Command Remediation Example	N/A

Setting Location in the vCenter Server Management Interface **Administration**

Configure File-Based Backup and Recovery

Configure File-Based Backup and Recovery so that you can recover your vCenter Server Appliance and its configuration using the vCenter Server installer. Backup and restore is an important part of protecting your environment.

Values	Installation Default Value: Not configured Baseline Suggested Value: Configured
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None

PowerCLI Command Assessment	N/A
PowerCLI Command Remediation Example	N/A
Setting Location in the vCenter Server Management Interface	Backup

Configure the Firewall to Only Allow Traffic from Authorized Networks

The vCenter Server Appliance must configure the firewall to only allow traffic from authorized networks.

Ensure that all incoming and outgoing network traffic is blocked unless explicitly allowed, reducing the attack surface and helping to prevent unauthorized access to the system. Outgoing (egress) traffic is not blocked, nor are related or established connections, so vCenter Server Appliance is still able to communicate with systems where it initiates the connection. Use perimeter firewalls to curtail those types of connections.

Values	Installation Default Value: Connections allowed from any IP address. Baseline Suggested Value: Connections allowed only from authorized infrastructure and administration workstations.
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Loss of connectivity. Ensure that you configure an allow rule for yourself prior to configuring a "deny all" rule.
PowerCLI Command Assessment	N/A
PowerCLI Command Remediation Example	N/A
Setting Location in the vCenter Server Management Interface	Firewall

Configure Remote Log Server

Configure a remote log server for the vCenter Server.

Remote logging to a central host enhances the security of vCenter Server by storing logs securely. Remote logging simplifies monitoring across hosts and supports aggregate analysis for detecting coordinated attacks. Centralized logging prevents tampering and serves as a reliable long-term audit record.

Values	Installation Default Value: Not configured Baseline Suggested Value: Site-specific log server
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	Note: You must first connect to the vCenter Server host by using the <code>Connect-CIServer</code> cmdlet. <pre>(Get-CisService -Name "com.vmware.appliance.logging.forwarding").get()</pre>
PowerCLI Command Remediation Example	N/A
Setting Location in the vSphere Client	N/A

Configure Time Synchronization

vCenter Server must have reliable time synchronization sources.

Cryptography, audit logging, cluster operations, and incident response and forensics heavily rely on synchronized time. Network Time Protocol (NTP) must have at least four sources. If you must choose between two sources and one source, one source is preferable.

Values	Installation Default Value: Undefined Baseline Suggested Value: Site-specific, or: 0.vmware.pool.ntp.org, 1.vmware.pool.ntp.org, 2.vmware.pool.ntp.org, 3.vmware.pool.ntp.org
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	Note: You must first connect to the vCenter Server host by using the <code>Connect-CIServer</code> cmdlet. <pre>(Get-CisService -Name "com.vmware.appliance.time-sync").get() (Get-CisService -Name "com.vmware.appliance.ntp").get()</pre>
PowerCLI Command Remediation Example	<pre>(Get-CisService -Name "com.vmware.appliance.time-sync").set("NTP") (Get-CisService -Name "com.vmware.appliance.ntp").set("0.vmware.pool.ntp.org,1.vmware.pool.ntp.org,2.vmware.pool.ntp.org,3.vmware.pool.ntp.org")</pre>
Setting Location in the vSphere Client	N/A

Install Software Updates

Ensure that vCenter Server has all software updates installed.

By keeping vCenter Server patches up to date, vulnerabilities can be mitigated. Attackers can exploit known vulnerabilities when attempting to gain unauthorized access or elevate privileges.

When applying updates, update vCenter Server first, if an update is available, then proceed with updating ESXi. This sequence ensures that the management layer is updated before updating the ESXi hosts.

Values	Installation Default Value: N/A Baseline Suggested Value: N/A
Recommended Action	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	N/A
PowerCLI Command Remediation Example	N/A
Setting Location in the vSphere Client	See the <i>Managing Host and Cluster Lifecycle</i> documentation.

Rotate the vpxuser Password

The vCenter Server must configure the vpxuser password to be rotated on an appropriate interval.

The `VirtualCenter.VimPasswordExpirationInDays` setting configures the rotation period. Ensure that the vCenter Server is properly rotating the password that it automatically sets on the ESXi hosts.

Values	Installation Default Value: 30 Baseline Suggested Value: 30
Recommended Action	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-AdvancedSetting -Entity \$VC -Name VirtualCenter.VimPasswordExpirationInDays</pre>
PowerCLI Command Remediation Example	<pre>Get-AdvancedSetting -Entity \$VC -Name VirtualCenter.VimPasswordExpirationInDays Set-AdvancedSetting -Value 30</pre>
Setting Location in the vSphere Client	Select vCenter Server > Configure > Advanced Settings

Virtual Machine Security Controls Reference

These security controls provide a baseline set of virtual machine security best practices. They are structured in a way that explains the benefits and tradeoffs of implementing the control. To change advanced system settings, you can use either the PowerCLI provided, or the vSphere Client (**Host > Configure > System > Advanced System Settings**).

Variable Used

The PowerCLI commands in this section use the following variable:

- `$VM = "virtual_machine_name"`

mks.enable3d

Deactivates 3D graphics features on virtual machines that do not need them to reduce potential attack vectors, enhancing overall system security.

Values	Installation Default Value: False Baseline Suggested Value: False
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Get-AdvancedSetting mks.enable3d</pre>
PowerCLI Command Remediation Example	<pre>Get-VM -Name \$VM Get-AdvancedSetting mks.enable3d Set-AdvancedSetting -Value FALSE</pre>

ethernet*.filter*.name

Limits access to virtual machines through the "dvFilter" network API.

The dvFilter interface is used by tools like NSX to filter and inspect network traffic. Other tools might use it, too. Ensure that those tools are authorized.

Values	Installation Default Value: Not present Baseline Suggested Value: Not present
Action Needed	Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

Legitimate network tools, including NSX, might require this functionality.

PowerCLI Command Assessment

```
Get-VM -Name $VM | Get-AdvancedSetting "ether-
net*.filter*.name*"
```

PowerCLI Command Remediation Example

```
Get-VM -Name $VM | Get-AdvancedSetting "ether-
net*.filter*.name*" | Remove-AdvancedSetting
```

Prevent Virtual Machines Booting from Unauthorized Sources

Virtual machines must prevent booting from unauthorized sources.

Unauthorized access to a virtual machine might occur when its primary boot volume is unavailable, and the EFI firmware seeks alternative boot sources, such as network boot. This can be mitigated through network controls, as well as with the advanced parameters `bios.bootDeviceClasses`, `bios.bootOrder`, and `bios.hddOrder`.

`bios.bootDeviceClasses` has the format "allow:XXXX" or "deny:XXXX," where XXXX is a comma-delimited list of boot classes. Boot classes are net (network PXE boot), usb (from attached USB devices); pcmcia (PCMCIA expansion cards, not used at present); cd (from attached virtual CD/DVD devices); hd (from attached virtual hard disks); fd (from attached virtual floppy devices); reserved (from unknown devices); efishell (into the EFI shell); and all, or any (same as all).

Use of allow or deny also implicitly states the opposite. For example:

- deny:all disallows all boot classes
- deny:net disallows network boot but allows all others
- allow:hd allows only hd boot and denies all others
- allow:hd,cd allows hd then cd device boot and denies all others

New virtual machines might require CD/DVD boot, and some dynamic environments, such as labs, might use network boot. Set those environments accordingly and document your rationale.

Values

Installation Default Value: allow:all

Baseline Suggested Value: allow:hd (once the guest operating system is installed)

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

The virtual machine can no longer boot from unspecified sources, which might negatively impact situations where PXE boot or recovery media is required. However, the parameter is easily changed at scale through PowerCLI. An alternate approach is to specify denied methods, such as "deny:net".

PowerCLI Command Assessment

```
Get-VM -Name $VM | Get-AdvancedSetting bios.bootDe-
viceClasses
```

PowerCLI Command Remediation Example

```
Get-VM -Name $VM | Get-AdvancedSetting bios.bootDe-
viceClasses | Set-AdvancedSetting -Value "allow:hd"
```

RemoteDisplay.maxConnections

Limits the number of console connections to a virtual machine.

Limiting virtual machine console sharing to one user prevents multiple observers, thus enhancing security. However, this might inadvertently create a potential pathway for service denial.

Values

Installation Default Value: -1

Action Needed**Potential Functional Impact if Default Value Is Changed****PowerCLI Command Assessment****PowerCLI Command Remediation Example**

Baseline Suggested Value: 1

Modify the installation default value.

Might create a denial-of-service condition where the console is unusable because the one user is connected, or a disconnected console session persists. Other products, such as VMware Cloud Director, might require that you set this option to a greater value.

```
Get-VM -Name $VM | Get-AdvancedSetting RemoteDisplay.maxConnections
```

```
Get-VM -Name $VM | Get-AdvancedSetting RemoteDisplay.maxConnections | Set-AdvancedSetting -Value 1
```

Limit PCI Device Passthrough Functionality

Virtual machines must limit PCI device passthrough functionality.

DirectPath I/O features enable virtual machines to directly access system hardware, impacting risk mitigation tools like vMotion, DRS, and High Availability. DirectPath I/O features also potentially give attackers privileged hardware access. Ensure that only necessary virtual machines have this privilege, compensating with guest operating system security controls.

Values

Installation Default Value: Not present

Baseline Suggested Value: Not present

Action Needed

Audit

Potential Functional Impact if Default Value Is Changed

Passthrough devices, such as GPUs, would be negatively affected if they are disconnected. Audit and document the business need for these virtual machines.

PowerCLI Command Assessment

```
Get-VM -Name $VM | Get-PassthroughDevice
```

PowerCLI Command Remediation Example

```
Get-VM -Name $VM | Get-PassthroughDevice | Remove-PassthroughDevice
```

Remove Unnecessary Virtual Machine Virtual Hardware Devices

Virtual machines must remove unnecessary virtual hardware.

To reduce potential attack surfaces, eliminate unnecessary virtual hardware from virtual machines. Rarely-used ports, temporary CD/DVD drives, and hardware introduced by migrations could be vulnerable. Removing these decreases the risk of software introduction or data exfiltration from a protected environment.

Values

Installation Default Value: Configured

Baseline Suggested Value: Not present

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

Removing the CD-ROM device might impact VMware Tools installation and maintenance. Removing XHCI controllers might impact console keyboard and mouse connectivity for some guest operating systems.

PowerCLI Command Assessment

```
$VMview = Get-VM -Name $VM | Get-View
```

```
$UnnecessaryHardware = "VirtualUSBController|VirtualUSBXHCIController|VirtualParallelPort|VirtualFloppy|VirtualSerialPort|VirtualHdAudioCard|VirtualAHCIController|VirtualEnsoniq1371|VirtualCdrom"
```

```
$VMview.Config.Hardware.Device | Where-Object
{ $_.GetType().Name -match $UnnecessaryHardware } |
Foreach-Object {
    $devname = $_.GetType().Name
    Write-Host "$VM`: [WARNING] VM has a $devname device. Please evaluate and consider removing." -ForegroundColor Yellow
}
```

PowerCLI Command Remediation Example

N/A

Setting Location in the vSphere Client**Virtual machine > Edit Settings > Virtual Hardware****tools.guestlib.enableHostInfo**

Prevents virtual machines from obtaining host information about the hypervisor.

Preventing virtual machines from obtaining host information about the hypervisor mitigates risk of advanced attacks, as it denies adversaries crucial details about the physical host.

Values

Installation Default Value: False

Baseline Suggested Value: False or undefined

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VM -Name $VM | Get-AdvancedSetting tool-
s.guestlib.enableHostInfo
```

PowerCLI Command Remediation Example

```
Get-VM -Name $VM | Get-AdvancedSetting tool-
s.guestlib.enableHostInfo | Remove-AdvancedSetting
```

Setting Location in the vSphere Client

N/A. This is a VMX file setting.

Set Encryption for Fault Tolerance

Virtual machines must require encryption for Fault Tolerance.

Requiring encryption for Fault Tolerance in virtual machines ensures secure data transmission. While the default 'opportunistic' encryption likely results in encryption due to ubiquitous AES-NI support in vSphere-compatible hardware, enforcing 'required' encryption guarantees no unencrypted operations.

Values

Installation Default Value: ftEncryptionOpportunistic

Baseline Suggested Value: ftEncryptionRequired

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
(Get-VM -Name $VM).ExtensionData.Config.FtEncryptionMode
```

PowerCLI Command Remediation Example

```
$VMview = Get-VM -Name $VM | Get-View
```

```
$ConfigSpec = New-Object VMware.Vim.VirtualMachineConfigSpec
$ConfigSpec.FtEncryptionMode = New-object
VMware.Vim.VirtualMachineConfigSpecEncryptedFtModes
$ConfigSpec.FtEncryptionMode = "ftEncryptionRequired"
$VMview.ReconfigVM_Task($ConfigSpec)
```

Setting Location in the vSphere Client

Virtual machine > Edit Settings > VM Options > Encryption

isolation.tools.copy.disable

Deactivates console copy operations on virtual machines.

Deactivating console copy operations on virtual machines prevents data copying between the virtual machine and the local client, regardless of whether the user is accessing through the Web Console, VMRC, or another method.

Values

Installation Default Value: True
Baseline Suggested Value: True or Undefined

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VM -Name $VM | Get-AdvancedSetting isolation.tools.copy.disable
```

PowerCLI Command Remediation Example

```
Get-VM -Name $VM | Get-AdvancedSetting isolation.tools.copy.disable | Remove-AdvancedSetting
```

isolation.tools.paste.disable

Deactivates console paste operations on virtual machines.

By disabling console paste operations on virtual machines, data transfer from the local client to the virtual machine is blocked, whether the user is using the Web Console, VMRC, or another console.

Values

Installation Default Value: True
Baseline Suggested Value: True or Undefined

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VM -Name $VM | Get-AdvancedSetting isolation.tools.paste.disable
```

PowerCLI Command Remediation Example

```
Get-VM -Name $VM | Get-AdvancedSetting isolation.tools.paste.disable | Remove-AdvancedSetting
```

isolation.tools.diskShrink.disable

Deactivates virtual disk shrinking on virtual machines.

Disabling virtual disk shrinking on virtual machines helps avoid disk unavailability issues. The ability to perform this operation is typically limited for non-administrative users in the guest environment.

Values

Installation Default Value: True
Baseline Suggested Value: True or Undefined

Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.tools.diskShrink.disable</pre>
PowerCLI Command Remediation Example	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.tools.diskShrink.disable Remove-AdvancedSetting</pre>
<u>isolation.tools.diskWiper.disable</u>	
Deactivates virtual disk wiping operations on virtual machines.	
Disabling virtual disk wiping on virtual machines helps avoid disk unavailability issues. The ability to perform this operation is typically limited for non-administrative users in the guest environment.	
Values	Installation Default Value: True Baseline Suggested Value: True or Undefined
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.tools.diskWiper.disable</pre>
PowerCLI Command Remediation Example	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.tools.diskWiper.disable Remove-AdvancedSetting</pre>
<u>isolation.device.connectable.disable</u>	
Prevents virtual machines from unauthorized removal, connection, and modification of devices.	
Preventing unauthorized device modifications in virtual machines blocks non-administrative users or processes from connecting, disconnecting, or adjusting device settings. This measure curbs unauthorized access and operation disruption, reducing denial of service risks, as well as some avenues for exfiltrating data.	
Values	Installation Default Value: True Baseline Suggested Value: True or Undefined
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.device.connectable.disable</pre>
PowerCLI Command Remediation Example	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.device.connectable.disable Remove-AdvancedSetting</pre>
<u>isolation.tools.dnd.disable</u>	
Deactivates drag and drop operations on virtual machine consoles.	
Deactivating drag and drop operations in the console of a virtual machine stops users from transferring data between the virtual machine and the local client, regardless of the console type, thus enhancing data security.	
Values	Installation Default Value: True

Action Needed	Baseline Suggested Value: True or Undefined Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.tools.dnd.disable</pre>
PowerCLI Command Remediation Example	<pre>Get-VM -Name \$VM Get-AdvancedSetting isolation.tools.dnd.disable Remove-AdvancedSetting</pre>

tools.setInfo.sizeLimit

Limits the informational messages passed from the virtual machine to the VMX file.

Limiting virtual machine informational messages to the VMX file prevents it from exceeding its default size of 1 MB. This option prevents potential denial of service situations that could arise if the datastore becomes full.

Values	Installation Default Value: 1048576 Baseline Suggested Value: 1048576 or Undefined
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Get-AdvancedSetting tools.setInfo.sizeLimit</pre>
PowerCLI Command Remediation Example	<pre>Get-VM -Name \$VM Get-AdvancedSetting tools.setInfo.sizeLimit Remove-AdvancedSetting</pre>

Activate Logging

Virtual machines must activate diagnostic logging.

Diagnostic logging for virtual machines helps with forensics and troubleshooting.

Values	Installation Default Value: True Baseline Suggested Value: True
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Where {\$_.ExtensionData.Config.Flags.EnableLogging -ne "True"}</pre>
PowerCLI Command Remediation Example	<pre>\$VMview = Get-VM -Name \$VM Get-View \$ConfigSpec = New-Object VMware.Vim.VirtualMachineConfigSpec \$ConfigSpec.Flags = New-Object VMware.Vim.VirtualMachineFlagInfo \$ConfigSpec.Flags.EnableLogging = \$true \$VMview.ReconfigVM_Task(\$ConfigSpec)</pre>

log.keepOld

Limits the number of retained virtual machine diagnostic logs.

By restricting the quantity of retained diagnostic logs, you avoid filling up the datastore without compromising diagnostic functionality.

Values	Installation Default Value: 10 Baseline Suggested Value: 10 or Undefined
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<code>Get-VM -Name \$VM Get-AdvancedSetting log.keepOld</code>
PowerCLI Command Remediation Example	<code>Get-VM -Name \$VM Get-AdvancedSetting log.keepOld Remove-AdvancedSetting</code>

log.rotateSize

Limits the size of virtual machine diagnostic logs.

Limiting the size of diagnostic logs on virtual machines prevents excessive space consumption, particularly on long-running virtual machines. The recommended minimum limit is 2 MB.

Values	Installation Default Value: 2048000 Baseline Suggested Value: 2048000 or Undefined
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<code>Get-VM -Name \$VM Get-AdvancedSetting log.rotateSize</code>
PowerCLI Command Remediation Example	<code>Get-VM -Name \$VM Get-AdvancedSetting log.rotateSize Remove-AdvancedSetting</code>

tools.guestlib.enableHostInfo

Deactivates the ability to send host information to guests.

By configuring a virtual machine to get detailed information about the physical host, an adversary could potentially use this information to inform further attacks on the host. As the default is the desired state, you can audit by verifying that this setting is either unset, or if it is set, that it is set to False.

Values	Installation Default Value: False Baseline Suggested Value: False or Undefined
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Inability to retrieve performance information about the host from inside the guest operating system could hamper troubleshooting.
PowerCLI Command Assessment	<code>Get-VM -Name \$VM Get-AdvancedSetting tools.guestlib.enableHostInfo</code>
PowerCLI Command Remediation Example	<code>Get-VM -Name \$VM Get-AdvancedSetting tools.guestlib.enableHostInfo Remove-AdvancedSetting</code>

tools.guest.desktop.autolock

Locks the virtual machine guest session when the remote console is disconnected.

Locking virtual machines when the last console connection closes can prevent potential unauthorized access from attackers exploiting logged-in console sessions.

Values	Installation Default Value: False Baseline Suggested Value: True
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>Get-VM -Name \$VM Get-AdvancedSetting tool- s.guest.desktop.autolock</pre>
PowerCLI Command Remediation Example	<pre>Get-VM -Name \$VM Get-AdvancedSetting tool- s.guest.desktop.autolock Remove-AdvancedSetting</pre>

Activate Encryption for vMotion

Virtual machines must require encryption for vMotion.

Requiring encryption for vMotion in virtual machines guarantees secure data transfer. The default 'opportunistic' encryption likely results in encryption due to widespread AES-NI support in vSphere-compatible hardware. However, enforcing 'required' encryption prevents any unencrypted operations.

Values	Installation Default Value: Opportunistic Baseline Suggested Value: Required
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None
PowerCLI Command Assessment	<pre>(Get-VM -Name \$VM).ExtensionData.Config.MigrateEn- ryption</pre>
PowerCLI Command Remediation Example	<pre>\$VMview = Get-VM -Name \$VM Get-View \$ConfigSpec = New-Object VMware.Vim.VirtualMa- chineConfigSpec \$ConfigSpec.MigrateEncryption = New-Object VMware.Vim.VirtualMachineConfigSpecEncryptedVMo- tionModes \$ConfigSpec.MigrateEncryption = "required" \$VMview.ReconfigVM_Task(\$ConfigSpec)</pre>

Guest Operating System Security Controls

These security controls provide a baseline set of guest operating system best practices. They are structured in a way that explains the benefits and tradeoffs of implementing the control. To make changes to these controls, use either the PowerCLI provided or the vSphere Client.

Variable Used

The PowerCLI commands in this section use the following variable:

- `$VM = "virtual_machine_name"`

VMware Tools Path

The default installation path for VMware Tools is `C:\Program Files\VMware\VMware Tools`.

Configure Guest Operating System Secure Boot

The guest operating system must enable Secure Boot.

Secure Boot, supported by all modern guest operating systems, uses public key cryptography to validate firmware, boot loader, drivers, and OS kernel. By preventing system booting with uncertain boot chain validity, Secure Boot effectively restricts malware.

Values	Installation Default Value: Site-specific Baseline Suggested Value: True
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Enabling Secure Boot after installing a guest operating system might involve more steps. Refer to your guest operating system documentation for instructions.
PowerCLI Command Assessment	<pre>(Get-VM -Name \$VM).ExtensionData.Config.BootOptions.EfiSecureBootEnabled</pre>
PowerCLI Command Remediation Example	<pre>\$VMObj = (Get-VM -Name \$VM) \$ConfigSpec = New-Object VMware.Vim.VirtualMachineConfigSpec \$bootOptions = New-Object VMware.Vim.VirtualMachineBootOptions \$bootOptions.EfiSecureBootEnabled = \$true \$ConfigSpec.BootOptions = \$bootOptions \$task = \$VMObj.ExtensionData.ReconfigVM_Task(\$ConfigSpec)</pre>
Setting Location in the vSphere Client	Virtual machine > Edit Settings > VM Options

Limit the Use of MSI Transforms

The guest operating must limit the use of MSI transforms when reconfiguring VMware Tools.

MSI transforms allow the installation database on Microsoft Windows guest operating systems to be changed. This can be helpful but also presents an opportunity to alter the security profile of the guest operating system from vSphere.

Values	Installation Default Value: False Baseline Suggested Value: False
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Administrators must use other methods to update and reconfigure VMware Tools when needed.
PowerCLI Command Assessment	<pre>VMwareToolboxCmd.exe config get autoupgrade allow-msi-transforms</pre>
PowerCLI Command Remediation Example	<pre>VMwareToolboxCmd.exe config set autoupgrade allow-msi-transforms false</pre>
Setting Location in the vSphere Client	N/A

Deactivate Appinfo

The guest operating system must deactivate Appinfo information gathering unless required.

Appinfo is a method to do application discovery through VMware Tools. If you are not using this tool, deactivate the module to reduce attack surface.

Values	Installation Default Value Value: False Baseline Suggested Value: True
Action Needed	Modify the installation default.
Potential Functional Impact if Default Value Is Changed	Products and services within the VMware ecosystem might require this functionality.
PowerCLI Command Assessment	<code>VMwareToolboxCmd.exe config get appinfo disabled</code>
PowerCLI Command Remediation Example	<code>VMwareToolboxCmd.exe config set appinfo disabled true</code>
Setting Location in thevSphere Client	N/A

Deactivate ContainerInfo

The guest operating system must deactivate ContainerInfo unless required.

The VMware Tools ContainerInfo plug-in for Linux gathers the list of running containers inside a Linux guest operating system.

Values	Installation Default Value: 21600 Baseline Suggested Value: 0
Action Needed	Modify the installation default.
Potential Functional Impact if Default Value Is Changed	Products and services within the VMware ecosystem might require this functionality.
PowerCLI Command Assessment	<code>VMwareToolboxCmd.exe config get containerinfo poll-interval</code>
PowerCLI Command Remediation Example	<code>VMwareToolboxCmd.exe config set containerinfo poll-interval 0</code>
Setting Location in thevSphere Client	N/A

Deactivate Guest Operations

Deactivate Guest Operations unless required.

Guest Operations are a set of functions that underpin most host-to-guest interaction. Deactivating them reduces attack surface but also drastically reduces functionality. Ensure that your environment does not require these functions. Do not deactivate guest operations on template virtual machines.

For a list of functions see the following documentation:

<https://vdc-download.vmware.com/vmwb-repository/dcr-public/fe08899f-1eec-4d8d-b3bc-a6664c168c2c/7fdf97a1-4c0d-4be0-9d43-2ceebbc174d9/doc/vim.vm.guest.GuestOperationsManager.html>

Values	Installation Default Value: False Baseline Suggested Value: True
Action Needed	Modify the installation default.
Potential Functional Impact if Default Value Is Changed	Products and services within the VMware ecosystem might require this functionality.

PowerCLI Command Assessment	VMwareToolboxCmd.exe config get guestoperations disabled
------------------------------------	--

PowerCLI Command Remediation Example	VMwareToolboxCmd.exe config set guestoperations disabled true
---	---

Setting Location in thevSphere Client	N/A
--	-----

Prevent Guest Operating System Recustomization

You must prevent the guest operating system on deployed and customized virtual machines from being recustomized.

The virtual machine deployment process offers many options for vSphere administrators to customize virtual machines by using scripts and running commands. These customization approaches can also be an avenue for an adversary to gain access to data inside a virtual machine, through cloning and recustomization. After deploying a virtual machine, prevent it from being customized again. You can always revert this change.

Values	Installation Default Value: True Baseline Suggested Value: False
Action Needed	Modify the installation default.
Potential Functional Impact if Default Value Is Changed	Once set, virtual machines are able to be customized when they are cloned. Do not make this change on template virtual machines. Making this change might impact disaster recovery processes that change IP addresses, through VMware Site Recovery Manager or VMware Cloud Disaster Recovery. For more information see the following documentation: https://techdocs.broadcom.com/us/en/vmware-cis/live-recovery/live-cyber-recovery/1-0/using-vmware-live-cyber-recovery/set-up-recovery-plans-in-vmware-live-cyber-recovery/configure-recovery-plans/configure-plan-mappings/ip-address-mapping.html

PowerCLI Command Assessment	VMwareToolboxCmd.exe config get deployPkg enable-customization
------------------------------------	--

PowerCLI Command Remediation Example	VMwareToolboxCmd.exe config set deployPkg enable-customization false
---	--

Setting Location in thevSphere Client	N/A
--	-----

Deactivate GuestStore Upgrade Operations

The guest operating must deactivate GuestStore Upgrade operations unless required.

The GuestStore feature provides a simple and flexible mechanism to distribute VMware-specific or custom content from a GuestStore repository to multiple guests simultaneously. If you are not using this feature, deactivate the plug-in to reduce the attack surface.

Values	Installation Default Value: Manual Baseline Suggested Value: Off
Action Needed	Modify the installation default.
Potential Functional Impact if Default Value Is Changed	Products and services within the VMware ecosystem might require this functionality.

PowerCLI Command Assessment	VMwareToolboxCmd.exe config get gueststoreupgrade policy
------------------------------------	--

PowerCLI Command Remediation Example

```
VMwareToolboxCmd.exe config set gueststoreupgrade
policy off
```

Setting Location in thevSphere Client

N/A

Deactivate Service Discovery

The guest operating system must deactivate Service Discovery unless required.

The VMware Tools Service Discovery plug-in connects to Aria Operations and provides additional data to that product about guest operating systems and workloads. If you are not using this feature, deactivate the plug-in to reduce the attack surface.

Values

Installation Default Value: False
Baseline Suggested Value: True

Action Needed

Modify the installation default.

Potential Functional Impact if Default Value Is Changed

Products and services within the VMware ecosystem might require this functionality.

PowerCLI Command Assessment

```
VMwareToolboxCmd.exe config get servicediscovery
disabled
```

PowerCLI Command Remediation Example

```
VMwareToolboxCmd.exe config set servicediscovery
disabled true
```

Setting Location in thevSphere Client

N/A

Activate VMware Tools Logging

The guest operating system must enable VMware Tools logging.

Ensure that VMware Tools logs information as appropriate. See <https://github.com/vmware/open-vm-tools/blob/master/open-vm-tools/tools.conf> for examples.

Values

Installation Default Value: True
Baseline Suggested Value: True

Action Needed

Audit the installation default.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
VMwareToolboxCmd.exe config get logging log
```

PowerCLI Command Remediation Example

```
VMwareToolboxCmd.exe config set logging log true
```

Setting Location in thevSphere Client

N/A

Send VMware Tools Logs to System Log Service

The guest operating system must send VMware Tools logs to the system log service.

By default, VMware Tools sends logs to a file on disk. Configure logs to be sent to syslog on Linux guests, and to the Windows Event Service on Microsoft Windows guests, for management and central archiving.

Values

Installation Default Value: file
Baseline Suggested Value: syslog

Action Needed

Modify the installation default.

Potential Functional Impact if Default Value Is Changed

Update the processes that rely on these files being in the default location.

PowerCLI Command Assessment

```
VMwareToolboxCmd.exe config get logging vmsvc.handler
VMwareToolboxCmd.exe config get logging toolbox-cmd.handler
VMwareToolboxCmd.exe config get logging vgaauthsvc.handler
VMwareToolboxCmd.exe config get logging vmtoolsd.handler
```

PowerCLI Command Remediation Example

```
VMwareToolboxCmd.exe config set logging vmsvc.handler syslog
VMwareToolboxCmd.exe config set logging toolbox-cmd.handler syslog
VMwareToolboxCmd.exe config set logging vgaauthsvc.handler syslog
VMwareToolboxCmd.exe config set logging vmtoolsd.handler syslog
```

Setting Location in thevSphere Client

N/A

Ensure That VMware Tools Version Is Up-To-Date

The guest operating system must ensure that VMware Tools are up-to-date.

VMware Tools are an important part of the VMware ecosystem. Using VMware Tools, you can perform guest operating system administration, such as:

- Graceful shutdown
- Lifecycle management
- Getting drivers for paravirtualized devices
- Customizing and deploying virtual machine templates

As with all software, you must manage and update VMware Tools as needed. Ensure that you are running a supported version for your guest operating system, whether it is delivered as part of the Linux distribution or installed by you for Microsoft Windows.

Values

Installation Default Value: N/A
Baseline Suggested Value: N/A

Action Needed

Audit the installation default value.

Potential Functional Impact if Default Value Is Changed

None

PowerCLI Command Assessment

```
Get-VM -Name $VM | Select-Object -Property Name,@{Name='ToolsVersion';Expression={$_.Guest.ToolsVersion}}
```

PowerCLI Command Remediation Example

Site-Specific. Multiple ways exist to update VMware Tools. Drivers for VMXNET3 and PVSCSI are also available through Windows Update, so ensure that you import them into tools such as WSUS.

Setting Location in thevSphere Client

Virtual machine > Virtual Machine Details > VMware Tools

Deactivate GlobalConf

The guest operating system must deactivate GlobalConf unless required.

The GlobalConf feature of VMware Tools provides the capability to push `tools.conf` file configurations to virtual machines.

Values	Installation Default Value: False Baseline Suggested Value: False
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Administrators need to use other methods to update and reconfigure VMware Tools when needed.
PowerCLI Command Assessment	<code>VMwareToolboxCmd.exe config get globalconf enabled</code>
PowerCLI Command Remediation Example	<code>VMwareToolboxCmd.exe config set globalconf enabled false</code>
Setting Location in the vSphere Client	N/A

Limit Automatic Renewal of VMware Tools Features

The guest operating system must limit the automatic removal of VMware Tools features.

VMware Tools automatic upgrade processes can add or remove features from the VMware Tools installation, which can be helpful but also presents an opportunity to alter the security profile of the guest operating system from vSphere.

Values	Installation Default Value: True Baseline Suggested Value: False
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Administrators need to use other methods to update and reconfigure VMware Tools when needed.
PowerCLI Command Assessment	<code>VMwareToolboxCmd.exe config get autoupgrade allow-remove-feature</code>
PowerCLI Command Remediation Example	<code>VMwareToolboxCmd.exe config set autoupgrade allow-remove-feature false</code>
Setting Location in the vSphere Client	N/A

Configure VMware Tools for Automatic Upgrades

The guest operating system must configure automatic VMware Tools upgrades as appropriate for the environment.

VMware Tools updates can be initiated by vSphere, which can be helpful for maintaining current VMware Tools versions. If you manage and update VMware Tools in other ways, deactivate this functionality. In general, leave automatic updates activated.

Values	Installation Default Value: True Baseline Suggested Value: True
Action Needed	Audit the installation default value.
Potential Functional Impact if Default Value Is Changed	Administrators need to use other methods to update and reconfigure VMware Tools when needed.

PowerCLI Command Assessment

```
VMwareToolboxCmd.exe config get autoupgrade al-
low-upgrade
```

PowerCLI Command Remediation Example

```
VMwareToolboxCmd.exe config set autoupgrade al-
low-upgrade true
```

Setting Location in the vSphere Client

N/A

Verify the Virtual Machine Hardware Version

The guest operating system must ensure that virtual machine hardware is version 19 or newer where supported.

Virtual machine hardware 19 is compatible with ESXi 7.0 Update 2 and later. Newer versions of virtual machine hardware enable new features and better performance. Consider upgrading to virtual machine hardware 20 if you are fully updated to vSphere 8.0 or later. As always, use caution when upgrading, and fully test the upgrade process before rolling out system-wide.

Consider all the locations where a virtual machine might run, or where you might need to restore the virtual machine. For example, users of the VMware Cloud Disaster Recovery service need to consider the vSphere levels of potential recovery SDDCs. While VMware Cloud runs atop vSphere, it might not have the same supported virtual hardware versions available.

Changes to the configuration of VMware-supplied virtual appliances are not supported and might cause service disruptions.

Note: When you upgrade the virtual machine hardware version, driver and other updates occur even though the impact is minimal.

Values

Installation Default Value: Site-specific
Baseline Suggested Value: vmx-19 or newer

Action Needed

Modify the installation default value.

Potential Functional Impact if Default Value Is Changed

Changing virtual machine hardware versions changes device versions inside the guest, which might have repercussions. Always test upgrading virtual hardware versions, and remember that snapshots capture the virtual machine version, too, so that you can revert versions if necessary.

Changes to the configuration of VMware-supplied virtual appliances are not supported and might cause service disruptions.

PowerCLI Command Assessment

```
(Get-VM -Name $VM | Get-View) | Select-Object -Prop-
erty Name,@{Name='HW Version';Expression={$_.Con-
fig.Version}}
```

PowerCLI Command Remediation Example

```
Set-VM -VM $VM -HardwareVersion vmx-19
```

Setting Location in the vSphere Client

When creating a virtual machine in the New Virtual Machine wizard, the **Select compatibility** selection sets the virtual machine hardware version.

vSAN Security Controls Reference

These security controls provide a baseline set of vSAN best practices. They are structured in a way that explains the benefits and tradeoffs of implementing the control. To make changes to these controls, see the *Administering VMware vSAN* documentation.

Protect Data at Rest

vSAN must protect data at rest.

vSAN Data-at-Rest encryption helps maintain the confidentiality of sensitive data while it resides on storage devices and reduce the risk of unauthorized access or exposure in the event of physical theft or loss.

You can change this configuration parameter while the cluster is operational. Enabling data-at-rest protections reformats disk groups (for vSAN OSA) and rewrites stored objects (for vSAN ESA), which might take considerable time, but it is done in the background. Workloads do not need to be powered off. vSAN ESA 8.0 Update 2 introduced the ability to enable data-at-rest protections on an existing vSAN ESA datastore. vSAN ESA 8.0 Update 3 introduces the ability to disable it again. Run the latest version of vSAN if using ESA.

Values	Installation Default Value: Deactivated Baseline Suggested Value: Enabled
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	All encryption comes at the cost of CPU cycles and potential storage latency. How much this impacts workloads depends on a variety of factors, such as the configuration of the underlying hardware and the type and frequency of storage I/O by the workload.

Protect Data While Traversing the Network

vSAN must protect data at rest, including storage-related network communications.

vSAN Data-in-Transit encryption helps ensure that sensitive data remains confidential while traversing the network, reducing the risk of unauthorized access or interception.

You can alter this configuration parameter while the cluster is operational.

Values	Installation Default Value: Deactivated Baseline Suggested Value: Enabled
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	All encryption comes at the cost of CPU cycles and potential storage latency. How much this impacts workloads depends on a variety of factors, such as the configuration of the underlying hardware and the type and frequency of storage I/O by the workload.

Restrict Access to NFS File Shares

NFS file shares on vSAN File Services must be configured to restrict access.

When configuring an NFS file share, select the "Customize net access" option and configure a restrictive set of permissions.

Values	Installation Default Value: No Access Baseline Suggested Value: Customize Net Access
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Loss of connectivity to clients.

Encrypt SMB Authentication

SMB file shares on vSAN File Services must accept only encrypted SMB authentication communications.

When configuring an SMB file share, activate the Protocol Encryption option.

Values	Installation Default Value: Deactivated Baseline Suggested Value: Activated
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	None.

Enable Bidirectional/Mutual CHAP Authentication

vSAN iSCSI target must enable bidirectional/mutual CHAP authentication.

Mutual CHAP provides an additional layer of protection by requiring both the initiator (client) and the target (server) to verify their identities to each other, thereby ensuring data transmitted between the two is not intercepted or altered by unauthorized entities.

Values	Installation Default Value: Deactivated Baseline Suggested Value: Activated
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Might be more difficult to configure clients.

Reserve Space to Complete Internal Maintenance Operations

vSAN must reserve space to complete internal maintenance operations.

vSAN Operations Reserve capacity setting helps ensure that vSAN always has sufficient free space to maintain the availability and reliability of the vSAN datastore and prevent potential data loss or service disruptions due to insufficient capacity during operations such as policy changes.

You can change this configuration parameter while the cluster is operational.

Values	Installation Default Value: Deactivated Baseline Suggested Value: Activated
Action Needed	Modify the installation default value.
Potential Functional Impact if Default Value Is Changed	Activating this option reduces the usable capacity of the vSAN datastore.

About the National Institute of Standards and Technology

The National Institute of Standards and Technology (NIST) is a non-regulatory government agency that develops technology, metrics, standards, and guidelines. Compliance with NIST standards and guidelines has become a top priority in many industries today.

The National Institute of Standards and Technology (NIST) was founded in 1901 and is now part of the U.S. Department of Commerce. NIST is one of the nation's oldest physical science laboratories. Today, NIST measurements support the smallest of technologies to the largest and most complex of human-made creations, from nanoscale devices, up to earthquake-resistant skyscrapers and global communication networks.

The Federal Information Security Management Act (FISMA) is a United States federal law passed in 2002 that made it a requirement for federal agencies to develop, document, and implement an information security and protection program.

NIST plays an important role in the FISMA implementation by producing key security standards and guidelines (for example, FIPS 199, FIPS 200, and SP 800 series).

Government and private organizations use NIST 800-53 to secure information systems. Cybersecurity and privacy controls are essential to protect organizational operations (including mission, functions, image, and reputation), organizational assets, and individuals from a diverse set of threats. Some of these threats include hostile cyber-attacks, natural disasters, structural failures, and human errors. VMware has enlisted a third-party audit partner to evaluate VMware products and solutions against the NIST 800-53 catalog of controls. For more information, visit the NIST webpage at <https://www.nist.gov/cyberframework>.

About DISA STIGs

The Defense Information Systems Agency (DISA) develops and publishes Security Technical Implementation Guides, or "STIGs." DISA STIGs provide technical guidance for hardening systems and reducing threats.

The Defense Information Systems Agency (DISA) is the U.S. Department of Defense (DoD) combat support agency responsible for maintaining the security posture of the DOD Information Network (DODIN). One of the ways DISA accomplishes this task is by developing, disseminating, and mandating the implementation of Security Technical Implementation Guides, or STIGs. In brief, STIGs are portable, standards-based guides for hardening systems. STIGs are mandatory for U.S. DoD IT systems and, as such, provide a vetted, secure baseline for non-DoD entities to measure their security posture.

Vendors such as VMware submit suggested security hardening guidance to DISA for evaluation, based on DISA protocols and feedback. Once that process is complete, the official STIG is published on the DISA organization's web site at <https://public.cyber.mil/stigs/>. VMware provides security baselines and hardening guidance for vSphere as part of the *vSphere Security Configuration Guide*. See <https://core.vmware.com/security>.

About NERC CIP

The North American Electric Reliability Corporation (NERC) is a nonprofit international regulatory authority tasked with ensuring the reliability and security of North America's bulk power system, including the United States, Canada, and part of Baja California in Mexico.

NERC develops, enforces, and monitors reliability standards, and educates, trains, and certifies industry personnel. The Critical Infrastructure Protection (CIP) standards, a key component of NERC's regulations, are designed to secure essential cyber assets of the bulk electric system. These standards mandate specific security management controls, system security management, information protection, and vulnerability assessments for utilities and other entities within the bulk power system.

For more information, visit the NERC CIP webpage at <https://core.vmware.com/nerc-cip>.

About VMware Security Development Lifecycle

The VMware Security Development Lifecycle (SDL) program identifies and mitigates security risk during the development phase of VMware software products. VMware also operates the VMware Security Response Center (VSRC) to conduct the analysis and remediation of software security issues in VMware products.

The SDL is the software development methodology that the VMware Security Engineering, Communication, and Response (vSECR) group, and VMware product development groups, use to help identify and mitigate security issues.

The VMware Security Response Center (VSRC) works with customers and the security research community to achieve the goals of addressing security issues and providing customers with actionable security information in a timely manner. For more information about the VSRC, see the webpage at <https://www.vmware.com/security/vsrc.html>.

Audit Logging in vSphere

Audit logging of network traffic, compliance alerts, firewall activity, operating system changes, and provisioning activities is considered a best practice for maintaining the security of any IT environment. In addition, logging is a specific requirement of many regulations and standards.

One of the first steps to take for ensuring that you are aware of changes to your infrastructure is to audit your environment. By default, vSphere includes tools that enable you to view and track changes. For example, you can use the Tasks and Events tab in the vSphere Client on any object in your vSphere hierarchy to see what changes have occurred. You can also use the PowerCLI to retrieve events and tasks. In addition, VMware Aria Operations for Logs offers audit logging to support collection and retention of important system events. Finally, many third-party tools are available that provide vCenter Server auditing.

Log files can also help determine who or what is accessing a host, a virtual machine, and so on. For more information, see [ESXi Log File Locations](#).

Single Sign-On Audit Events

Single Sign-On (SSO) audit events are records of user or system actions for accessing the SSO services.

vCenter Server 6.7 Update 2 and later improves VMware vCenter Single Sign-On auditing by adding events for the following operations:

- User management
- Login
- Group creation
- Identity source
- Policy updates

The supported identity sources are vsphere.local, Integrated Windows Authentication (IWA), and Active Directory over LDAP.

When a user logs in to vCenter Server through Single Sign-On, or makes changes that affect SSO, the following audit events are written to the SSO audit log file:

- **Login and Logout Attempts:** Events for all the successful and failed login and logout operations.
- **Privilege Change:** Event for change in a user role or permissions.
- **Account Change:** Event for change in the user account information, for example, user name, password, or any additional account information.
- **Security Change:** Event for change in a security configuration, parameter, or policy.
- **Account Enabled or Disabled:** Event for when an account is activated or deactivated.
- **Identity Source:** Event for adding, deleting, or editing an identity source.

In the vSphere Client, event data is displayed in the **Monitor** tab. See the *vSphere Monitoring and Performance* documentation.

SSO audit event data includes the following details:

- Timestamp of when the event occurred.
- User who performed the action.
- Description of the event.
- Severity of the event.
- IP address of client used to connect to vCenter Server, if available.

SSO Audit Event Log Overview

The vSphere Single-Sign On process writes audit events to the `audit_events.log` file in the `/var/log/audit/sso-events/` directory.

CAUTION: Never manually edit the `audit_events.log` file, as doing so might cause the audit logging to fail.

Keep the following in mind when working with the `audit_events.log` file:

- The log file is archived once it reaches 50 MB.
- A maximum of 10 archive files is kept. If the limit is reached, the oldest file is purged when a new archive is created.
- The archive files are named `audit_events-<index>.log.gz`, where the index is a numeral from 1 to 10. The first archive created is index 1, and is increased with each subsequent archive.
- The oldest events are in archive index 1. The highest indexed file is the latest archive.

Understanding Security and Compliance Next Steps

Conducting a security assessment is the first step in understanding any vulnerabilities in your infrastructure. A security assessment is part of a security audit, which looks at both systems and practices, including security compliance.

A security assessment generally refers to scanning your organization's physical infrastructure (firewalls, networks, hardware, and so on) to identify vulnerabilities and flaws. A security assessment is not the same as a security audit. A security audit includes not only a review of physical infrastructure but other areas such as policy and standard operating procedures, including security compliance. After you have the audit, you can decide on the steps to remedy the problems within the system.

You might ask these general questions when preparing to conduct a security audit:

1. Is our organization mandated to adhere to a compliance regulation? If so which one(s)?
2. What is our audit interval?
3. What is our internal self-assessment interval?
4. Do we have access to previous audit results and have we viewed them?
5. Do we use a third-party audit firm to help us prepare for an audit? If so, what is their level of comfort with virtualization?
6. Do we run vulnerability scans against the systems and applications? When and how often?
7. What are our internal cybersecurity policies?
8. Is your audit logging configured according to your needs? See [Audit Logging in vSphere](#).

In the absence of specific guidance or direction on where to begin, you can jumpstart securing your vSphere environment by:

- Keeping your environment up-to-date with the latest software and firmware patches
- Maintaining good password management and hygiene for all accounts
- Reviewing vendor-approved security recommendations
- Referring to the VMware Security Configuration Guides (see [vSphere Security Controls Reference](#))
- Using readily available and proven guidance from policy frameworks such as NIST, ISO, and so forth
- Following guidance from regulatory compliance frameworks such as PCI, DISA, and FedRAMP

vCenter Server and FIPS

In vSphere 7.0 Update 2 and later, you can enable FIPS-validated cryptography on the vCenter Server Appliance.

FIPS 140-2 is a U.S. and Canadian government standard that specifies security requirements for cryptographic modules. vSphere uses FIPS-validated cryptographic modules to match those specified by the FIPS 140-2 standard. The goal of vSphere FIPS support is to ease the compliance and security activities in various regulated environments.

In vSphere 6.7 and later, ESXi and vCenter Server use FIPS-validated cryptography to protect management interfaces and the VMware Certificate Authority (VMCA).

vSphere 7.0 Update 2 and later includes additional FIPS-validated cryptography to the vCenter Server Appliance.

Note: vSphere favors compatibility over FIPS, so some components have considerations to be aware of. See [Considerations When Using FIPS](#).

FIPS Modules Used in ESXi

A cryptographic module is a set of hardware, software, or firmware that implements security functions. ESXi uses several FIPS 140-2 validated cryptographic modules.

The following table shows the set of FIPS 140-2 validated cryptographic modules in use by ESXi.

Table 354: FIPS Modules

Cryptographic Module	Version	Algorithms (CAVP)	Certificate Number
VMkernel Cryptographic Module	2.0	AES-CBC, AES-CBC-CS3, AES-CTR, AES-ECB, AES-GCM, AES-XTS Testing Revision 2.0, Counter DRBG, HMAC-SHA-1, HMAC-SHA2-256, HMAC-SHA2-512, SHA-1, SHA2-256, SHA2-512 (A2792)	Certificate in progress
OpenSSL FIPS Object Module	3.0	AES, CKG, CVL, DRBG, DSA, ECDSA, HMAC, KAS-RSA-SSC, KAS-SSC, KBKDF, KDA, KMAC, KTS, KTS-RSA, PBKDF, RSA, SHA-3, SHS, Triple-DES (A1938)	Certificate #4282
VMware OpenSSL FIPS Object Module	2.0.20-vmw	AES, CKG, /drbg, DSA, ECDSA, HMAC, KAS-SSC, RSA, SHS, Triple-DES (C470)	Certificate #3857
VMware's ESXboot Cryptographic Module	1.0	HMAC-SHA2-224, RSA SigVer (FIPS186-4), SHA2-224, SHA2-256, SHA2-384, SHA2-512	Certificate #4442
VMware's Boring Crypto Module	6.0	AES, CVL, DRBG, ECDSA, HMAC, KAS, KAS-SSC, KTS, RSA, SHS, Triple-DES (A4970)	Certificate #4694

Activate and Deactivate FIPS on the vCenter Server Appliance

You can activate or deactivate FIPS-validated cryptography on the vCenter Server Appliance by using HTTP requests. FIPS-validated cryptography is deactivated by default.

You can use various ways to execute HTTP requests. This task shows how to use the Developer Center in the vSphere Client to activate and deactivate the FIPS-validated cryptography on the vCenter Server Appliance. See *VMware*

vCenter Server Management Programming Guide for more information about using APIs to work with the vCenter Server Appliance.

1. Log in to the vCenter Server system with the vSphere Client.
2. From the Menu, select **Developer Center**.
3. Click **API Explorer**.
4. From the **Select API** drop-down menu, select **appliance**.
5. Scroll down through the categories and expand **system/security/global_fips**.
6. Expand **GET** and click **Execute** under **Try it out**.
You can view the current setting under **Response**.
7. Change the setting.
 - a) To activate FIPS, expand **PUT**, enter the following in the `request_body`, and click **Execute**.


```
{
  "enabled":true
}
```
 - b) To deactivate FIPS, expand **PUT**, enter the following in the `request_body`, and click **Execute**.


```
{
  "enabled":false
}
```

The vCenter Server Appliance reboots after you activate or deactivate the FIPS-validated cryptography.

Considerations When Using FIPS

When activating FIPS on vCenter Server Appliance, some components present functional constraints currently.

You should see no differences after activating FIPS on vCenter Server, however there are some considerations to be aware of.

Table 355: FIPS Considerations

Product or Component	Consideration	Workaround
vSphere Single Sign-On	When you activate FIPS, vCenter Server supports only cryptographic modules for federated authentication. As a result, RSA SecureID and some CAC cards no longer function.	Use federated authentication. See the <i>vSphere Authentication</i> documentation for details.
Non-VMware and partner vSphere Client UI plug-ins	These plug-ins might not work with FIPS enabled.	Upgrade plug-ins to use conformant encryption libraries. See the topic titled "Preparing Local Plug-ins for FIPS Compliance" in the <i>vSphere Client SDK</i> documentation.
Certificates	Certificates with key sizes greater than 3072 bits have not been tested.	Generate certificates with keys using 2048 or 3072 bit sizes.

vSphere Resource Management

vSphere Resource Management describes resource management for VMware® ESXi and vCenter® Server environments.

This documentation focuses on the following topics.

- Resource allocation and resource management concepts
- Virtual machine attributes and admission control
- Resource pools and how to manage them
- Clusters, vSphere® Distributed Resource Scheduler (DRS), vSphere Distributed Power Management (DPM), and how to work with them
- Datastore clusters, Storage DRS, Storage I/O Control, and how to work with them
- Advanced resource management options
- Performance considerations

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

This information is for system administrators who want to understand how the system manages resources and how they can customize the default behavior. It's also essential for anyone who wants to understand and use resource pools, clusters, DRS, datastore clusters, Storage DRS, Storage I/O Control, or vSphere DPM.

This documentation assumes you have a working knowledge of VMware ESXi and of vCenter Server.

Note: In this document, "Memory" can refer to physical RAM or Persistent Memory.

Configure Your Resource Allocation Settings in vSphere

When available resource capacity does not meet the demands of the resource consumers (and virtualization overhead), administrators might need to customize the amount of resources that are allocated to virtual machines or resource pools in vSphere.

Use the resource allocation settings (shares, reservation, and limit) to determine the amount of CPU, memory, and storage resources provided for a virtual machine. In particular, administrators have several options for allocating resources.

- Reserve the physical resources of the host or cluster.
- Set an upper bound on the resources that can be allocated to a virtual machine.
- Guarantee that a particular virtual machine is always allocated a higher percentage of the physical resources than other virtual machines.

Note: In this chapter, "Memory" refers to physical RAM.

Resource Allocation Shares

Shares specify the relative importance of a virtual machine (or resource pool). If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when these two virtual machines are competing for resources.

Shares are typically specified as **High**, **Normal**, or **Low** and these values specify share values with a 4:2:1 ratio, respectively. You can also select **Custom** to assign a specific number of shares (which expresses a proportional weight) to each virtual machine.

Specifying shares makes sense only with regard to sibling virtual machines or resource pools, that is, virtual machines or resource pools with the same parent in the resource pool hierarchy. Siblings share resources according to their relative share values, bounded by the reservation and limit. When you assign shares to a virtual machine, you always specify the priority for that virtual machine relative to other powered-on virtual machines.

The following table shows the default CPU and memory share values for a virtual machine. For resource pools, the default CPU and memory share values are the same, but must be multiplied as if the resource pool were a virtual machine with four virtual CPUs and 16 GB of memory.

Table 356: Share Values

Setting	CPU share values	Memory share values
High	2000 shares per virtual CPU	20 shares per megabyte of configured virtual machine memory.
Normal	1000 shares per virtual CPU	10 shares per megabyte of configured virtual machine memory.
Low	500 shares per virtual CPU	5 shares per megabyte of configured virtual machine memory.

For example, an SMP virtual machine with two virtual CPUs and 1GB RAM with CPU and memory shares set to **Normal** has $2 \times 1000 = 2000$ shares of CPU and $10 \times 1024 = 10240$ shares of memory.

Virtual machines with more than one virtual CPU are called SMP (symmetric multiprocessing) virtual machines.

The relative priority represented by each share changes when a new virtual machine is powered on. This affects all virtual machines in the same resource pool. All of the virtual machines have the same number of virtual CPUs. Consider the following examples.

- Two CPU-bound virtual machines run on a host with 8GHz of aggregate CPU capacity. Their CPU shares are set to **Normal** and get 4GHz each.
- A third CPU-bound virtual machine is powered on. Its CPU shares value is set to **High**, which means it should have twice as many shares as the machines set to **Normal**. The new virtual machine receives 4GHz and the two other machines get only 2GHz each. The same result occurs if the user specifies a custom share value of 2000 for the third virtual machine.

Resource Allocation Reservation

A reservation specifies the guaranteed minimum allocation for a virtual machine.

vCenter Server or ESXi allows you to power on a virtual machine only if there are enough unreserved resources to satisfy the reservation of the virtual machine. The server guarantees that amount even when the physical server is heavily loaded. The reservation is expressed in concrete units (megahertz or megabytes).

For example, assume you have 2GHz available and specify a reservation of 1GHz for VM1 and 1GHz for VM2. Now each virtual machine is guaranteed to get 1GHz if it needs it. However, if VM1 is using only 500MHz, VM2 can use 1.5GHz.

Reservation defaults to 0. You can specify a reservation if you need to guarantee that the minimum required amounts of CPU or memory are always available for the virtual machine.

Resource Allocation Limit

Limit specifies an upper bound for CPU, memory, or storage I/O resources that can be allocated to a virtual machine.

A server can allocate more than the reservation to a virtual machine, but never allocates more than the limit, even if there are unused resources on the system. The limit is expressed in concrete units (megahertz, megabytes, or I/O operations per second).

CPU, memory, and storage I/O resource limits default to unlimited. When the memory limit is unlimited, the amount of memory configured for the virtual machine when it was created becomes its effective limit.

In most cases, it is not necessary to specify a limit. There are benefits and drawbacks:

- **Benefits** — Assigning a limit is useful if you start with a small number of virtual machines and want to manage user expectations. Performance deteriorates as you add more virtual machines. You can simulate having fewer resources available by specifying a limit.
- **Drawbacks** — You might waste idle resources if you specify a limit. The system does not allow virtual machines to use more resources than the limit, even when the system is underutilized and idle resources are available. Specify the limit only if you have good reasons for doing so.

Resource Allocation Settings Suggestions

Select resource allocation settings (reservation, limit and shares) that are appropriate for your ESXi environment.

The following guidelines can help you achieve better performance for your virtual machines.

- Use **Reservation** to specify the minimum acceptable amount of CPU or memory, not the amount you want to have available. The amount of concrete resources represented by a reservation does not change when you change the environment, such as by adding or removing virtual machines. The host assigns additional resources as available based on the limit for your virtual machine, the number of shares and estimated demand.
- When specifying the reservations for virtual machines, do not commit all resources (plan to leave at least 10% unreserved). As you move closer to fully reserving all capacity in the system, it becomes increasingly difficult to make changes to reservations and to the resource pool hierarchy without violating admission control. In a DRS-enabled cluster, reservations that fully commit the capacity of the cluster or of individual hosts in the cluster can prevent DRS from migrating virtual machines between hosts.
- If you expect frequent changes to the total available resources, use **Shares** to allocate resources fairly across virtual machines. If you use **Shares**, and you upgrade the host, for example, each virtual machine stays at the same priority (keeps the same number of shares) even though each share represents a larger amount of memory, CPU, or storage I/O resources.

Edit Settings in vSphere

Use the Edit Settings dialog box to change allocations for memory and CPU resources.

1. Browse to the virtual machine in the vSphere Client.
2. Right-click and select **Edit Settings**.
3. Edit the CPU Resources.

Option	Description
Shares	CPU shares for this resource pool with respect to the parent's total. Sibling resource pools share resources according to their relative share values bounded by the reservation and limit. Select Low , Normal , or High , which specify share values respectively in a 1:2:4 ratio. Select Custom to give each virtual machine a specific number of shares, which expresses a proportional weight.
Reservation	Guaranteed CPU allocation for this resource pool.
Limit	Upper limit for this resource pool's CPU allocation. Select Unlimited to specify no upper limit.

4. Edit the Memory Resources.

Option	Description
Shares	Memory shares for this resource pool with respect to the parent's total. Sibling resource pools share resources according to their relative share values bounded by the reservation and limit. Select Low , Normal , or High , which specify share values respectively in a 1:2:4 ratio. Select Custom to give each virtual machine a specific number of shares, which expresses a proportional weight.
Reservation	Guaranteed memory allocation for this resource pool.
Limit	Upper limit for this resource pool's memory allocation. Select Unlimited to specify no upper limit.


5. Click **OK**.

Changing Resource Allocation Settings—Example

The following example illustrates how you can change resource allocation settings to improve virtual machine performance.

Assume that on an ESXi host, you have created two new virtual machines—one each for your QA (VM-QA) and Marketing (VM-Marketing) departments.

Figure 125: Single Host with Two Virtual Machines



In the following example, assume that VM-QA is memory intensive and accordingly you want to change the resource allocation settings for the two virtual machines to:

- Specify that, when system memory is overcommitted, VM-QA can use twice as much CPU and memory resources as the Marketing virtual machine. Set the CPU shares and memory shares for VM-QA to **High** and for VM-Marketing set them to **Normal**.

- Ensure that the Marketing virtual machine has a certain amount of guaranteed CPU resources. You can do so using a reservation setting.
1. Browse to the virtual machines in the vSphere Client.
 2. Right-click **VM-QA**, the virtual machine for which you want to change shares, and select **Edit Settings**.
 3. Under **Virtual Hardware**, expand CPU and select **High** from the **Shares** drop-down menu.
 4. Under **Virtual Hardware**, expand Memory and select **High** from the **Shares** drop-down menu.
 5. Click **OK**.
 6. Right-click the marketing virtual machine (**VM-Marketing**) and select **Edit Settings**.
 7. Under **Virtual Hardware**, expand CPU and change the **Reservation** value to the desired number.
 8. Click **OK**.

Admission Control in vSphere

When you power on a virtual machine, the system checks the amount of CPU and memory resources that have not yet been reserved. Based on the available unreserved resources, the system determines whether it can guarantee the reservation for which the virtual machine is configured (if any). This process is called admission control.

If enough unreserved CPU and memory are available, or if there is no reservation, the virtual machine is powered on. Otherwise, an *Insufficient Resources* warning appears.

Note: In addition to the user-specified memory reservation, for each virtual machine there is also an amount of overhead memory. This extra memory commitment is included in the admission control calculation.

When the vSphere DPM feature is enabled, hosts might be placed in standby mode (that is, powered off) to reduce power consumption. The unreserved resources provided by these hosts are considered available for admission control. If a virtual machine cannot be powered on without these resources, a recommendation to power on sufficient standby hosts is made. For more information, see [Managing Power Resources](#).

CPU Virtualization in vSphere

CPU virtualization emphasizes performance and runs directly on the processor whenever possible. The underlying physical resources are used whenever possible and the virtualization layer runs instructions only as needed to make virtual machines operate as if they were running directly on a physical machine.

CPU virtualization is not the same thing as emulation. ESXi does not use emulation to run virtual CPUs. With emulation, all operations are run in software by an emulator. A software emulator allows programs to run on a computer system other than the one for which they were originally written. The emulator does this by emulating, or reproducing, the original computer's behavior by accepting the same data or inputs and achieving the same results. Emulation provides portability and runs software designed for one platform across several platforms.

When CPU resources are overcommitted, the ESXi host time-slices the physical processors across all virtual machines so each virtual machine runs as if it has its specified number of virtual processors. When an ESXi host runs multiple virtual machines, it allocates to each virtual machine a share of the physical resources. With the default resource allocation settings, all virtual machines associated with the same host receive an equal share of CPU per virtual CPU. This means that a single-processor virtual machines is assigned only half of the resources of a dual-processor virtual machine.

Software-Based CPU Virtualization

With software-based CPU virtualization, the guest application code runs directly on the processor, while the guest privileged code is translated and the translated code runs on the processor.

The translated code is slightly larger and usually runs more slowly than the native version. As a result, guest applications, which have a small privileged code component, run with speeds very close to native. Applications with a significant

privileged code component, such as system calls, traps, or page table updates can run slower in the virtualized environment.

Hardware-Assisted CPU Virtualization

Certain processors provide hardware assistance for CPU virtualization.

When using this assistance, the guest can use a separate mode of execution called guest mode. The guest code, whether application code or privileged code, runs in the guest mode. On certain events, the processor exits out of guest mode and enters root mode. The hypervisor executes in the root mode, determines the reason for the exit, takes any required actions, and restarts the guest in guest mode.

When you use hardware assistance for virtualization, there is no need to translate the code. As a result, system calls or trap-intensive workloads run very close to native speed. Some workloads, such as those involving updates to page tables, lead to a large number of exits from guest mode to root mode. Depending on the number of such exits and total time spent in exits, hardware-assisted CPU virtualization can speed up execution significantly.

Virtualization and Processor-Specific Behavior

Although VMware software virtualizes the CPU, the virtual machine detects the specific model of the processor on which it is running.

Processor models might differ in the CPU features they offer, and applications running in the virtual machine can make use of these features. Therefore, it is not possible to use vMotion® to migrate virtual machines between systems running on processors with different feature sets. You can avoid this restriction, in some cases, by using Enhanced vMotion Compatibility (EVC) with processors that support this feature. See the *vCenter Server and Host Management* documentation for more information.

Performance Implications of CPU Virtualization

CPU virtualization adds varying amounts of overhead depending on the workload and the type of virtualization used.

An application is CPU-bound if it spends most of its time executing instructions rather than waiting for external events such as user interaction, device input, or data retrieval. For such applications, the CPU virtualization overhead includes the additional instructions that must be executed. This overhead takes CPU processing time that the application itself can use. CPU virtualization overhead usually translates into a reduction in overall performance.

For applications that are not CPU-bound, CPU virtualization likely translates into an increase in CPU use. If spare CPU capacity is available to absorb the overhead, it can still deliver comparable performance in terms of overall throughput.

ESXi supports up to 128 virtual processors (CPUs) for each virtual machine.

Note: Deploy single-threaded applications on uniprocessor virtual machines, instead of on SMP virtual machines that have multiple CPUs, for the best performance and resource use.

Single-threaded applications can take advantage only of a single CPU. Deploying such applications in dual-processor virtual machines does not speed up the application. Instead, it causes the second virtual CPU to use physical resources that other virtual machines could otherwise use.

Administering CPU Resources in vSphere

You can configure virtual machines with one or more virtual processors, each with its own set of registers and control structures.

When a virtual machine is scheduled, its virtual processors are scheduled to run on physical processors. The VMkernel Resource Manager schedules the virtual CPUs on physical CPUs, thereby managing the virtual machine's access to physical CPU resources.

Note: In this chapter, "Memory" can refer to physical RAM or Persistent Memory.

View Processor Information

You can access information about current CPU configuration in the vSphere Client.

1. Browse to the host in the vSphere Client.
2. Under **Hardware**, expand **CPU** to view the information about the number and type of physical processors and the number of logical processors.

Note: In hyperthreaded systems, each hardware thread is a logical processor. For example, a dual-core processor with hyperthreading activated has two cores and four logical processors.

Specifying CPU Configuration

You can specify CPU configuration to improve resource management. However, if you do not customize CPU configuration, the ESXi host uses defaults that work well in most situations.

You can specify CPU configuration in the following ways:

- Use the attributes and special features available through the vSphere Client. The vSphere Client allows you to connect to the ESXi host or a vCenter Server system.
- Use advanced settings under certain circumstances.
- Use the vSphere SDK for scripted CPU allocation.
- Use hyperthreading.

Multicore Processors

Multicore processors provide many advantages for a host performing multitasking of virtual machines.

Note: In this topic, "Memory" can refer to physical RAM or Persistent Memory.

Intel and AMD have developed processors which combine two or more processor cores into a single integrated circuit (often called a package or socket). VMware uses the term socket to describe a single package which can have one or more processor cores with one or more logical processors in each core.

A dual-core processor, for example, provides almost double the performance of a single-core processor, by allowing two virtual CPUs to run at the same time. Cores within the same processor are typically configured with a shared last-level cache used by all cores, potentially reducing the need to access slower main memory. A shared memory bus that connects a physical processor to main memory can limit performance of its logical processors when the virtual machines running on them are running memory-intensive workloads which compete for the same memory bus resources.

Each logical processor of each processor core is used independently by the ESXi CPU scheduler to run virtual machines, providing capabilities similar to SMP systems. For example, a two-way virtual machine can have its virtual processors running on logical processors that belong to the same core, or on logical processors on different physical cores.

The ESXi CPU scheduler can detect the processor topology and the relationships between processor cores and the logical processors on them. It uses this information to schedule virtual machines and optimize performance.

The ESXi CPU scheduler can interpret processor topology, including the relationship between sockets, cores, and logical processors. The scheduler uses topology information to optimize the placement of virtual CPUs onto different sockets. This optimization can maximize overall cache usage, and to improve cache affinity by minimizing virtual CPU migrations.

Hyperthreading with vSphere

Hyperthreading technology allows a single physical processor core to behave like two logical processors. The processor can run two independent applications at the same time. To avoid confusion between logical and physical processors, Intel refers to a physical processor as a socket, and the discussion in this chapter uses that terminology as well.

Intel Corporation developed hyperthreading technology to enhance the performance of its Pentium IV and Xeon processor lines. Hyperthreading technology allows a single processor core to execute two independent threads simultaneously.

While hyperthreading does not double the performance of a system, it can increase performance by better utilizing idle resources leading to greater throughput for certain important workload types. An application running on one logical processor of a busy core can expect slightly more than half of the throughput that it obtains while running alone on a non-hyperthreaded processor. Hyperthreading performance improvements are highly application-dependent, and some applications might see performance degradation with hyperthreading because many processor resources (such as the cache) are shared between logical processors.

Note: On processors with Intel Hyper-Threading technology, each core can have two logical processors which share most of the core's resources, such as memory caches and functional units. Such logical processors are usually called threads.

Some processors do not support hyperthreading and as a result have only one thread per core. For such processors, the number of cores also matches the number of logical processors.

Hyperthreading and ESXi Hosts

A host that is enabled for hyperthreading should behave similarly to a host without hyperthreading. You might need to consider certain factors if you enable hyperthreading, however.

ESXi hosts manage processor time intelligently to guarantee that load is spread smoothly across processor cores in the system. Logical processors on the same core have consecutive CPU numbers, so that CPUs 0 and 1 are on the first core together, CPUs 2 and 3 are on the second core, and so on. Virtual machines are preferentially scheduled on two different cores rather than on two logical processors on the same core.

If there is no work for a logical processor, it is put into a halted state, which frees its execution resources and allows the virtual machine running on the other logical processor on the same core to use the full execution resources of the core. The VMware scheduler properly accounts for this halt time, and charges a virtual machine running with the full resources of a core more than a virtual machine running on a half core. This approach to processor management ensures that the server does not violate any of the standard ESXi resource allocation rules.

Consider your resource management needs before you enable CPU affinity on hosts using hyperthreading. For example, if you bind a high priority virtual machine to CPU 0 and another high priority virtual machine to CPU 1, the two virtual machines have to share the same physical core. In this case, it can be impossible to meet the resource demands of these virtual machines. Ensure that any custom affinity settings make sense for a hyperthreaded system.

Enable Hyperthreading

To enable hyperthreading, you must first enable it in your system's BIOS settings and then turn it on in the vSphere Client. Hyperthreading is enabled by default.

Consult your system documentation to determine whether your CPU supports hyperthreading.

1. Ensure that your system supports hyperthreading technology.
2. Enable hyperthreading in the system BIOS.
Some manufacturers label this option **Logical Processor**, while others call it **Enable Hyperthreading**.
3. Ensure that hyperthreading is enabled for the ESXi host.
 - a) Browse to the host in the vSphere Client.
 - b) Click **Configure**.
 - c) Under **System**, click **Advanced System Settings** and select **VMkernel.Boot.hyperthreading**.

You must restart the host for the setting to take effect. Hyperthreading is enabled if the value is `true`.

4. Under **Hardware**, click **Processors** to view the number of Logical processors.

Hyperthreading is enabled.

Using CPU Affinity with vSphere

By specifying a CPU affinity setting for each virtual machine, you can restrict the assignment of virtual machines to a subset of the available processors in multiprocessor systems. By using this feature, you can assign each virtual machine to processors in the specified affinity set.

CPU affinity specifies virtual machine-to-processor placement constraints and is different from the relationship created by a VM-VM or VM-Host affinity rule, which specifies virtual machine-to-virtual machine host placement constraints.

In this context, the term CPU refers to a logical processor on a hyperthreaded system and refers to a core on a non-hyperthreaded system.

The CPU affinity setting for a virtual machine applies to all of the virtual CPUs associated with the virtual machine and to all other threads (also known as worlds) associated with the virtual machine. Such virtual machine threads perform processing required for emulating mouse, keyboard, screen, CD-ROM, and miscellaneous legacy devices.

In some cases, such as display-intensive workloads, significant communication might occur between the virtual CPUs and these other virtual machine threads. Performance might degrade if the virtual machine's affinity setting prevents these additional threads from being scheduled concurrently with the virtual machine's virtual CPUs. Examples of this include a uniprocessor virtual machine with affinity to a single CPU or a two-way SMP virtual machine with affinity to only two CPUs.

For the best performance, when you use manual affinity settings, VMware recommends that you include at least one additional physical CPU in the affinity setting to allow at least one of the virtual machine's threads to be scheduled at the same time as its virtual CPUs. Examples of this include a uniprocessor virtual machine with affinity to at least two CPUs or a two-way SMP virtual machine with affinity to at least three CPUs.

Assign a Virtual Machine to a Specific Processor

Using CPU affinity, you can assign a virtual machine to a specific processor. This allows you to restrict the assignment of virtual machines to a specific available processor in multiprocessor systems.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Select **VMs**.
2. Right-click the virtual machine and click **Edit Settings**.
3. Under Virtual Hardware, expand **CPU**.
4. Under Scheduling Affinity, select physical processor affinity for the virtual machine.

Use '-' for ranges and ',' to separate values.

For example, "0, 2, 4-7" would indicate processors 0, 2, 4, 5, 6 and 7.
5. Select the processors where you want the virtual machine to run and click **OK**.

Potential Issues with CPU Affinity

Before you use CPU affinity, you might need to consider certain issues.

Potential issues with CPU affinity include:

- For multiprocessor systems, ESXi systems perform automatic load balancing. Avoid manual specification of virtual machine affinity to improve the scheduler's ability to balance load across processors.

- Affinity can interfere with the ESXi host's ability to meet the reservation and shares specified for a virtual machine.
- Because CPU admission control does not consider affinity, a virtual machine with manual affinity settings might not always receive its full reservation.
Virtual machines that do not have manual affinity settings are not adversely affected by virtual machines with manual affinity settings.
- When you move a virtual machine from one host to another, affinity might no longer apply because the new host might have a different number of processors.
- The NUMA scheduler might not be able to manage a virtual machine that is already assigned to certain processors using affinity. For more information, see [Using NUMA Systems with ESXi](#).
- Affinity can affect the host's ability to schedule virtual machines on multicore or hyperthreaded processors to take full advantage of resources shared on such processors.

Host Power Management Policies in ESXi

You can apply several power management features in ESXi that the host hardware provides to adjust the balance between performance and power. You can control how ESXi uses these features by selecting a power management policy.

Selecting a high-performance policy provides more absolute performance, but at lower efficiency and performance per watt. Low-power policies provide less absolute performance, but at higher efficiency.

You can select a policy for the host that you manage by using the VMware Host Client. If you do not select a policy, ESXi uses Balanced by default.

Table 357: CPU Power Management Policies

Power Management Policy	Description
High Performance	Do not use any power management features.
Balanced (Default)	Reduce energy consumption with minimal performance compromise
Low Power	Reduce energy consumption at the risk of lower performance
Custom	User-defined power management policy. Advanced configuration becomes available.

When a CPU runs at lower frequency, it can also run at lower voltage, which saves power. This type of power management is typically called Dynamic Voltage and Frequency Scaling (DVFS). ESXi attempts to adjust CPU frequencies so that virtual machine performance is not affected.

When a CPU is idle, ESXi can apply deep halt states, also known as C-states. The deeper the C-state, the less power the CPU uses, but it also takes longer for the CPU to start running again. When a CPU becomes idle, ESXi applies an algorithm to predict the idle state duration and chooses an appropriate C-state to enter. In power management policies that do not use deep C-states, ESXi uses only the shallowest halt state for idle CPUs, C1.

Select a CPU Power Management Policy

You set the CPU power management policy for a host using the vSphere Client.

Verify that the BIOS settings on the host system allow the operating system to control power management (for example, **OS Controlled**). If the host hardware does not allow the operating system to manage power, only the Not Supported policy is available. (On some systems, only the High Performance policy is available.)

1. Browse to the host in the vSphere Client.
2. Click **Configure**.
3. Under Hardware, select Overview and click **Edit Power Policy**.
4. Select a power management policy for the host and click **OK**.

The policy selection is saved in the host configuration and can be used again at boot time. You can change it at any time, and it does not require a server reboot.

Configure Custom Policy Parameters for Host Power Management

When you use the Custom policy for host power management, ESXi bases its power management policy on the values of several advanced configuration parameters.

Select **Custom** for the power management policy, as described in [Select a CPU Power Management Policy](#).

1. Browse to the host in the vSphere Client.
2. Click **Configure**.
3. Under **System**, select **Advanced System Settings**.
4. In the right pane, you can edit the power management parameters that affect the Custom policy.

Power management parameters that affect the Custom policy have descriptions that begin with **In Custom policy**. All other power parameters affect all power management policies.

5. Select the parameter and click the **Edit** button.

Note: The default values of power management parameters match the Balanced policy.

Parameter	Description
Power.UsePStates	Use ACPI P-states to save power when the processor is busy.
Power.MaxCpuLoad	Use P-states to save power on a CPU only when the CPU is busy for less than the given percentage of real time.
Power.MinFreqPct	Do not use any P-states slower than the given percentage of full CPU speed.
Power.UseStallCtr	Use a deeper P-state when the processor is frequently stalled waiting for events such as cache misses.
Power.TimerHz	Controls how many times per second ESXi reevaluates which P-state each CPU should be in.
Power.UseCStates	Use deep ACPI C-states (C2 or below) when the processor is idle.
Power.CStateMaxLatency	Do not use C-states whose latency is greater than this value.
Power.CStateResidencyCoef	When a CPU becomes idle, choose the deepest C-state whose latency multiplied by this value is less than the host's prediction of how long the CPU will remain idle. Larger values make ESXi more conservative about using deep C-states, while smaller values are more aggressive.
Power.CStatePredictionCoef	A parameter in the ESXi algorithm for predicting how long a CPU that becomes idle will remain idle. Changing this value is not recommended.
Power.PerfBias	Performance Energy Bias Hint (Intel-only). Sets an MSR on Intel processors to an Intel-recommended value. Intel recommends 0 for high performance, 6 for balanced, and 15 for low power. Other values are undefined.

6. Click **OK**.

Memory Virtualization with vSphere

Before you manage memory resources, you should understand how they are being virtualized and used by ESXi.

The VMkernel manages all physical RAM on the host. The VMkernel dedicates part of this managed physical RAM for its own use. The rest is available for use by virtual machines.

The virtual and physical memory space is divided into blocks called pages. When physical memory is full, the data for virtual pages that are not present in physical memory are stored on disk. Depending on processor architecture, pages are typically 4 KB or 2 MB. See [Advanced Memory Attributes](#).

Virtual Machine Memory

Each virtual machine consumes memory based on its configured size, plus additional overhead memory for virtualization.

The configured size is the amount of memory that is presented to the guest operating system. This is different from the amount of physical RAM that is allocated to the virtual machine. The latter depends on the resource settings (shares, reservation, limit) and the level of memory pressure on the host.

For example, consider a virtual machine with a configured size of 1GB. When the guest operating system boots, it detects that it is running on a dedicated machine with 1GB of physical memory. In some cases, the virtual machine might be allocated the full 1GB. In other cases, it might receive a smaller allocation. Regardless of the actual allocation, the guest operating system continues to behave as though it is running on a dedicated machine with 1GB of physical memory.

Shares

Specify the relative priority for a virtual machine if more than the reservation is available.

Reservation

Is a guaranteed lower bound on the amount of physical RAM that the host reserves for the virtual machine, even when memory is overcommitted. Set the reservation to a level that ensures the virtual machine has sufficient memory to run efficiently, without excessive paging.

After a virtual machine consumes all of the memory within its reservation, it is allowed to retain that amount of memory and this memory is not reclaimed, even if the virtual machine becomes idle. Some guest operating systems (for example, Linux) might not access all of the configured memory immediately after booting. Until the virtual machines consumes all of the memory within its reservation, VMkernel can allocate any unused portion of its reservation to other virtual machines. However, after the guest's workload increases and the virtual machine consumes its full reservation, it is allowed to keep this memory.

Limit

Is an upper bound on the amount of physical RAM that the host can allocate to the virtual machine. The virtual machine's memory allocation is also implicitly limited by its configured size.

Memory Overcommitment

For each running virtual machine, the system reserves physical RAM for the virtual machine's reservation (if any) and for its virtualization overhead.

The total configured memory sizes of all virtual machines may exceed the amount of available physical memory on the host. However, it doesn't necessarily mean memory is overcommitted. Memory is overcommitted when the combined working memory footprint of all virtual machines exceed that of the host memory sizes.

Because of the memory management techniques the ESXi host uses, your virtual machines can use more virtual RAM than there is physical RAM available on the host. For example, you can have a host with 2GB memory and run four virtual

machines with 1GB memory each. In that case, the memory is overcommitted. For instance, if all four virtual machines are idle, the combined consumed memory may be well below 2GB. However, if all 4GB virtual machines are actively consuming memory, then their memory footprint may exceed 2GB and the ESXi host will become overcommitted.

Overcommitment makes sense because, typically, some virtual machines are lightly loaded while others are more heavily loaded, and relative activity levels vary over time.

To improve memory utilization, the ESXi host transfers memory from idle virtual machines to virtual machines that need more memory. Use the Reservation or Shares parameter to preferentially allocate memory to important virtual machines. This memory remains available to other virtual machines if it is not in use. ESXi implements various mechanisms such as ballooning, memory sharing, memory compression and swapping to provide reasonable performance even if the host is not heavily memory overcommitted.

An ESXi host can run out of memory if virtual machines consume all reservable memory in a memory overcommitted environment. Although the powered on virtual machines are not affected, a new virtual machine might fail to power on due to lack of memory.

Note: All virtual machine memory overhead is also considered reserved.

In addition, memory compression is enabled by default on ESXi hosts to improve virtual machine performance when memory is overcommitted as described in [Memory Compression](#).

Memory Sharing

Memory sharing is a proprietary ESXi technique that can help achieve greater memory density on a host.

Memory sharing relies on the observation that several virtual machines might be running instances of the same guest operating system. These virtual machines might have the same applications or components loaded, or contain common data. In such cases, a host uses a proprietary Transparent Page Sharing (TPS) technique to eliminate redundant copies of memory pages. With memory sharing, a workload running on a virtual machine often consumes less memory than it might when running on physical machines. As a result, higher levels of overcommitment can be supported efficiently. The amount of memory saved by memory sharing depends on whether the workload consists of nearly identical machines which might free up more memory. A more diverse workload might result in a lower percentage of memory savings.

Note:

Due to security concerns, inter-virtual machine transparent page sharing is deactivated by default and page sharing is being restricted to intra-virtual machine memory sharing. Page sharing does not occur across virtual machines and only occurs inside a virtual machine. See [Sharing Memory Across Virtual Machines](#) for more information.

Memory Virtualization

Because of the extra level of memory mapping introduced by virtualization, ESXi can effectively manage memory across all virtual machines.

Some of the physical memory of a virtual machine might be mapped to shared pages or to pages that are unmapped, or swapped out.

A host performs virtual memory management without the knowledge of the guest operating system and without interfering with the guest operating system's own memory management subsystem.


The VMM for each virtual machine maintains a mapping from the guest operating system's physical memory pages to the physical memory pages on the underlying machine. (VMware refers to the underlying host physical pages as "machine" pages and the guest operating system's physical pages as "physical" pages.)

Each virtual machine sees a contiguous, zero-based, addressable physical memory space. The underlying machine memory on the server used by each virtual machine is not necessarily contiguous.

The guest virtual to guest physical addresses are managed by the guest operating system. The hypervisor is only responsible for translating the guest physical addresses to machine addresses. Hardware-assisted memory virtualization utilizes the hardware facility to generate the combined mappings with the guest's page tables and the nested page tables maintained by the hypervisor.

The diagram illustrates the ESXi implementation of memory virtualization.

Figure 126: ESXi Memory Mapping



- The boxes represent pages, and the arrows show the different memory mappings.
- The arrows from guest virtual memory to guest physical memory show the mapping maintained by the page tables in the guest operating system. (The mapping from virtual memory to linear memory for x86-architecture processors is not shown.)
- The arrows from guest physical memory to machine memory show the mapping maintained by the VMM.
- The dashed arrows show the mapping from guest virtual memory to machine memory in the shadow page tables also maintained by the VMM. The underlying processor running the virtual machine uses the shadow page table mappings.

Hardware-Assisted Memory Virtualization

Some CPUs, such as AMD SVM-V and the Intel Xeon 5500 series, provide hardware support for memory virtualization by using two layers of page tables.

Note: In this topic, "Memory" can refer to physical RAM or Persistent Memory.

The first layer of page tables stores guest virtual-to-physical translations, while the second layer of page tables stores guest physical-to-machine translation. The TLB (translation look-aside buffer) is a cache of translations maintained by the processor's memory management unit (MMU) hardware. A TLB miss is a miss in this cache and the hardware needs to go to memory (possibly many times) to find the required translation. For a TLB miss to a certain guest virtual address, the hardware looks at both page tables to translate guest virtual address to machine address. The first layer of page tables is maintained by the guest operating system. The VMM only maintains the second layer of page tables.

Performance Considerations

When you use hardware assistance, you eliminate the overhead for software memory virtualization. In particular, hardware assistance eliminates the overhead required to keep shadow page tables in synchronization with guest page tables. However, the TLB miss latency when using hardware assistance is significantly higher. By default the hypervisor uses large pages in hardware assisted modes to reduce the cost of TLB misses. As a result, whether or not a workload benefits by using hardware assistance primarily depends on the overhead the memory virtualization causes when using software memory virtualization. If a workload involves a small amount of page table activity (such as process creation, mapping the memory, or context switches), software virtualization does not cause significant overhead. Conversely, workloads with a large amount of page table activity are likely to benefit from hardware assistance.

By default the hypervisor uses large pages in hardware assisted modes to reduce the cost of TLB misses. The best performance is achieved by using large pages in both guest virtual to guest physical and guest physical to machine address translations.

The option `LPage.LPageAlwaysTryForNPT` can change the policy for using large pages in guest physical to machine address translations. For more information, see [Advanced Memory Attributes](#).

Support for Large Page Sizes

ESXi provides limited support for large page sizes.

x86 architecture allows system software to use 4KB, 2MB and 1GB pages. We refer to 4KB pages as small pages while 2MB and 1GB pages are referred to as large pages. Large pages relieve translation lookaside buffer (TLB) pressure and reduce the cost of page table walks, which results in improved workload performance.

In virtualized environments, large pages can be used by the hypervisor and the guest operating system independently. While the biggest performance impact is achieved if large pages are used by the guest and the hypervisor, in most cases a performance impact can be observed even if large pages are used only at the hypervisor level.

ESXi hypervisor uses 2MB pages for backing guest vRAM by default. vSphere ESXi provides a limited support for backing guest vRAM with 1GB pages. For more information, see *Backing Guest vRAM with 1GB Pages*.

Administer Your Memory Resources with vSphere

Using the vSphere Client you can view information about and make changes to memory allocation settings. To administer your memory resources effectively, you must also be familiar with memory overhead, idle memory tax, and how ESXi hosts reclaim memory.

When administering memory resources, you can specify memory allocation. If you do not customize memory allocation, the ESXi host uses defaults that work well in most situations.

You can specify memory allocation in several ways.

- Use the attributes and special features available through the vSphere Client. The vSphere Client allows you to connect to the ESXi host or vCenter Server system.
- Use advanced settings.
- Use the vSphere SDK for scripted memory allocation.

Note: In this chapter, "Memory" can refer to physical RAM or Persistent Memory.

Understanding Memory Overhead

Virtualization of memory resources has some associated overhead.

ESXi virtual machines can incur two kinds of memory overhead.

- The additional time to access memory within a virtual machine.
- The extra space needed by the ESXi host for its own code and data structures, beyond the memory allocated to each virtual machine.

ESXi memory virtualization adds little time overhead to memory accesses. Because the processor's paging hardware uses page tables (shadow page tables for software-based approach or two level page tables for hardware-assisted approach) directly, most memory accesses in the virtual machine can execute without address translation overhead.

The memory space overhead has two components.

- A fixed, system-wide overhead for the VMkernel.
- Additional overhead for each virtual machine.

Overhead memory includes space reserved for the virtual machine frame buffer and various virtualization data structures, such as shadow page tables. Overhead memory depends on the number of virtual CPUs and the configured memory for the guest operating system.

Overhead Memory on Virtual Machines

Virtual machines require a certain amount of available overhead memory to power on. You should be aware of the amount of this overhead.

The amount of overhead memory needed for a virtual machine depends on a large number of factors, including the number of vCPUs and memory size, the number and types of devices, the execution mode that the monitor is using and the hardware version of the virtual machine. The version of vSphere you are using can also effect the amount of memory needed. VMX automatically calculates the amount of overhead memory needed for a virtual machine.

In order to find out how much overhead memory is needed for your specific configuration, first power on the virtual machine in question. Look in the `vmware.log` file. When the virtual machine powers on, the amount of overhead memory it needs is printed to the log. Search within the log for `VMMEM` to see the initial and precise amount of overhead memory reserved for the virtual machine.

How ESXi Hosts Allocate Memory

A host allocates the memory specified by the `Limit` parameter to each virtual machine, unless memory is overcommitted. ESXi never allocates more memory to a virtual machine than its specified physical memory size.

For example, a 1GB virtual machine might have the default limit (unlimited) or a user-specified limit (for example 2GB). In both cases, the ESXi host never allocates more than 1GB, the physical memory size that was specified for it.

When memory is overcommitted, each virtual machine is allocated an amount of memory somewhere between what is specified by **Reservation** and what is specified by **Limit**. The amount of memory granted to a virtual machine above its reservation usually varies with the current memory load.

A host determines allocations for each virtual machine based on the number of shares allocated to it and an estimate of its recent working set size.

- **Shares** — ESXi hosts use a modified proportional-share memory allocation policy. Memory shares entitle a virtual machine to a fraction of available physical memory.
- **Working set size** — ESXi hosts estimate the working set for a virtual machine by monitoring memory activity over successive periods of virtual machine execution time. Estimates are smoothed over several time periods using techniques that respond rapidly to increases in working set size and more slowly to decreases in working set size. This approach ensures that a virtual machine from which idle memory is reclaimed can ramp up quickly to its full share-based allocation when it starts using its memory more actively. Memory activity is monitored to estimate the working set sizes for a default period of 60 seconds. To modify this default, adjust the `Mem.SamplePeriod` advanced setting. See [How Do You Set Advanced Host Attributes in the vSphere Client](#).

Memory Tax for Idle Virtual Machines

If a virtual machine is not actively using all of its currently allocated memory, ESXi charges more for idle memory than for memory that is in use. This is done to help prevent virtual machines from hoarding idle memory.

The idle memory tax is applied in a progressive fashion. The effective tax rate increases as the ratio of idle memory to active memory for the virtual machine rises. (In earlier versions of ESXi that did not support hierarchical resource pools, all idle memory for a virtual machine was taxed equally.)

You can modify the idle memory tax rate with the `Mem.IdleTax` option. Use this option, together with the `Mem.SamplePeriod` advanced attribute, to control how the system determines target memory allocations for virtual machines. See [How Do You Set Advanced Host Attributes in the vSphere Client](#).

Note: In most cases, changes to `Mem.IdleTax` are not necessary nor appropriate.

VMX Swap Files

Virtual machine executable (VMX) swap files allow the host to greatly reduce the amount of overhead memory reserved for the VMX process.

Note: VMX swap files are not related to the swap to host swap cache feature or to regular host-level swap files.

ESXi reserves memory per virtual machine for a variety of purposes. Memory for the needs of certain components, such as the virtual machine monitor (VMM) and virtual devices, is fully reserved when a virtual machine is powered on. However, some of the overhead memory that is reserved for the VMX process can be swapped. The VMX swap feature reduces the VMX memory reservation significantly (for example, from about 50MB or more per virtual machine to about 10MB per virtual machine). This allows the remaining memory to be swapped out when host memory is overcommitted, reducing overhead memory reservation for each virtual machine.

The host creates VMX swap files automatically, provided there is sufficient free disk space at the time a virtual machine is powered on.

Memory Reclamation

ESXi hosts can reclaim memory from virtual machines.


A host allocates the amount of memory specified by a reservation directly to a virtual machine. Anything beyond the reservation is allocated using the host's physical resources or, when physical resources are not available, handled using special techniques such as ballooning or swapping. Hosts can use two techniques for dynamically expanding or contracting the amount of memory allocated to virtual machines.

- ESXi systems use a memory balloon driver (`vmmemctl`), loaded into the guest operating system running in a virtual machine. See [Memory Balloon Driver](#).
- ESXi system swaps out a page from a virtual machine to a server swap file without any involvement by the guest operating system. Each virtual machine has its own swap file.

Memory Balloon Driver

The memory balloon driver (`vmmemctl`) collaborates with the server to reclaim pages that are considered least valuable by the guest operating system.

The driver uses a proprietary ballooning technique that provides predictable performance that closely matches the behavior of a native system under similar memory constraints. This technique increases or decreases memory pressure on the guest operating system, causing the guest to use its own native memory management algorithms. When memory is tight, the guest operating system determines which pages to reclaim and, if necessary, swaps them to its own virtual disk.

Figure 127: Memory Ballooning in the Guest Operating System

Note: You must configure the guest operating system with sufficient swap space. Some guest operating systems have additional limitations.

If necessary, you can limit the amount of memory `vmxmemctl` reclaims by setting the **`sched.mem.maxmemctl`** parameter for a specific virtual machine. This option specifies the maximum amount of memory that can be reclaimed from a virtual machine in megabytes (MB). See [Set Advanced Virtual Machine Attributes](#).

Sharing Memory Across Virtual Machines

Many ESXi workloads present opportunities for sharing memory across virtual machines (as well as within a single virtual machine).

ESXi memory sharing runs as a background activity that scans for sharing opportunities over time. The amount of memory saved varies over time. For a fairly constant workload, the amount generally increases slowly until all sharing opportunities are exploited.

To determine the effectiveness of memory sharing for a given workload, try running the workload, and use `resxtop` or `esxstop` to observe the actual savings. Find the information in the `PSHARE` field of the interactive mode in the Memory page.

Use the `Mem.ShareScanTime` and `Mem.ShareScanGHz` advanced settings to control the rate at which the system scans memory to identify opportunities for sharing memory.

You can also configure sharing for individual virtual machines by setting the `sched.mem.pshare.enable` option.

Due to security concerns, inter-virtual machine transparent page sharing is deactivated by default and page sharing is being restricted to intra-virtual machine memory sharing. This means page sharing does not occur across virtual machines and only occurs inside of a virtual machine. The concept of salting has been introduced to help address concerns system administrators may have over the security implications of transparent page sharing. Salting can be used to allow more granular management of the virtual machines participating in transparent page sharing than was previously possible. With the new salting settings, virtual machines can share pages only if the salt value and contents of the pages are identical. A new host config option `Mem.ShareForceSalting` can be configured to activate or deactivate salting.

See [Advanced Attributes in vSphere](#) for information on how to set advanced options.

Memory Compression

ESXi provides a memory compression cache to improve virtual machine performance when you use memory overcommitment. Memory compression is activated by default. When a host's memory becomes overcommitted, ESXi compresses virtual pages and stores them in memory.

Because accessing compressed memory is faster than accessing memory that is swapped to disk, memory compression in ESXi allows you to overcommit memory without significantly hindering performance. When a virtual page needs to be swapped, ESXi first attempts to compress the page. Pages that can be compressed to 2 KB or smaller are stored in the virtual machine's compression cache, increasing the capacity of the host.

You can set the maximum size for the compression cache and deactivate memory compression using the Advanced Settings dialog box in the vSphere Client.

Activate or Deactivate the Memory Compression Cache

Memory compression is activated by default. You can use Advanced System Settings in the vSphere Client to activate or deactivate memory compression for a host.

1. Browse to the host in the vSphere Client.
2. Click **Configure**.
3. Under **System**, select **Advanced System Settings**.
4. Locate Mem.MemZipEnable and click the **Edit** button.
5. Enter 1 to activate or enter 0 to deactivate the memory compression cache.
6. Click **OK**.

Set the Maximum Size of the Memory Compression Cache

You can set the maximum size of the memory compression cache for the host's virtual machines.

You set the size of the compression cache as a percentage of the memory size of the virtual machine. For example, if you enter 20 and a virtual machine's memory size is 1000 MB, ESXi can use up to 200MB of host memory to store the compressed pages of the virtual machine.

If you do not set the size of the compression cache, ESXi uses the default value of 10 percent.

1. Browse to the host in the vSphere Client.
2. Click **Configure**.
3. Under **System**, select **Advanced System Settings**.
4. Locate Mem.MemZipMaxPct and click the **Edit** button.
The value of this attribute determines the maximum size of the compression cache for the virtual machine.
5. Enter the maximum size for the compression cache.
The value is a percentage of the size of the virtual machine and must be between 5 and 100 percent.
6. Click **OK**.

Measuring and Differentiating Types of Memory Usage


The **Performance** tab of the vSphere Client displays several metrics that can be used to analyze memory usage.

Some of these memory metrics measure guest physical memory while other metrics measure machine memory. For instance, two types of memory usage that you can examine using performance metrics are guest physical memory and machine memory. You measure guest physical memory using the Memory Granted metric (for a virtual machine) or Memory Shared (for a host). To measure machine memory, however, use Memory Consumed (for a virtual machine) or Memory Shared Common (for a host). Understanding the conceptual difference between these types of memory usage is important for knowing what these metrics are measuring and how to interpret them.

The VMkernel maps guest physical memory to machine memory, but they are not always mapped one-to-one. Multiple regions of guest physical memory might be mapped to the same region of machine memory (when memory sharing) or specific regions of guest physical memory might not be mapped to machine memory (when the VMkernel swaps out or balloons guest physical memory). In these situations, calculations of guest physical memory usage and machine memory usage for an individual virtual machine or a host differ.

Consider the example in the following figure, which shows two virtual machines running on a host. Each block represents 4 KB of memory and each color/letter represents a different set of data on a block.

Figure 128: Memory Usage Example



The performance metrics for the virtual machines can be determined as follows:

- To determine Memory Granted (the amount of guest physical memory that is mapped to machine memory) for virtual machine 1, count the number of blocks in virtual machine 1's guest physical memory that have arrows to machine memory and multiply by 4 KB. Since there are five blocks with arrows, Memory Granted is 20 KB.
- Memory Consumed is the amount of machine memory allocated to the virtual machine, accounting for savings from shared memory. First, count the number of blocks in machine memory that have arrows from virtual machine 1's guest physical memory. There are three such blocks, but one block is shared with virtual machine 2. So count two full blocks plus half of the third and multiply by 4 KB for a total of 10 KB Memory Consumed.

The important difference between these two metrics is that Memory Granted counts the number of blocks with arrows at the guest physical memory level and Memory Consumed counts the number of blocks with arrows at the machine memory level. The number of blocks differs between the two levels due to memory sharing and so Memory Granted and Memory Consumed differ. Memory is being saved through sharing or other reclamation techniques.

A similar result is obtained when determining Memory Shared and Memory Shared Common for the host.

- Memory Shared for the host is the sum of each virtual machine's Memory Shared. Calculate shared memory by looking at each virtual machine's guest physical memory and counting the number of blocks that have arrows to machine memory blocks that themselves have more than one arrow pointing at them. There are six such blocks in the example, so Memory Shared for the host is 24 KB.
- Memory Shared Common is the amount of machine memory shared by virtual machines. To determine common memory, look at the machine memory and count the number of blocks that have more than one arrow pointing at them. There are three such blocks, so Memory Shared Common is 12 KB.

Memory Shared is concerned with guest physical memory and looks at the origin of the arrows. Memory Shared Common, however, deals with machine memory and looks at the destination of the arrows.

The memory metrics that measure guest physical memory and machine memory might appear contradictory. In fact, they are measuring different aspects of a virtual machine's memory usage. By understanding the differences between these metrics, you can better use them to diagnose performance issues.

Memory Reliability

Memory reliability, also known as error isolation, allows ESXi to stop using parts of memory when it determines that a failure might occur, as well as when a failure did occur.

When enough corrected errors are reported at a particular address, ESXi stops using this address to prevent the corrected error from becoming an uncorrected error.

Memory reliability provides better VMkernel reliability despite corrected and uncorrected errors in RAM. It also enables the system to avoid using memory pages that might contain errors.

Correcting an Error Isolation Notification

With memory reliability, VMkernel stops using pages that receive an error isolation notification.

The user receives an event in the vSphere Client when VMkernel recovers from an uncorrectable memory error, when VMkernel retires a significant percentage of system memory due to a large number of correctable errors, or if there are a large number of pages that are unable to retire.

1. Vacate the host.
2. Migrate the virtual machines.
3. Run memory related hardware tests.

System Swap

System swap is a memory reclamation process that can take advantage of unused memory resources across an entire system.

System swap allows the system to reclaim memory from memory consumers that are not virtual machines. When system swap is enabled you have a tradeoff between the impact of reclaiming the memory from another process and the ability to assign the memory to a virtual machine that can use it. The amount of space required for the system swap is 1GB.

Memory is reclaimed by taking data out of memory and writing it to background storage. Accessing the data from background storage is slower than accessing data from memory, so it is important to carefully select where to store the swapped data.

ESXi determines automatically where the system swap should be stored, this is the **Preferred swap file location**. This decision can be aided by selecting a certain set of options. The system selects the best possible enabled option. If none of the options are feasible then system swap is not activated.

The available options are:

- Datastore - Allow the use of the datastore specified. Please note that a vSAN datastore or a VMware vSphere® Virtual Volumes™ datastore cannot be specified for system swap files.
- Host Swap Cache - Allow the use of part of the host swap cache.
- Preferred swap file location - Allow the use of the preferred swap file location configured for the host.

Configure System Swap

You can customize the options that determine the system swap location.

Select the **Enabled** check box in the **Edit System Swap Settings** dialog box.

1. Browse to the host in the vSphere Client.
2. Click **Configure**.
3. Under **System**, select **System Swap**.
4. Click **Edit**.
5. Select the check boxes for each option that you want to enable.
6. If you select the **datastore** option, select a datastore from the drop-down menu.
7. Click **OK**.

Using Swap Files with vSphere

You can specify the location of your guest swap file, reserve swap space when memory is overcommitted, and delete a swap file.

ESXi hosts use swapping to forcibly reclaim memory from a virtual machine when the `vmmemctl` driver is not available or is not responsive.

- It was never installed.
- It is explicitly deactivated.
- It is not running (for example, while the guest operating system is booting).
- It is temporarily unable to reclaim memory quickly enough to satisfy current system demands.
- It is functioning properly, but maximum balloon size is reached.

Standard demand-paging techniques swap pages back in when the virtual machine needs them.

Swap File Location

By default, the swap file is created in the same location as the virtual machine's configuration file, which may either be on a VMFS datastore, a vSAN datastore or a VMware vSphere® Virtual Volumes™ datastore. On a vSAN datastore or a vVols datastore, the swap file is created as a separate vSAN or vVols object.

The ESXi host creates a swap file when a virtual machine is powered on. If this file cannot be created, the virtual machine cannot power on. Instead of accepting the default, you can also:

- Use per-virtual machine configuration options to change the datastore to another shared storage location.
- Use host-local swap, which allows you to specify a datastore stored locally on the host. This allows you to swap at a per-host level, saving space on the SAN. However, it can lead to a slight degradation in performance for vSphere vMotion because pages swapped to a local swap file on the source host must be transferred across the network to the destination host. Currently vSAN and vVols datastores cannot be specified for host-local swap.

Enable Host-Local Swap for a DRS Cluster

Host-local swap allows you to specify a datastore stored locally on the host as the swap file location. You can enable host-local swap for a DRS cluster.

1. Browse to the cluster in the vSphere Client.
2. Click **Configure**.
3. Under **Configuration**, select **General** to view the swap file location and click **Edit** to change it.
4. Select the **Datastore specified by host** option and click **OK**.
5. Browse to one of the hosts in the cluster in the vSphere Client.
6. Click **Configure**.
7. Under Virtual Machines, select **Swap file location**.
8. Click **Edit** and select the local datastore to use and click **OK**.
9. Repeat 5 through 8 for each host in the cluster.

Host-local swap is now enabled for the DRS cluster.

Enable Host-Local Swap for a Standalone Host

Host-local swap allows you to specify a datastore stored locally on the host as the swap file location. You can enable host-local swap for a standalone host.

1. Browse to the host in the vSphere Client.
2. Click **Configure**.
3. Under **Virtual Machines**, select **Swap file location**.
4. Click **Edit** and select **Selected Datastore**.
5. Select a local datastore from the list and click **OK**.

Host-local swap is now enabled for the standalone host.

Swap Space and Memory Overcommitment

You must reserve swap space for any unreserved virtual machine memory (the difference between the reservation and the configured memory size) on per-virtual machine swap files.

This swap reservation is required to ensure that the ESXi host is able to preserve virtual machine memory under any circumstances. In practice, only a small fraction of the host-level swap space might be used.

If you are overcommitting memory with ESXi, to support the intra-guest swapping induced by ballooning, ensure that your guest operating systems also have sufficient swap space. This guest-level swap space must be greater than or equal to the difference between the virtual machine's configured memory size and its Reservation.

CAUTION: If memory is overcommitted, and the guest operating system is configured with insufficient swap space, the guest operating system in the virtual machine can fail.

To prevent virtual machine failure, increase the size of the swap space in your virtual machines.

- Windows guest operating systems— Windows operating systems refer to their swap space as paging files. Some Windows operating systems try to increase the size of paging files automatically, if there is sufficient free disk space. See your Microsoft Windows documentation or search the Windows help files for “paging files.” Follow the instructions for changing the size of the virtual memory paging file.

- Linux guest operating system — Linux operating systems refer to their swap space as swap files. For information on increasing swap files, see the following Linux man pages:
 - `mkswap` — Sets up a Linux swap area.
 - `swapon` — Enables devices and files for paging and swapping.

Guest operating systems with a lot of memory and small virtual disks (for example, a virtual machine with 8GB RAM and a 2GB virtual disk) are more susceptible to having insufficient swap space.

Note: Do not store swap files on thin-provisioned LUNs. Running a virtual machine with a swap file that is stored on a thin-provisioned LUN can cause swap file growth failure, which can lead to termination of the virtual machine.

When you create a large swap file (for example, larger than 100GB), the amount of time it takes for the virtual machine to power on can increase significantly. To avoid this, set a high reservation for large virtual machines.

You can also place swap files on less costly storage using host-local swap files.

Configure Virtual Machine Swapfile Properties for the Host

Configure a swapfile location for the host to determine the default location for virtual machine swapfiles in the vSphere Client.

Required privilege: **Host machine > Configuration > Storage partition configuration**

By default, swapfiles for a virtual machine are located on a datastore in the folder that contains the other virtual machine files. However, you can configure your host to place virtual machine swapfiles on an alternative datastore.

You can use this option to place virtual machine swapfiles on lower-cost or higher-performance storage. You can also override this host-level setting for individual virtual machines.

Setting an alternative swapfile location might cause migrations with vMotion to complete more slowly. For best vMotion performance, store the virtual machine on a local datastore rather than in the same directory as the virtual machine swapfiles. If the virtual machine is stored on a local datastore, storing the swapfile with the other virtual machine files will not improve vMotion.

- Browse to the host in the vSphere Client.
- Click **Configure**.
- Under **Virtual Machines**, click **Swap file location**.

The selected swapfile location is displayed. If configuration of the swapfile location is not supported on the selected host, the tab indicates that the feature is not supported.

If the host is part of a cluster, and the cluster settings specify that swapfiles are to be stored in the same directory as the virtual machine, you cannot edit the swapfile location from the host under **Configure**. To change the swapfile location for such a host, edit the cluster settings.

- Click **Edit**.
- Select where to store the swapfile.

Option	Description
Virtual machine directory	Stores the swapfile in the same directory as the virtual machine configuration file.
Use a specific datastore	Stores the swapfile in the location you specify. If the swapfile cannot be stored on the datastore that the host specifies, the swapfile is stored in the same folder as the virtual machine.

6. Optional: If you select **Use a specific datastore**, select a datastore from the list.
7. Click **OK**.

The virtual machine swapfile is stored in the location you selected.

Configure a Virtual Machine Swap File Location for a Cluster

By default, swap files for a virtual machine are on a datastore in the folder that contains the other virtual machine files. However, you can instead configure the hosts in your cluster to place virtual machine swap files on an alternative datastore of your choice.

Before you configure a virtual machine swap file location for a cluster, you must configure the virtual machine swap file locations for the hosts in the cluster as described in [Configure Virtual Machine Swapfile Properties for the Host](#).

You can configure an alternative swap file location to place virtual machine swap files on either lower-cost or higher-performance storage, depending on your needs.

1. Browse to the cluster.
2. Click **Configure**.
3. Select **Configuration > General**.
4. Next to swap file location, click **Edit**.
5. Select where to store the swap file.

Option	Description
Virtual machine directory	Stores the swap file in the same directory as the virtual machine configuration file.
Datastore specified by host	Stores the swap file in the location specified in the host configuration. If the swap file cannot be stored on the datastore that the host specifies, the swap file is stored in the same folder as the virtual machine.

6. Click **OK**.

Delete Swap Files

If a host fails, and that host had running virtual machines that were using swap files, those swap files continue to exist and consume many gigabytes of disk space. You can delete the swap files to eliminate this problem.

1. Restart the virtual machine that was on the host that failed.
2. Stop the virtual machine.

The swap file for the virtual machine is deleted.

Using Persistent Memory with vSphere

Persistent Memory (PMem), also known as Non-Volatile Memory (NVM), is capable of maintaining data even after a power outage. PMem can be used by applications that are sensitive to downtime and require high performance.

VMs can be configured to use PMem on a standalone host, or in a cluster. PMem is treated as a local datastore. Persistent memory significantly reduces storage latency. In ESXi you can create VMs that are configured with PMem, and applications inside these VMs can take advantage of this increased speed. Once a VM is initially powered on, PMem is reserved for it regardless of whether it is powered on or off. This PMem stays reserved until the VM is migrated or removed.

Persistent memory can be consumed by virtual machines in two different modes. Legacy guest OSes can still take advantage of virtual persistent memory disk feature.

- **Virtual Persistent Memory (vPMem)**

Using vPMem, the memory is exposed to a guest OS as a virtual NVDIMM. This enables the guest OS to use PMem in byte addressable random mode.

Note: You must use VM hardware version 14 and a guest OS that supports NVM technology.

Note: You must use VM hardware version 19 when you configure vSphere HA for PMem VMs. For more information, see [Configure vSphere HA for PMem VMs](#).

- **Virtual Persistent Memory Disk (vPMemDisk)**

Using vPMemDisk, the memory can be accessed by the guest OS as a virtual SCSI device, but the virtual disk is stored in a PMem datastore.

When you create a VM with PMem, memory is reserved for it at the time of Hard disk creation. Admission control is also done at the time of Hard disk creation. For more information, see [vSphere HA Admission Control PMem Reservation](#).

In a cluster, each VM has some capacity for PMem. The total amount of PMem must not be greater than the total amount available in the cluster. The consumption of PMem includes both powered on and powered off VMs. If a VM is configured to use PMem and you do not use DRS, then you must manually pick a host that has sufficient PMem to place the VM on.

NVDIMM and traditional storage

NVDIMM is accessed as memory. When you use traditional storage, software exists between applications and storage devices which can cause a delay in processing time. When you use PMem, the applications use the storage directly. This means that PMem performance is better than traditional storage. Storage is local to the host. However, since system software cannot track the changes, solutions such as backups do not currently work with PMem.

Solutions such as vSphere HA have limited scope if vPMem is used in a mode that is not write-through to a non-PMem datastore. When vSphere HA is activated for vPMem VMs with failover enabled, the VM can be failed over to different host. When this happens, the VM is using the PMem resources on the new host. To free up the resources on the old host, a garbage collector periodically identifies and frees up these resources for use by other VMs.

Name spaces

Name spaces for PMem are configured before ESXi starts. Name spaces are similar to disks on the system. ESXi reads name spaces and combines multiple name spaces into one logical volume by writing GPT headers. This is formatted automatically by default, if you have not previously configured it. If it has already been formatted, ESXi attempts to mount the PMem.

GPT tables

If the data in PMem storage is corrupted it can cause ESXi to fail. To avoid this, ESXi checks for errors in the metadata during PMem mount time.

PMem regions

PMem regions are a continuous byte stream that represent a single vNVDimm or vPMemDisk. Each PMem volume belongs to a single host. This could be difficult to manage if an administrator has to manage each host in a cluster with a large number of hosts. However, you do not have to manage each individual datastore. Instead you can think of the entire PMem capacity in the cluster as one datastore.

VC and DRS automate initial placement of PMem datastores. Select a local PMem storage profile when the VM is created or when the device is added to the VM. The rest of the configuration is automated. One limitation is that ESXi does not allow you to put the VM home on a PMem datastore. This is because it takes up valuable space to store VM log and stat files. These regions are used to represent the VM data and can be exposed as byte addressable nvDimms, or VpMem disks.

Migration

Since PMem is a local datastore, if you want to move a VM you must use storage vMotion. A VM with vPMem can only be migrated to an ESX host with PMem resource. A VM with vPMemDisk can be migrated to an ESX host without a PMem resource.

Error handling and NVDimm management

Host failures can result in a loss of availability on vPMem VMs which are not in write-through mode. In the case of catastrophic errors, you may lose all data and must take manual steps to reformat the PMem.

vSphere Persistent Memory with the vSphere Client

For a brief conceptual introduction to Persistent Memory, see:

Enhancements to Working with PMEM in the vSphere Client

For a brief overview of enhancements in the HTML5-based vSphere Client when working with PMem, see:

Migrating and Cloning VMs Using PMEM in the vSphere Client

For a brief overview of migrating and cloning virtual machines that use PMem, see:

Configure vSphere HA for PMem VMs

You can configure vSphere HA for PMem VMs in write-through mode, so that when a host fails VMs can be restored on another functioning host.

- You must select Hardware version 19.
- PMem VMs with vPMemDisks are not supported.

1. When creating a new VM in the **New Virtual Machine** wizard, select **Customize hardware**.

- a) Click **ADD NEW DEVICE** and select **Add NVDIMM** from the drop-down menu.
- b) Click the checkbox **Allow failover on another host for all NVDIMM devices**.
- c) Click **NEXT** and complete the **New Virtual Machine** wizard.

On Host failure, NVDIMM PMem data cannot be recovered. By default, HA will not attempt to restart this virtual machine on another host. Allowing HA on host failure to failover the virtual machine, will restart the virtual machine on another host with a new, empty NVDIMM.

2. To enable HA on an existing VM, browse to the VM.

- a) Under **VM Hardware**, click **EDIT**.
- b) Select the NVDIMM.
- c) Click the checkbox **Allow failover on another host for all NVDIMM devices**.
- d) Click **OK**.

On host failure, HA will restart this virtual machine on another host with new, empty NVDIMMs.

vSphere HA Admission Control PMem Reservation

Admission control is a policy used by vSphere HA to ensure failover capacity within a cluster.

Raising the number of potential host failures to tolerate will increase the availability restraints and capacity reserved. You can reserve a percentage of Persistent Memory for Host failover capacity. This is actual storage capacity that is blocked and must be considered for host power off.

Under **Edit Cluster Settings** you can select **Admission Control** to specify the number of failures the host will tolerate.

If you select CPU/Memory reservation defined by:

- **Cluster resource percentage**, some amount of persistent memory capacity in the cluster is dedicated for failover purpose even if the virtual machines in the cluster are not using persistent memory currently. This percentage can

either be specified through an override, or it is automatically calculated according to the **host failures to tolerate** setting. When PMem admission control is enabled, PMem capacity is reserved across the cluster even if there are VMs using PMem as disks.

- **Slot Policy (powered-on VMs)**, persistent memory admission control overrides the Slot Policy with the Cluster Resource Percentage policy, for persistent memory resources only. The percentage value is automatically calculated from the **host failures cluster tolerates** setting and cannot be overridden.
- **Dedicated failover hosts**, the persistent memory of the dedicated failover hosts is dedicated for failover purpose and you cannot provision virtual machines with persistent memory on these hosts.

Note: After you select an admission control policy, you must also click the **Reserve Persistent Memory failover capacity** checkbox to enable PMem admission control.


vSphere Memory Monitoring and Remediation

vMMR collects data and provides visibility of performance statistics so you can determine if your application workload is regressed due to Memory Mode.

Intel Optane Persistent Memory can be configured in BIOS settings in App Direct or Memory Mode. In App Direct mode, persistent memory can be accessed as byte addressable, persistent memory along with DRAM. In Memory Mode, DRAM becomes the hardware cache and the larger PMem becomes volatile and appears as system memory.

Memory Mode is invisible and transparent to VMs. Once you configure the system in Memory Mode, the system appears as a traditional system with DRAM. A cluster can have a mix of hosts with different configurations. vSphere shows additional information about the system being in Memory Mode. ESXi programs performance counters that gather information about host level and VM level statistics. These performance statistics are used to create alarms. Statistics are also tracked in performance charts.

You can find out if the system is in Memory Mode in the host **Summary** tab, **Memory Tiering: Hardware with some additional details**.



The screenshot shows the vSphere host configuration page with the 'Summary' tab selected. The 'Memory Tiering' section is expanded, showing 'Hardware' as the tiering type. A 'DETAILS' link is visible next to 'Hardware'. A tooltip box is open, displaying the message: 'Intel Optane™ Persistent Memory configured in Memory Mode.' The tooltip has a close button (X) in the top right corner.

Summary	Monitor	Configure	Permissions	VMs	Datastores	Networks	Updates
Logical Processors:	96						
NICs:	3						
Virtual Machines:	1						
Memory Tiering:	Hardware						
State:	Connected						
Uptime:	6 hours						

You can also view the size of DRAM and PMEM under **Configure > Hardware > Overview > Memory**.

Memory	
Total	503.68 GB
System	385.17 MB
Virtual machines	503.3 GB
Memory Tiering	Hardware
Tier 0	256 GB DRAM (Cache)
Tier 1	503.67 GB PMem (Memory)

ESXi gathers and exposes two kinds of memory statistics:

- **Host Level Statistics:** A memory sub-component measures DRAM and PMem performance by programming performance counters. The host level statistics are total, read/write bandwidth, read/write latency and miss rate for the different memory types (DRAM, PMem).
- **VM Level Statistics:** vSphere monitors performance counters to get data on DRAM and PMEM read bandwidth of the VM.

Both Host and VM have new Memory pane under Performance charts. This shows memory details like Memory Utilization and Memory Reclamation along with new statistics. On the ESXi host level, you can monitor the Memory Bandwidth and Memory Miss Rate charts. On the VM level, you can view the PMem read bandwidth and DRAM read bandwidth.

From the **VMs** tab of an ESXi host, you can view a list containing performance information about all virtual machines that reside on the host. To display information about the Memory Mode impact on a virtual machine, click the view columns

icon and select the Active Memory, DRAM Read Bandwidth, and PMem Read Bandwidth metrics.

There are two preconfigured default alarms, one at the host level (Host Memory Mode High Active DRAM Usage) and another at the VM level (Virtual Machine High PMem Bandwidth Usage). If the alarm condition is met, an event will be published to trigger the corresponding alarm. You can also create custom alarms based on performance metrics. vMMR alarms only work on hosts configured with Memory Mode.

When DRS is enabled and fully automated in the cluster, if the active memory utilization of the host is above certain percentage of the size of the DRAM cache, DRS might move some VMs out of the host in order to balance the load.

For more information, see *vSphere Monitoring and Performance*.

Note: vMMR is supported on Intel Broadwell, Skylake, Cascade Lake and Ice Lake platforms. Host Level DRAM statistics are available on these platforms. Host and VM Level PMem statistics are only available in Cascade Lake and Ice Lake hosts configured in Memory Mode.

Configure Virtual Graphics on vSphere

You can edit graphics settings for supported graphics implementations.

vSphere supports multiple graphics implementations.

- VMware supports 3d graphics solutions from AMD, Intel and NVIDIA.
- NVIDIA GRID support.
- Allows single NVIDIA vib to support both vSGA and vGPU implementations.

- Provides vCenter GPU performance charts for Intel and NVIDIA.
- Enables graphics for Horizon View VDI desktops.
- You can create a VM with and up to 16 vGPU devices. HWv21 must be enabled on ESXi and vCenter.

You can configure host graphics settings, and customize vGPU graphics settings on a per VM basis.

Note: In this chapter, "Memory" refers to physical RAM.

View GPU Statistics

You can view detailed information for a host graphics card.

You can see GPU temperature, utilization, and memory usage.

Note: These statistics are only displayed when the GPU driver is installed on the host.

1. In the vSphere Client, navigate to the host.
2. Click the **Monitor** tab and click **Performance**.
3. Click **Advanced** and select **GPU** from the drop-down menu.

Add an NVIDIA GRID vGPU to a Virtual Machine

If an ESXi host has an NVIDIA GRID GPU graphics device, you can configure a virtual machine to use the NVIDIA GRID virtual GPU (vGPU) technology.

- Verify that an NVIDIA GRID GPU graphics device with an appropriate driver is installed on the host. See the *vSphere Upgrade* documentation.
- Verify that the virtual machine is compatible with ESXi 6.0 and later.

NVIDIA GRID GPU graphics devices are designed to optimize complex graphics operations and enable them to run at high performance without overloading the CPU.

1. Right-click a virtual machine and select **Edit Settings**.
2. On the **Virtual Hardware** tab, select **Add New Device** and select **New PCI Device** from the drop-down menu.
3. Expand the **New PCI device**, and select the NVIDIA GRID vGPU passthrough device to which to connect your virtual machine.

Note: Full memory reservation will be applied automatically, it's required for PCI device.

4. Select a GPU profile.
A GPU profile represents the vGPU type.
5. Click **OK**.

The virtual machine can access the device.

Configuring Host Graphics

You can customize the graphics options on a per host basis.

Virtual machines should be powered off.

1. Select a host and select **Configure > Hardware > Graphics**.
2. Under **Host Graphics**, select **Edit**.
3. In the **Edit Host Graphics Settings** window, select:

Option	Description
Shared	VMware shared virtual graphics

Option	Description
Shared Direct	Vendor shared passthrough graphics

4. Select a shared passthrough GPU assignment policy.
 - a) Spread VMs across GPUs (best performance)
 - b) Group VMs on GPU until full (GPU Consolidation)

5. Click **OK**.

After clicking **OK**, you must restart Xorg on the host.

Configuring Graphics Devices

You can edit graphics type for a video card.

Virtual machines must be powered off.

1. Under **Graphics Devices**, select a graphics card and click **Edit**.
 - a) Select **Shared** for VMware shared virtual graphics.
 - b) Select **Shared Direct** for Vendor shared passthrough graphics.
2. Click **OK**.

If you select a device, it shows which virtual machines are using that device if they are active.

After clicking **OK**, you must restart Xorg on the host.

Migrating VMs with vGPUs

You can migrate VM's that have vGPUs.

Starting with vSphere 8.0 U2, DRS can automatically migrate VMs with vGPUs. When the DRS Cluster Advanced Options are set and the Estimated VM Device's Stun Time for a VM is lower than the VM Device's vMotion Stun Time limit, DRS will automate VM migrations.

To enable this functionality, make sure the infrastructure meets the following requirements:

- Healthy vSphere Lifecycle Services (See [KB 91891](#))
- Configure the VM's vGPU devices through vCenter only
- Healthy vMotion network ([vMotion NICs setup](#))

If these requirements are met, add the following DRS Cluster Advanced Options:

DRS Cluster Advanced Options	Value
PassthroughDrsAutomation	1
LBMaxVmotionPerHost	1

For vGPU VMs with Stun Times exceeding the "vMotion Stun Time Limit" (Default 100 seconds), a VI Admin can add the following DRS Cluster Advanced Option:

DRS Cluster Advanced Options	Value
VmDevicesStunTimeTolerated	<number of seconds, greater than any VM's Estimated Stun Time in the Cluster> (Default 100 seconds)

Or you can modify the "vMotion Stun Time Limit" in the **VM Configuration > > VM Options tab > Advanced section**.

For earlier releases:

For Maintenance Mode evacuations, please refer to KB 88271 [vGPU Virtual Machine automated migration for Host Maintenance Mode in a DRS Cluster](#).

If VM placement issues arise, reduce DRS Automation to **Partially Automated** please see [Edit Cluster Settings](#) for more information.

You can manually migrate VMs with vGPUs to another host. See [Using vMotion to Migrate vGPU Virtual Machines](#) in the vCenter Server and Host Management guide for more information.

Configuring vGPU Size

You can adjust the vGPU size in vCenter.

vGPU profiles on a single physical GPU can have different types or sizes. This means, for example, you can run vGPU VM profiles with 16GB compute and 4GB graphics simultaneously on a single device. This allows you greater flexibility with vGPU workloads and better utilization of devices with mixed vGPU workloads.

You can activate or deactivate heterogeneous sizes on a per-device basis. Activating this feature reduces the maximum number of instances of some profiles. For example, the A16-1Q profile supports 24 instances when the feature is deactivated, but it supports 16 instances when it is activated. You may prefer to deactivate the feature on some devices in order to maximize capacity for their workloads.

1. Browse to the ESXi host.
2. You can adjust vGPU size under: **Configure > Hardware > Graphics > Graphics Devices**.
3. Select **vGPU Mode** and choose either **SameSize** or **MixedSize**. The default value is **SameSize**.
 - **SameSize** indicates the sizes are homogeneous.
 - **MixedSize** indicates the sizes can be heterogeneous.

When editing a graphics device, the following message is displayed.

Settings will take effect after restarting the host or xorg service.

This is required when you edit vGPU size.

Manage Storage I/O Resources with vSphere

vSphere Storage I/O Control allows cluster-wide storage I/O prioritization, which allows better workload consolidation and helps reduce extra costs associated with over provisioning.

Storage I/O Control extends the constructs of shares and limits to handle storage I/O resources. You can control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which ensures that more important virtual machines get preference over less important virtual machines for I/O resource allocation.

When you enable Storage I/O Control on a datastore, ESXi begins to monitor the device latency that hosts observe when communicating with that datastore. When device latency exceeds a threshold, the datastore is considered to be congested and each virtual machine that accesses that datastore is allocated I/O resources in proportion to their shares. You set shares per virtual machine. You can adjust the number for each based on need.

The I/O filter framework (VAIO) allows VMware and its partners to develop filters that intercept I/O for each VMDK and provides the desired functionality at the VMDK granularity. VAIO works along Storage Policy-Based Management (SPBM) which allows you to set the filter preferences through a storage policy that is attached to VMDKs.

Configuring Storage I/O Control is a two-step process:

1. Enable Storage I/O Control for the datastore.
2. Set the number of storage I/O shares and upper limit of I/O operations per second (IOPS) allowed for each virtual machine.

By default, all virtual machine shares are set to Normal (1000) with unlimited IOPS.

Note: Storage I/O Control is enabled by default on Storage DRS-enabled datastore clusters.

Note: vSphere 8.0 U1 supports seamless hot-add and hot-remove of a LWD I/O filter during VM reconfiguration.

Note: In this chapter, "Memory" refers to physical RAM.

Virtual Machine Storage Policies

Virtual machine storage policies are essential to virtual machine provisioning. The policies control which type of storage is provided for the virtual machine, how the virtual machine is placed within the storage, and which data services are offered for the virtual machine.

vSphere includes default storage policies. However, you can define and assign new policies.

You use the VM Storage Policies interface to create a storage policy. When you define the policy, you specify various storage requirements for applications that run on virtual machines. You can also use storage policies to request specific data services, such as caching or replication, for virtual disks.

You apply the storage policy when you create, clone, or migrate the virtual machine. After you apply the storage policy, the Storage Policy Based Management (SPBM) mechanism places the virtual machine in a matching datastore and, in certain storage environments, determines how the virtual machine storage objects are provisioned and allocated within the storage resource to guarantee the required level of service. The SPBM also enables requested data services for the virtual machine. vCenter Server monitors policy compliance and sends an alert if the virtual machine is in breach of the assigned storage policy.

See *vSphere Storage* for more information.

I/O Filters

I/O filters that are associated with virtual disks gain direct access to the virtual machine I/O path regardless of the underlying storage topology.

VMware offers certain categories of I/O filters. In addition, the I/O filters can be created by third-party vendors. Typically, they are distributed as packages that provide an installer to deploy filter components on vCenter Server and ESXi host clusters.

When I/O filters are deployed on the ESXi cluster, vCenter Server automatically configures and registers an I/O filter storage provider, also called a VASA provider, for each host in the cluster. The storage providers communicate with vCenter Server and make data services offered by the I/O filter visible in the VM Storage Policies interface. You can reference these data services when defining common rules for a VM policy. After you associate virtual disks with this policy, the I/O filters are enabled on the virtual disks.

See *vSphere Storage* for more information.

Storage I/O Control Requirements

Storage I/O Control has several requirements and limitations.

- Datastores that are Storage I/O Control-enabled must be managed by a single vCenter Server system.
- Storage I/O Control is supported on Fibre Channel-connected, iSCSI-connected, and NFS-connected storage. Raw Device Mapping (RDM) is not supported.
- Storage I/O Control does not support datastores with multiple extents.

- Before using Storage I/O Control on datastores that are backed by arrays with automated storage tiering capabilities, check the *VMware Storage/SAN Compatibility Guide* to verify whether your automated tiered storage array has been certified to be compatible with Storage I/O Control.

Automated storage tiering is the ability of an array (or group of arrays) to migrate LUNs/volumes or parts of LUNs/volumes to different types of storage media (SSD, FC, SAS, SATA) based on user-set policies and current I/O patterns. No special certification is required for arrays that do not have these automatic migration/tiering features, including those that provide the ability to manually migrate data between different types of storage media.

Storage I/O Control Resource Shares and Limits

You allocate the number of storage I/O shares and upper limit of I/O operations per second (IOPS) allowed for each virtual machine. When storage I/O congestion is detected for a datastore, the I/O workloads of the virtual machines accessing that datastore are adjusted according to the proportion of virtual machine shares each virtual machine has.

Storage I/O shares are similar to shares used for memory and CPU resource allocation, which are described in [Resource Allocation Shares](#). These shares represent the relative importance of a virtual machine regarding the distribution of storage I/O resources. Under resource contention, virtual machines with higher share values have greater access to the storage array. When you allocate storage I/O resources, you can limit the IOPS allowed for a virtual machine. By default, IOPS are unlimited.

The benefits and drawbacks of setting resource limits are described in [Resource Allocation Limit](#). If the limit you want to set for a virtual machine is in terms of MB per second instead of IOPS, you can convert MB per second into IOPS based on the typical I/O size for that virtual machine. For example, to restrict a back up application with 64 KB IOs to 10 MB per second, set a limit of 160 IOPS.

View Storage I/O Control Shares and Limits

You can view the shares and limits for all virtual machines running on a datastore. Viewing this information allows you to compare the settings of all virtual machines that are accessing the datastore, regardless of the cluster in which they are running.

1. Browse to the datastore in the vSphere Client.
2. Click the **VMs** tab.

The tab displays each virtual machine running on the datastore and the associated shares value, and percentage of datastore shares.

Monitor Storage I/O Control Shares

Use the datastore **Performance** tab to monitor how Storage I/O Control handles the I/O workloads of the virtual machines accessing a datastore based on their shares.

Datastore performance charts allow you to monitor the following information:

- Average latency and aggregated IOPS on the datastore
- Latency among hosts
- Queue depth among hosts
- Read/write IOPS among hosts
- Read/write latency among virtual machine disks

- Read/write IOPS among virtual machine disks

1. Browse to the datastore in the vSphere Client.
2. Under the **Monitor** tab, click **Performance**.
3. Select **Advanced**.

Set Storage I/O Control Resource Shares and Limits

Allocate storage I/O resources to virtual machines based on importance by assigning a relative amount of shares to the virtual machine.

See *vSphere Storage* for information on creating VM storage policies and defining common rules for VM storage policies.

Unless virtual machine workloads are very similar, shares do not necessarily dictate allocation in terms of I/O operations or megabytes per second. Higher shares allow a virtual machine to keep more concurrent I/O operations pending at the storage device or datastore compared to a virtual machine with lower shares. Two virtual machines might experience different throughput based on their workloads.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and click **Edit Settings**.
3. Click the **Virtual Hardware** tab and select a virtual hard disk from the list. Expand **Hard disk**.
4. Select a **VM storage policy** from the drop-down menu.
If you select a storage policy, do not manually configure **Shares** and **Limit - IOPS**.
5. Under **Shares**, click the drop-down menu and select the relative amount of shares to allocate to the virtual machine (Low, Normal, or High).
You can select **Custom** to enter a user-defined shares value.
6. Under **Limit - IOPS**, click the drop-down menu and enter the upper limit of storage resources to allocate to the virtual machine.
IOPS are the number of I/O operations per second. By default, IOPS are unlimited. You select Low (500), Normal (1000), or High (2000), or you can select Custom to enter a user-defined number of shares.
7. Click **OK**.

Enable Storage I/O Control

When you enable Storage I/O Control, ESXi monitors datastore latency and throttles the I/O load if the datastore average latency exceeds the threshold.

1. Browse to the datastore in the vSphere Client.
2. Click the **Configure** tab.
3. Click **Settings** and click **General**.
4. Click **Edit** for **Datastore Capabilities**.
5. Select the **Enable Storage I/O Control** check box.
6. Click **OK**.

Under **Datastore Capabilities**, Storage I/O Control is enabled for the datastore.

Set Storage I/O Control Threshold Value

The congestion threshold value for a datastore is the upper limit of latency that is allowed for a datastore before Storage I/O Control begins to assign importance to the virtual machine workloads according to their shares.

Verify that Storage I/O Control is enabled.

You do not need to adjust the threshold setting in most environments.

CAUTION: Storage I/O Control might not function correctly if you share the same spindles on two different datastores.

If you change the congestion threshold setting, set the value based on the following considerations.

- A higher value typically results in higher aggregate throughput and weaker isolation. Throttling will not occur unless the overall average latency is higher than the threshold.
- If throughput is more critical than latency, do not set the value too low. For example, for Fibre Channel disks, a value below 20ms could lower peak disk throughput. A very high value (above 50ms) might allow very high latency without any significant gain in overall throughput.
- A lower value will result in lower device latency and stronger virtual machine I/O performance isolation. Stronger isolation means that the shares controls are enforced more often. Lower device latency translates into lower I/O latency for the virtual machines with the highest shares, at the cost of higher I/O latency experienced by the virtual machines with fewer shares.
- A very low value (lower than 20ms) will result in lower device latency and isolation among I/Os at the potential cost of a decrease in aggregate datastore throughput.
- Setting the value extremely high or extremely lowly results in poor isolation.

1. Browse to the datastore in the vSphere Client.
2. Click the **Configure** tab and click **Settings**.
3. Click **General**.
4. Click **Edit** for **Datastore Capabilities**.

5. Select the **Enable Storage I/O Control** check box.

Storage I/O Control automatically sets the latency threshold that corresponds to the estimated latency when the datastore is operating at 90% of its peak throughput.

6. Optional: Adjust the **Congestion Threshold**.

- Select a value from the **Percentage of peak throughput** drop-down menu.

The percentage of peak throughput value indicates the estimated latency threshold when the datastore is using that percentage of its estimated peak throughput.

- Select a value from the **Manual** drop-down menu.

The value must be between 5ms and 100ms. Setting improper congestion threshold values can be detrimental to the performance of the virtual machines on the datastore.

7. Optional: Click **Reset to defaults** to restore the congestion threshold setting to the default value (30ms).
8. Click **OK**.

Storage DRS Integration with Storage Profiles

Storage Policy Based Management (SPBM) allows you to specify the policy for a virtual machine which is enforced by Storage DRS. A datastore cluster can have set of datastores with different capability profiles. If the virtual machines have storage profiles associated with them, Storage DRS can enforce placement based on underlying datastore capabilities.

By default, Storage DRS will not enforce storage policies associated with a virtual machine. Please configure `EnforceStorageProfiles` option according to your requirements. The options are Default (0), Soft (1) or Hard (2).

As part of Storage DRS integration with storage profiles, the Storage DRS cluster level advanced option `EnforceStorageProfiles` is introduced. Advanced option `EnforceStorageProfiles` takes one of these integer values: 0, 1 or 2. The default value is 0. When the option is set to 0, it indicates that there is no storage profile or policy enforcement on the Storage DRS cluster. When the option is set to 1, it indicates that there is a storage profile or policy soft enforcement on the Storage DRS cluster. This is analogous with DRS soft rules. Storage DRS will comply with storage profile or policy in the optimum level. Storage DRS will violate the storage profile compliant if it is required to do so. Storage DRS affinity rules will have higher precedence over storage profiles only when storage profile enforcement is set to 1. When the option is set to 2, it indicates that there is a storage profile or policy hard enforcement on the Storage DRS cluster. This is analogous with DRS hard rules. Storage DRS will not violate the storage profile or policy compliant. Storage profiles will have higher precedence over affinity rules. Storage DRS will generate fault: `could not fix anti-affinity rule violation`

1. Log in to the vSphere Client as an Administrator.
2. In the vSphere Client, click on the Storage DRS cluster, then select **Manage > Settings > Storage DRS**.
3. Click **Edit > Advanced Options**.
4. Click in the area under the Option heading and type `EnforceStorageProfiles`
5. Click in the area under the Value heading to the right of the previously entered advanced option name and type the value of either 0, 1 or 2.
6. Select **Add**.
7. Click **OK**.

Managing Resource Pools with vSphere


A resource pool is a logical abstraction for flexible management of resources. Resource pools can be grouped into hierarchies and used to hierarchically partition available CPU and memory resources.

Each standalone host and each DRS cluster has an (invisible) root resource pool that groups the resources of that host or cluster. The root resource pool does not appear because the resources of the host (or cluster) and the root resource pool are always the same.

Users can create child resource pools of the root resource pool or of any user-created child resource pool. Each child resource pool owns some of the parent's resources and can, in turn, have a hierarchy of child resource pools to represent successively smaller units of computational capability.

A resource pool can contain child resource pools, virtual machines, or both. You can create a hierarchy of shared resources. The resource pools at a higher level are called parent resource pools. Resource pools and virtual machines that are at the same level are called siblings. The cluster itself represents the root resource pool. If you do not create child resource pools, only the root resource pools exist.

In the following example, RP-QA is the parent resource pool for RP-QA-UI. RP-Marketing and RP-QA are siblings. The three virtual machines immediately below RP-Marketing are also siblings.

Figure 129: Parents, Children, and Siblings in Resource Pool Hierarchy

For each resource pool, you specify reservation, limit, shares, and whether the reservation should be expandable. The resource pool resources are then available to child resource pools and virtual machines.

Note: In this chapter, "Memory" refers to physical RAM.

Why Use Resource Pools?


Resource pools allow you to delegate control over resources of a host (or a cluster), but the benefits are evident when you use resource pools to compartmentalize all resources in a cluster. Create multiple resource pools as direct children of the host or cluster and configure them. You can then delegate control over the resource pools to other individuals or organizations.

Using resource pools can result in the following benefits.

- Flexible hierarchical organization—Add, remove, or reorganize resource pools or change resource allocations as needed.
- Isolation between pools, sharing within pools—Top-level administrators can make a pool of resources available to a department-level administrator. Allocation changes that are internal to one departmental resource pool do not unfairly affect other unrelated resource pools.
- Access control and delegation—When a top-level administrator makes a resource pool available to a department-level administrator, that administrator can then perform all virtual machine creation and management within the boundaries of the resources to which the resource pool is entitled by the current shares, reservation, and limit settings. Delegation is usually done in conjunction with permissions settings.
- Separation of resources from hardware—If you are using clusters enabled for DRS, the resources of all hosts are always assigned to the cluster. That means administrators can perform resource management independently of the actual hosts that contribute to the resources. If you replace three 2GB hosts with two 3GB hosts, you do not need to make changes to your resource allocations.
This separation allows administrators to think more about aggregate computing capacity and less about individual hosts.
- Management of sets of virtual machines running a multitier service—Group virtual machines for a multitier service in a resource pool. You do not need to set resources on each virtual machine. Instead, you can control the aggregate allocation of resources to the set of virtual machines by changing settings on their enclosing resource pool.

For example, assume a host has a number of virtual machines. The marketing department uses three of the virtual machines and the QA department uses two virtual machines. Because the QA department needs larger amounts of CPU and memory, the administrator creates one resource pool for each group. The administrator sets **CPU Shares to High** for the QA department pool and to **Normal** for the Marketing department pool so that the QA department users can run automated tests. The second resource pool with fewer CPU and memory resources is sufficient for the lighter load of the marketing staff. Whenever the QA department is not fully using its allocation, the marketing department can use the available resources.

The numbers in the following figure show the effective allocations to the resource pools.

Figure 130: Allocating Resources to Resource Pools

Create a Resource Pool

You can create a child resource pool of any ESXi host, resource pool, or DRS cluster.

The vSphere Client is connected to the vCenter Server system.

Note: If a host has been added to a cluster, you cannot create child resource pools of that host. If the cluster is enabled for DRS, you can create child resource pools of the cluster.

When you create a child resource pool, you are prompted for resource pool attribute information. The system uses admission control to make sure you cannot allocate resources that are not available. If you want your shares to scale dynamically when adding or removing VMs, you can select scalable shares.

Note: Shares are scaled at the parent level. All descendant resource pools created from a parent with scalable shares have scalable shares by default.

1. In the vSphere Client, select a parent object for the resource pool (a host, another resource pool, or a DRS cluster).
2. Right-click the object and select **New Resource Pool**.
3. Type a name to identify the resource pool.
4. Select the checkbox if you want to enable scalable shares.
5. Specify how to allocate CPU and memory resources.

The CPU resources for your resource pool are the guaranteed physical resources the host reserves for a resource pool. Normally, you accept the default and let the host handle resource allocation.

Option	Description
Shares	Specify shares for this resource pool with respect to the parent's total resources. Sibling resource pools share resources according to their relative share values bounded by the reservation and limit. • Select Low, Normal, or High to specify share values respectively in a 1:2:4 ratio. • Select Custom to give each virtual machine a specific number of shares, which expresses a proportional weight.
Reservation	Specify a guaranteed CPU or memory allocation for this resource pool. Defaults to 0.

Option	Description
	A nonzero reservation is subtracted from the unreserved resources of the parent (host or resource pool). The resources are considered reserved, regardless of whether virtual machines are associated with the resource pool.
Expandable Reservation	When the check box is selected (default), expandable reservations are considered during admission control. If you power on a virtual machine in this resource pool, and the combined reservations of the virtual machines are larger than the reservation of the resource pool, the resource pool can use resources from its parent or ancestors.
Limit	Specify the upper limit for this resource pool's CPU or memory allocation. You can usually accept the default (Unlimited). To specify a limit, deselect the Unlimited check box.

6. Click **OK**.

After you create a resource pool, you can add virtual machines to it. A virtual machine's shares are relative to other virtual machines (or resource pools) with the same parent resource pool.

Creating Resource Pools

Assume that you have a host that provides 6GHz of CPU and 3GB of memory that must be shared between your marketing and QA departments. You also want to share the resources unevenly, giving one department (QA) a higher priority. This can be accomplished by creating a resource pool for each department and using the **Shares** attribute to prioritize the allocation of resources.

The example shows how to create a resource pool with the ESXi host as the parent resource.

1. In the **New Resource Pool** dialog box, type a name for the QA department's resource pool (for example, RP-QA).
2. Specify **Shares** of **High** for the CPU and memory resources of RP-QA.
3. Create a second resource pool, RP-Marketing.
Leave Shares at **Normal** for CPU and memory.
4. Click **OK**.

If there is resource contention, RP-QA receives 4GHz and 2GB of memory, and RP-Marketing 2GHz and 1GB. Otherwise, they can receive more than this allotment. Those resources are then available to the virtual machines in the respective resource pools.

Edit a Resource Pool

After you create the resource pool, you can edit its CPU and memory resource settings.

1. Browse to the resource pool in the vSphere Client.
2. Select **Edit Resource Settings** from the **Actions** drop-down menu.
3. Optional: You can change all attributes of the selected resource pool as described in [Create a Resource Pool](#).
 - Select the checkbox if you want to enable scalable shares.
Note: Shares are scaled at the parent level. All descendant resource pools created from a parent with scalable shares have scalable shares by default.
 - Under **CPU**, select CPU resource settings.
 - Under **Memory**, select memory resource settings.

4. Click **OK** to save your changes.

Add a Virtual Machine to a Resource Pool

When you create a virtual machine, you can specify a resource pool location as part of the creation process. You can also add an existing virtual machine to a resource pool.

When you move a virtual machine to a new resource pool:

- The virtual machine's reservation and limit do not change.
- If the virtual machine's shares are high, medium, or low, %Shares adjusts to reflect the total number of shares in use in the new resource pool.
- If the virtual machine has custom shares assigned, the share value is maintained.

Note: Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares when you move it into a resource pool so that the virtual machine's shares are consistent with the relative values in the new resource pool. A warning appears if a virtual machine would receive a very large (or very small) percentage of total shares.

- Under **Monitor**, the information displayed in the **Resource Reservation** tab about the resource pool's reserved and unreserved CPU and memory resources changes to reflect the reservations associated with the virtual machine (if any).

Note: If a virtual machine has been powered off or suspended, it can be moved but overall available resources (such as reserved and unreserved CPU and memory) for the resource pool are not affected.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and click **Migrate**.
 - You can move the virtual machine to another host.
 - You can move the virtual machine's storage to another datastore.
 - You can move the virtual machine to another host and move its storage to another datastore.
3. Select a resource pool in which to run the virtual machine.
4. Review your selections and click **Finish**.

If a virtual machine is powered on, and the destination resource pool does not have enough CPU or memory to guarantee the virtual machine's reservation, the move fails because admission control does not allow it. An error dialog box displays available and requested resources, so you can consider whether an adjustment might resolve the issue.

Remove a Virtual Machine from a Resource Pool

You can remove a virtual machine from a resource pool either by moving the virtual machine to another resource pool or deleting it.

When you remove a virtual machine from a resource pool, the total number of shares associated with the resource pool decreases, so that each remaining share represents more resources. For example, assume you have a pool that is entitled to 6GHz, containing three virtual machines with shares set to **Normal**. Assuming the virtual machines are CPU-

bound, each gets an equal allocation of 2GHz. If one of the virtual machines is moved to a different resource pool, the two remaining virtual machines each receive an equal allocation of 3GHz.

1. Browse to the resource pool in the vSphere Client.
2. Choose one of the following methods to remove the virtual machine from a resource pool.
 - Right-click the virtual machine and select **Move To...** to move the virtual machine to another resource pool.
You do not need to power off the virtual machine before you move it.
 - Right-click the virtual machine and select **Delete from Disk**.
You must power off the virtual machine before you can completely remove it.

Remove a Resource Pool

You can remove a resource pool from the inventory.

1. In the vSphere Client, right-click the resource pool and Select **Delete**.
A confirmation dialog box appears.
2. Click **Yes** to remove the resource pool.

Resource Pool Admission Control

When you power on a virtual machine in a resource pool, or try to create a child resource pool, the system performs additional admission control to ensure the resource pool's restrictions are not violated.

Before you power on a virtual machine or create a resource pool, ensure that sufficient resources are available using the **Resource Reservation** tab in the vSphere Client. The **Available Reservation** value for CPU and memory displays resources that are unreserved.

How available CPU and memory resources are computed and whether actions are performed depends on the **Reservation Type**.

Table 358: Reservation Types

Reservation Type	Description
Fixed	The system checks whether the selected resource pool has sufficient unreserved resources. If it does, the action can be performed. If it does not, a message appears and the action cannot be performed.
Expandable (default)	The system considers the resources available in the selected resource pool and its direct parent resource pool. If the parent resource pool also has the Expandable Reservation option selected, it can borrow resources from its parent resource pool. Borrowing resources occurs recursively from the ancestors of the current resource pool as long as the Expandable Reservation option is selected. Leaving this option selected offers more flexibility, but, at the same time provides less protection. A child resource pool owner might reserve more resources than you anticipate.

The system does not allow you to violate preconfigured **Reservation** or **Limit** settings. Each time you reconfigure a resource pool or power on a virtual machine, the system validates all parameters so all service-level guarantees can still be met.

Expandable Reservations Example 1

This example shows you how a resource pool with expandable reservations works.

Assume an administrator manages pool P, and defines two child resource pools, S1 and S2, for two different users (or groups).

The administrator knows that users want to power on virtual machines with reservations, but does not know how much each user will need to reserve. Making the reservations for S1 and S2 expandable allows the administrator to more flexibly share and inherit the common reservation for pool P.

Without expandable reservations, the administrator needs to explicitly allocate S1 and S2 a specific amount. Such specific allocations can be inflexible, especially in deep resource pool hierarchies and can complicate setting reservations in the resource pool hierarchy.

Expandable reservations cause a loss of strict isolation. S1 can start using all of P's reservation, so that no memory or CPU is directly available to S2.


Expandable Reservations Example 2

This example shows how a resource pool with expandable reservations works.

Assume the following scenario, as shown in the figure.


- Parent pool RP-MOM has a reservation of 6GHz and one running virtual machine VM-M1 that reserves 1GHz.
- You create a child resource pool RP-KID with a reservation of 2GHz and with **Expandable Reservation** selected.
- You add two virtual machines, VM-K1 and VM-K2, with reservations of 2GHz each to the child resource pool and try to power them on.
- VM-K1 can reserve the resources directly from RP-KID (which has 2GHz).
- No local resources are available for VM-K2, so it borrows resources from the parent resource pool, RP-MOM. RP-MOM has 6GHz minus 1GHz (reserved by the virtual machine) minus 2GHz (reserved by RP-KID), which leaves 3GHz unreserved. With 3GHz available, you can power on the 2GHz virtual machine.

Figure 131: Admission Control with Expandable Resource Pools: Successful Power-On



Now, consider another scenario with VM-M1 and VM-M2.

- Power on two virtual machines in RP-MOM with a total reservation of 3GHz.
- You can still power on VM-K1 in RP-KID because 2GHz are available locally.
- When you try to power on VM-K2, RP-KID has no unreserved CPU capacity so it checks its parent. RP-MOM has only 1GHz of unreserved capacity available (5GHz of RP-MOM are already in use—3GHz reserved by the local virtual machines and 2GHz reserved by RP-KID). As a result, you cannot power on VM-K2, which requires a 2GHz reservation.

Figure 132: Admission Control with Expandable Resource Pools: Power-On Prevented

vSphere Cluster Services

vSphere Cluster Services (vCLS) is activated by default and runs in all vSphere clusters. vCLS ensures that if vCenter Server becomes unavailable, cluster services remain available to maintain the resources and health of the workloads that run in the clusters. vCenter Server is still required to run DRS and HA.

vCLS is activated when you upgrade to vSphere 7.0 Update 3 or when you have a new vSphere 7.0 Update 3 or later deployment. vCLS is upgraded as part of vCenter Server upgrade.

vCLS uses agent virtual machines to maintain cluster services health. The vCLS agent virtual machines (vCLS VMs) are created when you add hosts to clusters. Up to three vCLS VMs are required to run in each vSphere cluster, distributed within a cluster. vCLS is also activated on clusters which contain only one or two hosts. In these clusters the number of vCLS VMs is one and two, respectively.

New anti-affinity rules are applied automatically. Every three minutes a check is performed, if multiple vCLS VMs are located on a single host they will be automatically redistributed to different hosts.

Table 359: Number of vCLS Agent VMs in Clusters

Number of Hosts in a Cluster	Number of vCLS Agent VMs
1	1
2	2
3 or more	3

vCLS VMs run in every cluster even if cluster services like vSphere DRS or vSphere HA are not activated on the cluster. The life cycle operations of vCLS VMs are managed by vCenter Server services like ESX Agent Manager and Workload Control Plane. vCLS VMs do not support NICs.

A cluster activated with vCLS can contain ESXi hosts of different versions if the ESXi versions are compatible with vCenter Server. vCLS works with vSphere Lifecycle Manager clusters.

Embedded vCLS

In vSphere 8.0 U3 Embedded vCLS is introduced, which has new features and improvements.

Differences between External vCLS and Embedded vCLS

In External vCLS the term `quorum` signifies the minimum number of powered-on vCLS VMs necessary for DRS to function, which is 1. In Embedded vCLS we use the term `minimum vCLS VM count`, its purpose and value of 1 are unchanged. In order to improve availability of DRS, both types of vCLS provide redundancy for VMs. If a host with one vCLS VM fails, there are other vCLS VMs available to preserve the functionality of DRS. The total number of vCLS VMs deployed in a cluster, assuming there are enough available hosts and resources to deploy redundant vCLS VMs, is called the desired redundancy count. For External vCLS the desired redundancy count is 3. For Embedded vCLS it is reduced to 2.

What is a vCLS VM?

The vCLS VM introduced in vSphere 8.0 Update 3 is based on vSphere Pod technology. The vCLS service is deployed on the host or cluster as a virtual machine on an ESXi host. However, this is a special VM, with a very minimal operating system and uses a container run-time, which is what allows the vCLS VM to run on vSphere. The new vCLS VMs deliver a secure and high-performance run-times.

vCLS VM power off

It is possible to power off Embedded vCLS VMs. When such VM is powered off, vSphere interprets this condition as a failure and proceeds to restart the VM and re-register vCLS VM with hostd.

When the vCLS VM is powered off, it waits for vSphere to restart that VM. If vSphere doesn't restart that VM after timeout, vCLS VM reconfiguration is initiated.

Whenever an Embedded vCLS VM is evacuated by the system, such as for a host that is entering Maintenance Mode, it is destroyed and replaced on another host, instead of using vMotion.

Note: This behavior is different from External vCLS behavior where vCLS detects VM power-off and immediately proceeds to power on that VM.

vCLS VMs

Visibility and access to vCLS VMs in Embedded vCLS has parity with External vCLS.

- vCLS VMs are visible in VC.
- In VC inventory hierarchy vCLS VMs reside in dedicated VM folder named **vCLS**.
- Some operations on vCLS VMs are unavailable.
- Query requests to vCLS VMs are available.
- The naming of an Embedded vCLS VM is the same as with External vCLS VMs: `vCLS-{UUID}` where UUID is UUID of the ESX host (`summary.hardware.uuid`).

VM ExtraConfig

To identify vCLS VMs, Embedded vCLS continues to use a specific option in VM ExtraConfig.

The extra config option identifying Embedded vCLS VM is `vCLSCRX.agent` and the value is `true`.

Note: This extra config option set in addition to the External vCLS VM extra config option, which is `HDCS.agent`. Embedded and External vCLS VMs are similar enough that they should both be detected for exclusion from most operations using the presence of `HDCS.agent`, but different enough that some operations still need to differentiate them using the presence of `vCLSCRX.agent`.

Embedded vCLS changes the extension that manages these VMs to be the internal/system `vpzd` extension.

- The property `managedBy` in VM configuration of Embedded vCLS VM is set to the key of that VC extension. Having this property helps to identify Embedded vCLS VMs and differentiate them from External vCLS VMs, because this property appears in vSphere.

- The property `config.managedBy.extensionKey` is set to `VirtualCenter`.
- The property `config.managedBy.type` is set to `vcls-entity`.

Disconnected vCLS VMs

A vCLS VM may become disconnected in VC. There are various scenarios where this can happen.

- Have ESX host running vCLS VM
- Simulate connection problem by stopping `vpxa` on the host; connection state transitions to `notResponding`.
- Stop vCLS VM on the host.
- Restart `vpxa` on the host; connection state transitions to `connected`.

After the above steps, the vCLS VM is in VC inventory but not on the ESX host. Hence this VM is marked as `orphaned`. `HdcsManager` is responsible for detecting and deleting disconnected vCLS VMs. In External vCLS, vCLS detects and deletes disconnected vCLS VMs.

Cluster virtual property `vclsVmType`

vSphere provides information about whether cluster is running vCLS VMs managed by External vCLS or Embedded vCLS managed by vSphere. The value of this property is `Embedded` for Embedded vCLS VM cluster, and `external` for External vCLS VM cluster.

Cleaning up vCLS VM

If a cluster ESX host running vCLS VMs loses connectivity to VC and is removed from the VC inventory, the vCLS VM continues to run on that ESX host. In other words, an unmanaged ESX host may end up running vCLS VM.

When an unmanaged ESX host with running vCLS VMs is added as standalone host, `vpxd` detects unexpected presence of vCLS VM and stops it.

When an unmanaged ESX host with running vCLS VM is added to VC cluster, vCLS does not stop vCLS VM prior to adding the host to the cluster. Adding the host triggers vCLS reconfiguration workflow. That workflow stops extraneous vCLS VMs.

Deactivate vCLS

Note: In External vCLS the facility to deactivate vCLS (globally or per cluster) is called `Retreat Mode`.

External vCLS implements facility that allows to deactivate vCLS VM deployment on per-cluster basis. In a sense, this facility is a stop switch to deactivate External vCLS functionality in case of unforeseen circumstances. The several sample scenarios where this functionality can be useful:

- Cluster has both DRS and HA deactivated. In this case, vCLS VMs provide no use, and you do not want the presence of system VMs in the VC inventory.
- Cluster has DRS deactivated and HA activated. In this case vCLS VM do provide some help for optimally failing over VMs, but you might not want the presence of system VMs in VC inventory for marginal benefit.
- Cluster has DRS activated. In this case you may want to temporarily deactivate External vCLS deployment to resolve some transient configuration or transient run-time issues. For example, you might want to deactivate External vCLS until VSAN datastore becomes available, in order to ensure that vCLS VMs are deployed on that VSAN datastore.

Embedded vCLS preserves the same functionality to activate or deactivate vCLS for a cluster.

vSphere DRS and vCLS VMs

vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere cluster. DRS depends on the availability of vCLS VMs.

Note: If you try to activate DRS on a cluster where there are issues with the vCLS VMs, a warning message is displayed on the **Cluster Summary** page.

Note: If DRS is on but there are issues with the vCLS VMs, you must resolve these issues for DRS to operate. A warning message is displayed on the **Cluster Summary** page.

If DRS is non-functional this does not mean that DRS is deactivated. vCLS health turns **Unhealthy** only in a DRS activated cluster when vCLS VMs are not running and the first instance of DRS is skipped because of this. vCLS health will stay **Degraded** on a non-DRS activated cluster when at least one vCLS VM is not running. vCenter has measures to avoid entering a degraded state by accident. For example, cluster Maintenance Mode recommendations do not include options that put all vCLS-capable hosts into Maintenance Mode at once.

Datastore selection for External vCLS

The datastore for vCLS VMs is automatically selected based on ranking all the datastores connected to the hosts inside the cluster.

A datastore is more likely to be selected if there are hosts in the cluster with free reserved DRS slots connected to the datastore. The algorithm tries to place vCLS VMs in a shared datastore if possible before selecting a local datastore. A datastore with more free space is preferred and the algorithm tries not to place more than one vCLS VM on the same datastore. You can only change the datastore of vCLS VMs after they are deployed and powered on.

If you want to move the VMDKs for vCLS VMs to a different datastore or attach a different storage policy, you can reconfigure vCLS VMs. A warning message is displayed when you perform this operation.

You can perform a storage vMotion to migrate vCLS VMs to a different datastore. You can tag vCLS VMs or attach custom attributes if you want to group them separately from workload VMs, for instance if you have a specific meta-data strategy for all VMs that run in a data center.

Note: When a datastore is placed in maintenance mode, if the datastore hosts vCLS VMs, you must manually apply storage vMotion to the vCLS VMs to move them to a new location or put the cluster in retreat mode. A warning message is displayed.

The enter maintenance mode task will start but cannot finish because there is 1 virtual machine residing on the datastore. You can always cancel the task in your Recent Tasks if you decide to continue.

The selected datastore might be storing vSphere Cluster Services VMs which cannot be powered off. To ensure the health of vSphere Cluster Services, these VMs have to be manually vMotioned to a different datastore within the cluster prior to taking this datastore down for maintenance. Refer to this KB article: KB 79892.

Select the check box **Let me migrate storage for all virtual machines and continue entering maintenance mode after migration.** to proceed.

Note: Embedded vCLS does not utilize datastores and ignores the placement settings.

External vCLS Datastore Placement

You can override default vCLS VM datastore placement.

vSphere Cluster Services (vCLS) VM datastore location is chosen by a default datastore selection logic. To override the default vCLS VM datastore placement for a cluster, you can specify a set of allowed datastores by browsing to the cluster and clicking **ADD** under **Configure > vSphere Cluster Service > Datastores**. Some datastores cannot be selected for vCLS because they are blocked by solutions like SRM or vSAN maintenance mode where vCLS cannot be configured. Users cannot add or remove Solution blocked datastores for vCLS VMs.

Monitoring vSphere Cluster Services

You can monitor the resources consumed by vCLS VMs and their health status.

vCLS VMs are not displayed in the inventory tree in the **Hosts and Clusters** tab. vCLS VMs from all clusters within a data center are placed inside a separate VMs and templates folder named **vCLS**. This folder and the vCLS VMs are visible only in the **VMs and Templates** tab of the vSphere Client. These VMs are identified by a different icon than regular workload VMs. You can view information about the purpose of the vCLS VMs in the **Summary** tab of the vCLS VMs.

You can monitor the resources consumed by vCLS VMs in the **Monitor** tab.

Table 360: vCLS VM Resource Allocation

Property	Size
VMDK size	245 MB (thin disk)
Memory	128 MB
CPU	1 vCPU
Hard disk	2 GB
Storage on datastore	480 MB (thin disk)

Note: Each vCLS VM has 100MHz and 100MB capacity reserved in the cluster. Depending on the number of vCLS VMs running in the cluster, a max of 400 MHz and 400 MB of capacity can be reserved for these VMs.

You can monitor the health status of vCLS in the **Cluster Services** portlet displayed in the **Summary** tab of the cluster.

Table 361: Health status of vCLS

Status	Color Coding	Summary
Healthy	Green	If there is at least one vCLS VM running, the status remains healthy, regardless of the number of hosts in the cluster.
Degraded	Yellow	If there is no vCLS VM running for less than 3 minutes (180 seconds), the status is degraded.
Unhealthy	Red	If there is no vCLS VM running for 3 minutes or more, the status is unhealthy in a DRS enabled cluster.

Maintaining Health of vSphere Cluster Services

vCLS VMs are always powered-on because vSphere DRS depends on the availability of these VMs. These VMs should be treated as system VMs. Only administrators can perform selective operations on vCLS VMs. To avoid failure of cluster services, avoid performing any configuration or operations on the vCLS VMs.

vCLS VMs are protected from accidental deletion. Cluster VMs and folders are protected from modification by users, including administrators.

Only users which are part of the Administrators SSO group can perform the following operations:

- ReadOnly access for vCLS VMs
- Console access to vCLS VMs
- Relocate vCLS VMs to either new storage, compute resource or both using cold or hot migration
- Use tags and custom attributes for vCLS VMs

Operations that might disrupt the healthy functioning of vCLS VMs:

- Changing the power state of the vCLS VMs
- Resource reconfiguration of the vCLS VMs such as changing CPU, Memory, Disk size, Disk placement
- VM encryption
- Triggering vMotion of the vCLS VMs
- Changing the BIOS
- Removing the vCLS VMs from the inventory
- Deleting the vCLS VMs from disk
- Enabling FT of vCLS VMs
- Cloning vCLS VMs
- Configuring PMem
- Moving vCLS VM to a different folder
- Renaming the vCLS VMs
- Renaming the vCLS folders
- Enabling DRS rules and overrides on vCLS VMs
- Enabling HA admission control policy on vCLS VMs
- Enabling HA overrides on vCLS VMs
- Moving vCLS VMs to a resource pool
- Recovering vCLS VMs from a snapshot

When you perform any disruptive operation on the vCLS VMs, a warning dialog box appears.

Troubleshooting:

The health of vCLS VMs, including power state, is managed by VMware ESX Agent Manager and Workload Control Plane services. In case of power on failure of vCLS VMs, or if the first instance of DRS for a cluster is skipped due to lack of quorum of vCLS VMs, a banner appears in the cluster summary page along with a link to a Knowledge Base article to help troubleshoot the error state.

Because vCLS VMs are treated as system VMs, you do not need to backup or snapshot these VMs. The health state of these VMs is managed by vCenter Server services.

Putting a Cluster in Retreat Mode

When a datastore is placed in maintenance mode, if the datastore hosts vCLS VMs, you must manually storage vMotion the vCLS VMs to a new location or put the cluster in retreat mode.

This task explains how to put a cluster in retreat mode.

1. Login to the vSphere Client.
2. Navigate to the cluster on which vCLS must be deactivated.
3. Navigate to the vCenter Server **Configure** tab.
4. Under **Configuration**, select **General**.
5. Select either the default **System Managed** option or **Retreat Mode**, which deactivates vCLS.
6. Click **OK**.

If the cluster has DRS activated, it stops functioning and an additional warning is displayed in the **Cluster Summary**. DRS is not functional, even if it is activated, until vCLS is reconfigured by removing it from Retreat Mode.

vSphere HA does not perform optimal placement during a host failure scenario. HA depends on DRS for placement recommendations. HA will still power on the VMs but these VMs might be powered on in a less optimal host.

Retrieving Password for External vCLS

You can retrieve the password to login to the vCLS VMs.

To ensure cluster services health, avoid accessing the vCLS VMs. This document is intended for explicit diagnostics on vCLS VMs.

1. Use SSH to login to the vCenter Server Appliance.
2. Run the following python script:
`/usr/lib/vmware-wcp/decrypt_clustervm_pw.py`

3. Read the output for the password.

```
pwd-script-output
```

```
Read key from file
```

```
Connected to PSQL
```

```
PWD: (password displayed here)
```

With the retrieved password, you can log into the vCLS VMs.

vCLS VM Anti-Affinity Policies

vSphere supports anti-affinity between vCLS VMs and another group of workload VMs.

Compute policies provide a way to specify how the vSphere Distributed Resource Scheduler (DRS) should place VMs on hosts in a resource pool. Use the vSphere Compute Policies editor to create and delete compute policies. You can create or delete, but not modify, a compute policy. If you delete a category tag used in the definition of the policy, the policy is also deleted. Open the **VM Summary** page in vSphere to view the compute policies that apply to a VM and its compliance status with each policy. You can create a compute policy for a group of workload VMs that is anti-affine to the group of vCLS VMs. A vCLS anti-affinity policy can have a single user visible tag for a group of workload VMs, and the other group of vCLS VMs is internally recognized.

Create or Delete a vCLS VM Anti-Affinity Policy

A vCLS VM anti-affinity policy describes a relationship between a category of VMs and vCLS system VMs.

A vCLS VM anti-affinity policy discourages placement of vCLS VMs and application VMs on the same host. This kind of policy can be useful when you do not want vCLS VMs and virtual machines running critical workload to run on the same host. Some best practices for running critical workloads such as SAP HANA require dedicated hosts. After the policy is created, the placement engine attempts to place vCLS VMs on the hosts where policy VMs are not running.

Enforcement of a vCLS VM anti-affinity policy can be affected in several ways:

- If the policy applies to multiple VMs on different hosts and it is not possible to have enough hosts to distribute vCLS VMs, vCLS VMs are consolidated into the hosts without policy VMs.

- If a provisioning operation specifies a destination host, that specification is always honored even if it violates the policy. DRS will try to move the vCLS VMs to a compliant host in a subsequent remediation cycle.
1. Create a category and tag for each group of VMs that you want to include in a vCLS VM anti-affinity policy.
 2. Tag the VMs that you want to include.
 3. Create a vCLS VM anti-affinity policy.
 - a) From the vSphere, click **Policies and Profiles > Compute Policies**.
 - b) Click **Add** to open the **New Compute Policy Wizard**.
 - c) Fill in the policy **Name** and choose **vCLS VM anti affinity** from the **Policy type** drop-down control.
The policy **Name** must be unique.
 - d) Provide a **Description** of the policy, then use **VM tag** to choose the **Category** and **Tag** to which the policy applies.
Unless you have multiple VM tags associated with a category, the wizard fills in the VM tag after you select the tag **Category**.
 - e) Click **Create** to create the policy.
 4. Optional: To delete a compute policy, open vSphere, click **Policies and Profiles > Compute Policies** to show each policy as a card. Click **DELETE** to delete a policy.

Create a vSphere DRS Cluster

A cluster is a collection of ESXi hosts and associated virtual machines with shared resources and a shared management interface. Before you can obtain the benefits of cluster-level resource management you must create a cluster and activate DRS.

Depending on whether or not Enhanced vMotion Compatibility (EVC) is activated, DRS behaves differently when you use vSphere Fault Tolerance (vSphere FT) virtual machines in your cluster.

Table 362: DRS Behavior with vSphere FT Virtual Machines and EVC

EVC	DRS (Load Balancing)	DRS (Initial Placement)
Enabled	Enabled (Primary and Secondary VMs)	Enabled (Primary and Secondary VMs)
Disabled	Disabled (Primary and Secondary VMs)	Disabled (Primary VMs) Fully Automated (Secondary VMs)

Admission Control and Initial Placement

When you attempt to power on a single virtual machine or a group of virtual machines in a DRS-enabled cluster, vCenter Server performs admission control. It checks that there are enough resources in the cluster to support the virtual machine(s).

If the cluster does not have sufficient resources to power on a single virtual machine, or any of the virtual machines in a group power-on attempt, a message appears. Otherwise, for each virtual machine, DRS generates a recommendation of a host on which to run the virtual machine and takes one of the following actions

- Automatically executes the placement recommendation.
- Displays the placement recommendation, which the user can then choose to accept or override.
Note: No initial placement recommendations are given for virtual machines on standalone hosts or in non-DRS clusters. When powered on, they are placed on the host where they currently reside.
- DRS considers network bandwidth. By calculating host network saturation, DRS is able to make better placement decisions. This can help avoid performance degradation of virtual machines with a more comprehensive understanding of the environment.

Single Virtual Machine Power On

In a DRS cluster, you can power on a single virtual machine and receive initial placement recommendations.

When you power on a single virtual machine, you have two types of initial placement recommendations:

- A single virtual machine is being powered on and no prerequisite steps are needed.
The user is presented with a list of mutually exclusive initial placement recommendations for the virtual machine. You can select only one.
- A single virtual machine is being powered on, but prerequisite actions are required.
These actions include powering on a host in standby mode or the migration of other virtual machines from one host to another. In this case, the recommendations provided have multiple lines, showing each of the prerequisite actions. The user can either accept this entire recommendation or cancel powering on the virtual machine.

Group Power-on

You can attempt to power on multiple virtual machines at the same time (group power-on).

Virtual machines selected for a group power-on attempt do not have to be in the same DRS cluster. They can be selected across clusters but must be within the same data center. It is also possible to include virtual machines located in non-DRS clusters or on standalone hosts. These virtual machines are powered on automatically and not included in any initial placement recommendation.

The initial placement recommendations for group power-on attempts are provided on a per-cluster basis. If all the placement-related actions for a group power-on attempt are in automatic mode, the virtual machines are powered on with no initial placement recommendation given. If placement-related actions for any of the virtual machines are in manual mode, the powering on of all the virtual machines (including the virtual machines that are in automatic mode) is manual. These actions are included in an initial placement recommendation.

For each DRS cluster that the virtual machines being powered on belong to, there is a single recommendation, which contains all the prerequisites (or no recommendation). All such cluster-specific recommendations are presented together under the **Power On Recommendations** tab.

When a nonautomatic group power-on attempt is made, and virtual machines not subject to an initial placement recommendation (that is, the virtual machines on standalone hosts or in non-DRS clusters) are included, vCenter Server attempts to power them on automatically. If these power-ons are successful, they are listed under the **Started Power-Ons** tab. Any virtual machines that fail to power-on are listed under the **Failed Power-Ons** tab.

Group Power-on

The user selects three virtual machines in the same data center for a group power-on attempt. The first two virtual machines (VM1 and VM2) are in the same DRS cluster (Cluster1), while the third virtual machine (VM3) is on a standalone host. VM1 is in automatic mode and VM2 is in manual mode. For this scenario, the user is presented with an initial placement recommendation for Cluster1 (under the **Power On Recommendations** tab) which consists of actions for powering on VM1 and VM2. An attempt is made to power on VM3 automatically and, if successful, it is listed under the **Started Power-Ons** tab. If this attempt fails, it is listed under the **Failed Power-Ons** tab.


Virtual Machine Migration

Although DRS performs initial placements so that load is balanced across the cluster, changes in virtual machine load and resource availability can cause the cluster to become unbalanced. To correct such imbalances, DRS generates migration recommendations.

If DRS is enabled on the cluster, load can be distributed more uniformly to reduce the degree of this imbalance. For example, the three hosts on the left side of the following figure are unbalanced. Assume that Host 1, Host 2, and Host 3 have identical capacity, and all virtual machines have the same configuration and load (which includes reservation, if set).

However, because Host 1 has six virtual machines, its resources might be overused while ample resources are available on Host 2 and Host 3. DRS migrates (or recommends the migration of) virtual machines from Host 1 to Host 2 and Host 3. On the right side of the diagram, the properly load balanced configuration of the hosts that results appears.

Figure 133: Load Balancing



When a cluster becomes unbalanced, DRS makes recommendations or migrates virtual machines, depending on the default automation level:

- If the cluster or any of the virtual machines involved are manual or partially automated, vCenter Server does not take automatic actions to balance resources. Instead, the Summary page indicates that migration recommendations are available and the DRS Recommendations page displays recommendations for changes that make the most efficient use of resources across the cluster.
- If the cluster and virtual machines involved are all fully automated, vCenter Server migrates running virtual machines between hosts as needed to ensure efficient use of cluster resources.

Note: Even in an automatic migration setup, users can explicitly migrate individual virtual machines, but vCenter Server might move those virtual machines to other hosts to optimize cluster resources.

By default, automation level is specified for the whole cluster. You can also specify a custom automation level for individual virtual machines.

DRS Migration Threshold

The DRS migration threshold allows you to specify which recommendations are generated and then applied (when the virtual machines involved in the recommendation are in fully automated mode) or shown (if in manual mode). This threshold is a measure of how aggressive DRS is in recommending migrations to improve VM happiness.

You can move the threshold slider to use one of five settings, ranging from Conservative to Aggressive. The higher the aggressiveness setting, the more frequently DRS might recommend migrations to improve VM happiness. The Conservative setting generates only priority-one recommendations (mandatory recommendations).

After a recommendation receives a priority level, this level is compared to the migration threshold you set. If the priority level is less than or equal to the threshold setting, the recommendation is either applied (if the relevant virtual machines are in fully automated mode) or displayed to the user for confirmation (if in manual or partially automated mode.)

DRS Score

Each migration recommendation is computed using the VM happiness metric which measures execution efficiency. This metric is displayed as DRS Score in the cluster's Summary tab in the vSphere Client. DRS load balancing recommendations attempt to improve the DRS score of a VM. The Cluster DRS score is a weighted average of the VM DRS Scores of all the powered on VMs in the cluster. The Cluster DRS Score is shown in the gauge component. The color of the filled in section changes depending on the value to match the corresponding bar in the VM DRS Score histogram. The bars in the histogram show the percentage of VMs that have a DRS Score in that range. You can view the list with server-side sorting and filtering by selecting the Monitor tab of the cluster and selecting vSphere DRS, which shows a list of the VMs in the cluster sorted by their DRS score in ascending order.

Migration Recommendations

If you create a cluster with a default manual or partially automated mode, vCenter Server displays migration recommendations on the DRS Recommendations page.

The system supplies as many recommendations as necessary to enforce rules and balance the resources of the cluster. Each recommendation includes the virtual machine to be moved, current (source) host and destination host, and a reason for the recommendation. The reason can be one of the following:

- Balance average CPU loads or reservations.
- Balance average memory loads or reservations.
- Satisfy resource pool reservations.
- Satisfy an affinity rule.
- Host is entering maintenance mode or standby mode.

Note: If you are using the vSphere Distributed Power Management (DPM) feature, in addition to migration recommendations, DRS provides host power state recommendations.

DRS Cluster Requirements

Hosts that are added to a DRS cluster must meet certain requirements to use cluster features successfully.

Note: vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere Cluster. Starting with vSphere 7.0 Update 1, DRS depends on the availability of vCLS VMs. See [vSphere Cluster Services](#) for more information.

Shared Storage Requirements

A DRS cluster has certain shared storage requirements.

Ensure that the managed hosts use shared storage. Shared storage is typically on a SAN, but can also be implemented using NAS shared storage.

See the *vSphere Storage* documentation for information about other shared storage.

Shared VMFS Volume Requirements

A DRS cluster has certain shared VMFS volume requirements.

Configure all managed hosts to use shared VMFS volumes.

- Place the disks of all virtual machines on VMFS volumes that are accessible by source and destination hosts.
- Ensure the VMFS volume is sufficiently large to store all virtual disks for your virtual machines.
- Ensure all VMFS volumes on source and destination hosts use volume names, and all virtual machines use those volume names for specifying the virtual disks.

Note: Virtual machine swap files also need to be on a VMFS accessible to source and destination hosts (just like `.vmdk` virtual disk files). This requirement does not apply if all source and destination hosts are ESX Server 3.5 or higher and using host-local swap. In that case, vMotion with swap files on unshared storage is supported. Swap files are placed on a VMFS by default, but administrators might override the file location using advanced virtual machine configuration options.

Processor Compatibility Requirements

A DRS cluster has certain processor compatibility requirements.

To avoid limiting the capabilities of DRS, you should maximize the processor compatibility of source and destination hosts in the cluster.

vMotion transfers the running architectural state of a virtual machine between underlying ESXi hosts. vMotion compatibility means that the processors of the destination host must be able to resume execution using the equivalent instructions where the processors of the source host were suspended. Processor clock speeds and cache sizes might vary, but processors must come from the same vendor class (Intel versus AMD) and the same processor family to be compatible for migration with vMotion.

Processor families are defined by the processor vendors. You can distinguish different processor versions within the same family by comparing the processors' model, stepping level, and extended features.

Sometimes, processor vendors have introduced significant architectural changes within the same processor family (such as 64-bit extensions and SSE3). VMware identifies these exceptions if it cannot guarantee successful migration with vMotion.

vCenter Server provides features that help ensure that virtual machines migrated with vMotion meet processor compatibility requirements. These features include:

- **Enhanced vMotion Compatibility (EVC)** – You can use EVC to help ensure vMotion compatibility for the hosts in a cluster. EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. This prevents migrations with vMotion from failing due to incompatible CPUs. Configure EVC from the Cluster Settings dialog box. The hosts in a cluster must meet certain requirements for the cluster to use EVC. For information about EVC and EVC requirements, see the *vCenter Server and Host Management* documentation.
- **CPU compatibility masks** – vCenter Server compares the CPU features available to a virtual machine with the CPU features of the destination host to determine whether to allow or disallow migrations with vMotion. By applying CPU compatibility masks to individual virtual machines, you can hide certain CPU features from the virtual machine and potentially prevent migrations with vMotion from failing due to incompatible CPUs.

vMotion Requirements for DRS Clusters

A DRS cluster has certain vMotion requirements.

To enable the use of DRS migration recommendations, the hosts in your cluster must be part of a vMotion network. If the hosts are not in the vMotion network, DRS can still make initial placement recommendations.

To be configured for vMotion, each host in the cluster must meet the following requirements:

- vMotion does not support raw disks or migration of applications clustered using Microsoft Cluster Service (MSCS).
- vMotion requires a private Gigabit Ethernet migration network between all of the vMotion enabled managed hosts. When vMotion is enabled on a managed host, configure a unique network identity object for the managed host and connect it to the private migration network.

Configuring DRS with Virtual Flash

DRS can manage virtual machines that have virtual flash reservations.

Virtual flash capacity appears as a statistic that is regularly reported from the host to the vSphere Client. Each time DRS runs, it uses the most recent capacity value reported.

You can configure one virtual flash resource per host. This means that during virtual machine power-on time, DRS does not need to select between different virtual flash resources on a given host.

DRS selects a host that has sufficient available virtual flash capacity to start the virtual machine. If DRS cannot satisfy the virtual flash reservation of a virtual machine, it cannot be powered-on. DRS treats a powered-on virtual machine with a virtual flash reservation as having a soft affinity with its current host. DRS will not recommend such a virtual machine for vMotion except for mandatory reasons, such as putting a host in maintenance mode, or to reduce the load on an over utilized host.

Create a Cluster

A cluster is a group of hosts. When a host is added to a cluster, the host's resources become part of the cluster's resources. The cluster manages the resources of all hosts within it.

- Verify that you have sufficient permissions to create a cluster object.
- Verify that a data center exists in the inventory.
- If you want to use vSAN, it must be enabled before you configure vSphere HA.

Clusters enable the vSphere High Availability (HA) and vSphere Distributed Resource Scheduler (DRS) solutions.

Note: vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere Cluster. Starting with vSphere 7.0 Update 1, DRS depends on the availability of vCLS VMs. See [vSphere Cluster Services](#) for more information.

1. Browse to a data center in the vSphere Client.
2. Right-click the data center and select **New Cluster**.
3. Enter a name for the cluster.
4. Select DRS and vSphere HA cluster features.

Option	Description
To use DRS with this cluster	1. Select the DRS Turn ON check box. 2. Select an automation level and a migration threshold.
To use HA with this cluster	1. Select the vSphere HA Turn ON check box. 2. Select whether to enable host monitoring and admission control. 3. If admission control is enabled, specify a policy. 4. Select a VM Monitoring option. 5. Specify the virtual machine monitoring sensitivity.

5. Select an Enhanced vMotion Compatibility (EVC) setting.
EVC ensures that all hosts in a cluster present the same CPU feature set to virtual machines, even if the actual CPUs on the hosts differ. This prevents migrations with vMotion from failing due to incompatible CPUs.
6. Click **OK**.

The cluster is added to the inventory.

Add hosts and resource pools to the cluster.

Note: Under the **Cluster Summary** page, you can see **Cluster Services** which displays vSphere Cluster Services health status.

Edit Cluster Settings

When you add a host to a DRS cluster, the host's resources become part of the cluster's resources. In addition to this aggregation of resources, with a DRS cluster you can support cluster-wide resource pools and enforce cluster-level resource allocation policies.

You can create a cluster without a special license, but you must have a license to enable a cluster for vSphere DRS or vSphere HA.

Note: vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere Cluster. Starting with vSphere 7.0 Update 1, DRS depends on the availability of vCLS VMs. See [vSphere Cluster Services](#) for more information.

The following cluster-level resource management capabilities are also available.

Load Balancing

The distribution and usage of CPU and memory resources for all hosts and virtual machines in the cluster are continuously monitored. DRS compares these metrics to an ideal resource usage given the attributes of the cluster's resource pools and virtual machines, the current demand, and the imbalance target. DRS then provides recommendations or performs virtual machine migrations accordingly. See [Virtual Machine Migration](#). When you power on a virtual machine in the cluster, DRS attempts to maintain proper load balancing by either placing the virtual machine on an appropriate host or making a recommendation. See [Admission Control and Initial Placement](#).

Power management

When the vSphere Distributed Power Management (DPM) feature is enabled, DRS compares cluster and host-level capacity to the demands of the cluster's virtual machines, including recent historical demand. DRS then recommends you place hosts in standby, or places hosts in standby power mode when sufficient excess capacity is found. DRS powers-on hosts if capacity is needed. Depending on the resulting host power state recommendations, virtual machines might need to be migrated to and from the hosts as well. See [Managing Power Resources](#).

Affinity Rules

You can control the placement of virtual machines on hosts within a cluster, by assigning affinity rules. See [Using Affinity Rules with vSphere DRS](#).

1. Browse to a cluster in the vSphere Client.
2. Click the **Configure** tab and click **Services**.
3. Under **vSphere DRS** click **Edit**.
4. Under **DRS Automation**, select a default automation level for DRS.

Automation Level	Action
Manual	- Initial placement: Recommended host is displayed. - Migration: Recommendation is displayed.
Partially Automated	- Initial placement: Automatic. - Migration: Recommendation is displayed.
Fully Automated	- Initial placement: Automatic. - Migration: Recommendation is run automatically.

5. Set the **Migration Threshold** for DRS.
6. Select the **Predictive DRS** check box. In addition to real-time metrics, DRS responds to forecasted metrics provided by vRealize Operations server. You must also configure **Predictive DRS** in a version of vRealize Operations that supports this feature.
7. Select **Virtual Machine Automation** check box to enable individual virtual machine automation levels. Override for individual virtual machines can be set from the VM Overrides page.
8. Under **Additional Options**, select a check box to enforce one of the default policies.

Option	Description
VM Distribution	For availability, distribute a more even number of virtual machines across hosts. This is secondary to DRS load balancing.
Memory Metric for Load Balancing	Load balance based on consumed memory of virtual machines rather than active memory. This setting is only recommended for clusters where host memory is not over-committed. Note: This setting is no longer supported and will not be displayed in vCenter 7.0.
CPU Over-Commitment	Control CPU over-commitment in the cluster.
Scalable Shares	Enable scalable shares for the resource pools on this cluster.

9. Under **Power Management**, select Automation Level.

10. If DPM is enabled, set the **DPM Threshold**.

11. Click **OK**.

Note: Under the **Cluster Summary** page, you can see **Cluster Services** which displays vSphere Cluster Services health status.

You can view memory utilization for DRS in the vSphere Client. To find out more, see: Viewing Distributed Resource Scheduler Memory Utilization

Set a Custom Automation Level for a Virtual Machine

After you create a DRS cluster, you can customize the automation level for individual virtual machines to override the cluster's default automation level.

For example, you can select **Manual** for specific virtual machines in a cluster with full automation, or **Partially Automated** for specific virtual machines in a manual cluster.

If a virtual machine is set to **Disabled**, vCenter Server does not migrate that virtual machine or provide migration recommendations for it.

1. Browse to the cluster in the vSphere Client.
2. Click the **Configure** tab and click **Services**.
3. Under Services, select **vSphere DRS** and click **Edit**. Expand DRS Automation.
4. Select the **Enable individual virtual machine automation levels** check box.
5. To temporarily deactivate any individual virtual machine overrides, deselect the **Enable individual virtual machine automation levels** check box.
Virtual machine settings are restored when the check box is selected again.

6. To temporarily suspend all vMotion activity in a cluster, put the cluster in manual mode and deselect the **Enable individual virtual machine automation levels** check box.
7. Select one or more virtual machines.
8. Click the **Automation Level** column and select an automation level from the drop-down menu.

Option	Description
Manual	Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.
Fully Automated	Placement and migration recommendations run automatically.
Partially Automated	Initial placement is performed automatically. Migration recommendations are displayed, but do not run.
Disabled	vCenter Server does not migrate the virtual machine or provide migration recommendations for it.

9. Click **OK**.

Note:

Other VMware products or features, such as vSphere vApp and vSphere Fault Tolerance, might override the automation levels of virtual machines in a DRS cluster. Refer to the product-specific documentation for details.

Deactivate DRS

You can turn off DRS for a cluster.

When DRS is deactivated:

- DRS affinity rules are not removed but are not applied until DRS is reactivated.
- Host and VM groups are not removed but are not applied until DRS is reactivated.
- Resource pools are permanently removed from the cluster. To avoid losing the resource pools, save a snapshot of the resource pool tree on your local machine. You can use the snapshot to restore the resource pool when you activate DRS.

Note: Do not deactivate DRS if Workload Management is activated, this makes WCP services unrecoverable.

1. Browse to the cluster in the vSphere Client.
2. Click the **Configure** tab and click **Services**.
3. Under **vSphere DRS**, click **Edit**.
4. Deselect the **Turn On vSphere DRS** check box.
5. Click **OK** to turn off DRS.
6. Optional: Choose an option to save the resource pool.
 - Click **Yes** to save a resource pool tree snapshot on a local machine.
 - Click **No** to turn off DRS without saving a resource pool tree snapshot.

DRS is turned off.

Note: vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere Cluster. Starting with vSphere 7.0 Update 1, DRS depends on the availability of vCLS VMs. See [vSphere Cluster Services](#) for more information.

Restore a Resource Pool Tree

You can restore a previously saved resource pool tree snapshot.

- vSphere DRS must be turned ON.
- You can restore a snapshot only on the same cluster that it was taken.
- No other resource pools are present in the cluster.
- Backup and restore must always be performed on the same version of vCenter and ESXi.

1. Browse to the cluster in the vSphere Client.
2. Select the cluster and click the **Configure** tab.
3. Under **System**, click **vSphere DRS**.
4. Click **Restore Resource Pool Tree**.
5. Click **Browse**, and locate the snapshot file on your local machine.
6. Click **Open**.
7. Click **OK** to restore the resource pool tree.

DRS Awareness of vSAN Stretched Cluster

DRS Awareness of vSAN Stretched Cluster is available on stretched clusters with DRS enabled. A vSAN stretched cluster has read locality, where the VM reads data from a local site. Fetching reads from a remote site can affect VM performance. With DRS Awareness of vSAN Stretched Cluster, DRS is now fully aware of VM read locality and will place the VM on a site that can fully satisfy the read locality. This is automatic, there are no configurable options. DRS Awareness of vSAN Stretched Cluster works with existing affinity rules. It also works with VMware Cloud on AWS.

vSAN Stretched Cluster with vSphere HA and vSphere DRS provide resiliency by having two copies of data spread across two fault domains and a witness node in a third fault domain in case of failures. The two active fault domains provide replication of data so that both fault domains have a current copy of the data.

vSAN Stretched Cluster provides an automated method of moving workloads within the two fault domains. In case of full site failures, VMs are restarted on the secondary site by vSphere HA. This ensures that there is no downtime for critical production workloads. Once the primary site is back online, DRS immediately rebalances the VMs back to the primary site with soft affinity hosts. This process causes the VM to read and write from the secondary site while the VM data components are still rebuilding and might reduce VM performance.

In releases prior to vSphere 7.0 U2, we recommend that you change DRS from fully automated to partially automated mode, to avoid VMs migrating while resynchronization is in progress to the primary site. Set DRS back to fully automated only after the resynchronization is complete.

DRS Awareness of vSAN Stretched Cluster introduces a fully automated read locality solution for recovering from failures on a vSAN stretched cluster. The read locality information indicates the hosts the VM has full access to, and DRS uses this information when placing a VM on a host on vSAN Stretched Clusters. DRS prevents VMs from failing back to the primary site when vSAN resynchronization is still in progress during the site recovery phase. DRS automatically migrates a VM back to the primary affined site when its data components have achieved full read locality. This allows you to operate DRS in fully automatic mode in case of full site failures.

In case of partial site failures, if a VM loses read locality due to loss of data components greater than or equal to its Failures to Tolerate vSphere DRS will identify the VMs that consume a very high read bandwidth and try to rebalance them to the secondary site. This ensures that VMs with read-heavy workloads do not decrease during partial site failures. Once the primary site is back online and the data components have completed resynchronization, the VM is moved back to the site it is affined to.

DRS Placement of vGPUs

DRS distributes vGPU VMs across a cluster's hosts.

DRS will distribute vGPU VMs in a breadth-first manner across a cluster's hosts. Fractional vGPU profile allocation for a VM may be subject to homogeneous profile mutual exclusion rules.

- Manually migrate vGPU VMs to a desired host to open up unused Physical GPU capacity.
- Use the same vGPU profile configuration in all vGPU VMs in a Cluster.
- Enable Host "GPU Consolidation", please see [Configuring Host Graphics](#) for more information.
- If DRS Automation is active, consider putting the Cluster or VM into **Partially Automated** mode, please see [Edit Cluster Settings](#) for more information.

DRS Overhead Memory Management for VMs

In vSphere 8.0 U3 DRS has enhanced overhead memory management for VMs which are being reconfigured.

In VMware vSphere, overhead memory refers to the amount of memory that is used by ESXi to manage a virtual machine (VM). This memory is necessary for ESXi to perform its functions and is separate from the guest memory allocated to the VMs. The amount of overhead memory depends on several factors, including the number of virtual CPUs (vCPUs), the amount of memory allocated to the VM, and the VM's configuration and hardware version. More vCPUs and larger memory allocations lead to higher overhead memory consumption. In vSphere, DRS works with ESXi memory management to ensure the VM has optimal overhead memory usage. DRS manages the overhead memory by setting up the VM overhead memory limit, and the ESXi is allowed to consume the overhead memory within the limit.

Reconfiguring a VM in VMware vSphere can directly impact the overhead memory that ESXi requires to manage a VM. When you change a VM's configuration, such as modifying the number of vCPUs, the amount of assigned RAM, or adding virtual hardware like network adapters or disk controllers, the overhead memory requirements can change. For example, reconfiguring a VM reservation from 250GB to 0GB requires about 25MB of overhead memory usage. ESXi allocates 25MB of additional overhead memory to manage the page table mapping between virtual and physical pages. vSphere monitors and manages these changes. However, previous vSphere releases were not customized to accommodate these overhead memory increases. If the new overhead memory increase exceeds the overhead limit, it can cause a reconfiguration failure.

In vSphere 8.0 U3, DRS proactively updates the VM's overhead memory limit before any reconfiguration. DRS checks various factors, including the VM's resource specifications, IO filters, and other elements impacting overhead memory. DRS ensures that the new overhead limit accommodates the expected increase in overhead memory due to the updated VM specifications after reconfiguration, which optimizes VM performance and stability.

DRS enhanced overhead memory management can help to prevent reconfiguration failures by intelligently managing overhead memory limits before a VM undergoes reconfiguration, which dramatically reduces the risk of reconfiguration failures. This proactive approach ensures a more reliable experience. By optimizing VM performance and stability, virtual environments can run efficiently and without interruption, particularly during critical reconfiguration processes. This enhancement seamlessly integrates into your existing vSphere environment, while increasing performance and reliability.

vSphere DRS Maintenance Mode Functionality with ROBO Enterprise License

VMware's large Remote Office Branch Office (ROBO) Enterprise license supports automatic VM evacuations when a host enters maintenance mode.

In a ROBO Enterprise cluster, DRS is deactivated by default and you cannot make changes to DRS configuration. When a host in a ROBO Enterprise cluster enters maintenance mode, VMs are automatically evacuated from the host by DRS. Before evacuating the VMs from the host, DRS creates VM-host affinity mappings to track where the VMs were placed. When the host exits maintenance mode, the VMs that were running on the host are migrated back to the host. VM-host affinity mappings are cleared after the migration.

Limitations of DRS Maintenance Mode with ROBO Enterprise License

DRS Functionality with ROBO Enterprise License is not full DRS functionality.

There are some limitations you should be aware of before you initiate maintenance mode on a ROBO Enterprise cluster. On a ROBO Enterprise cluster, DRS is deactivated by default. If you have migrated from a DRS supported license to a ROBO Enterprise license, you may have VMs with affinity or anti-affinity rules present in the system. You must deactivate or delete VMs with affinity or anti-affinity rules or the ROBO Enterprise maintenance mode operation is deactivated. ROBO Enterprise maintenance mode operation is deactivated if DRS is not set to fully automated mode. DRS automation level must be set to fully automated mode in order to evacuate VMs automatically through the host maintenance workflow. If a VM overrides DRS fully automated mode, you must evacuate the VM manually.

Using DRS Maintenance Mode with ROBO Enterprise License

vSphere supports limited DRS Maintenance Mode functionality with ROBO Enterprise License.

- Check if all hosts in a cluster have the ROBO Enterprise License installed. If they do not, you must install the license.
 - Check whether any DRS rules are configured and activated. If they are, you must deactivate or delete them to use ROBO Enterprise maintenance mode operation.
1. In order for DRS Maintenance Mode to work with ROBO Enterprise License, make sure that each host on the cluster has the ROBO Enterprise License installed.
 - If the license is not installed, go to step 2.
 - If the license is installed, go to step 3.
 2. Install the ROBO Enterprise License
 - a) Browse to the host in the vSphere Client.
 - b) Under the **Configure** tab, select **Licensing**.
 - c) Click **ASSIGN LICENSE**.
 - d) Enter your ROBO Enterprise License key and click **OK**.

You must repeat these steps for all hosts in the cluster.
 3. Select the host in the cluster, right-click and select **Enter Maintenance Mode** and click **OK**.
The VMs on the host are automatically evacuated.

After the host exits Maintenance Mode, the VMs are automatically migrated back to the host. The host is restored to the original state. However, if a host is overloaded DRS cannot migrate VMs back to the original host. DRS attempts to restore the host to the original state, but it cannot make a host overloaded.

If you need to deactivate DRS Maintenance Mode with ROBO Enterprise License, you can edit the vpxd.cfg file. Open the vpxd.cfg file. Under the `<cluster>` option, change `<roboMMEnabled>true</roboMMEnabled>` to `<roboMMEnabled>>false</roboMMEnabled>`. This is runtime configuration, so you do not need to restart vpxd after updating the configuration.

Troubleshooting DRS Maintenance Mode with ROBO Enterprise License

If you experience issues using maintenance mode with your ROBO Enterprise cluster, consider the following.

In order for maintenance mode to function correctly with a ROBO Enterprise cluster:

- Check if all hosts in a cluster have the ROBO Enterprise License installed. If they do not, you must install the license.
- Check whether any DRS rules are configured and activated. If they are, you must deactivate or delete them to use ROBO Enterprise maintenance mode operation.
- If the compatibility check fails, make sure that the other hosts are compatible with the VM.

Use DRS Clusters to Manage Resources with vSphere

After you create a DRS cluster, you can customize it and use it to manage resources.

To customize your DRS cluster and the resources it contains you can configure affinity rules and you can add and remove hosts and virtual machines. When a cluster's settings and resources have been defined, you should ensure that it is and remains a valid cluster. You can also use a valid DRS cluster to manage power resources and interoperate with vSphere HA.

Note: In this chapter, "Memory" can refer to physical RAM or Persistent Memory.

Adding Hosts to a Cluster

The procedure for adding hosts to a cluster is different for hosts managed by the same vCenter Server (managed hosts) than for hosts not managed by that server.

After a host has been added, the virtual machines deployed to the host become part of the cluster and DRS can recommend migration of some virtual machines to other hosts in the cluster.

Note: vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere Cluster. Starting with vSphere 7.0 Update 1, DRS depends on the availability of vCLS VMs. See [vSphere Cluster Services](#) for more information.

Add a Managed Host to a Cluster

When you add a standalone host already being managed by vCenter Server to a DRS cluster, the host's resources become associated with the cluster.

You can decide whether you want to associate existing virtual machines and resource pools with the cluster's root resource pool or graft the resource pool hierarchy.

Note: If a host has no child resource pools or virtual machines, the host's resources are added to the cluster but no resource pool hierarchy with a top-level resource pool is created.

1. Browse to the host in the vSphere Client.
2. Right-click the host and select **Move To....**
3. Select a cluster.
4. Click **OK** to apply the changes.
5. Select what to do with the host's virtual machines and resource pools.
 - **Put this host's virtual machines in the cluster's root resource pool**
vCenter Server removes all existing resource pools of the host and the virtual machines in the host's hierarchy are all attached to the root. Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares after selecting this option, which destroys the resource pool hierarchy.
 - **Create a resource pool for this host's virtual machines and resource pools**
vCenter Server creates a top-level resource pool that becomes a direct child of the cluster and adds all children of the host to that new resource pool. You can supply a name for that new top-level resource pool. The default is **Grafted from <host_name>**.

The host is added to the cluster.

Add an Unmanaged Host to a Cluster

You can add an unmanaged host to a cluster. Such a host is not currently managed by the same vCenter Server system as the cluster and it is not visible in the vSphere Client.

1. Browse to the cluster in the vSphere Client.
2. Right-click the cluster and select **Add Host**.
3. Enter the host name, user name, and password, and click **Next**.
4. View the summary information and click **Next**.
5. Assign an existing or a new license key and click **Next**.
6. Optional: You can enable lockdown mode to prevent remote users from logging directly into the host.
If you do not enable lockdown mode, you can configure this option later by editing Security Profile in host settings.
7. Select what to do with the host's virtual machines and resource pools.
 - **Put this host's virtual machines in the cluster's root resource pool**
vCenter Server removes all existing resource pools of the host and the virtual machines in the host's hierarchy are all attached to the root. Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares after selecting this option, which destroys the resource pool hierarchy.
 - **Create a resource pool for this host's virtual machines and resource pools**
vCenter Server creates a top-level resource pool that becomes a direct child of the cluster and adds all children of the host to that new resource pool. You can supply a name for that new top-level resource pool. The default is **Grafted from <host_name>**.
8. Review settings and click **Finish**.

The host is added to the cluster.

Adding Virtual Machines to a Cluster

You can add a virtual machine to a cluster in a number of ways.

- When you add a host to a cluster, all virtual machines on that host are added to the cluster.
- When a virtual machine is created, the **New Virtual Machine** wizard prompts you for the location to place the virtual machine. You can select a standalone host or a cluster and you can select any resource pool inside the host or cluster.
- You can migrate a virtual machine from a standalone host to a cluster or from a cluster to another cluster using the **Migrate Virtual Machine** wizard. To start this wizard, right-click the virtual machine name and select **Migrate**.

Move a Virtual Machine to a Cluster

You can move a virtual machine to a cluster.

1. Find the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click on the virtual machine and select **Move To...**
3. Select a cluster.
4. Click **OK**.

Removing Virtual Machines from a Cluster

You can remove virtual machines from a cluster.

You can remove a virtual machine from a cluster in two ways.

- When you remove a host from a cluster, all of the powered-off virtual machines that you do not migrate to other hosts are removed as well. You can remove a host only if it is in maintenance mode or disconnected. If you remove a host from a DRS cluster, the cluster can become yellow because it is overcommitted.
- You can migrate a virtual machine from a cluster to a standalone host or from a cluster to another cluster using the **Migrate** wizard. To start this wizard right-click the virtual machine name and select **Migrate**.

Move a Virtual Machine Out of a Cluster

You can move a virtual machine out of a cluster.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and select **Migrate**.
3. Select **Change datastore** and click **Next**.
4. Select a datastore and click **Next**.
5. Click **Finish**.

If the virtual machine is a member of a DRS cluster rules group, vCenter Server displays a warning before it allows the migration to proceed. The warning indicates that dependent virtual machines are not migrated automatically. You have to acknowledge the warning before migration can proceed.

Removing a Host from a Cluster

When you remove a host from a DRS cluster, you affect resource pool hierarchies, virtual machines, and you might create invalid clusters. Consider the affected objects before you remove the host.

- **Resource Pool Hierarchies** – When you remove a host from a cluster, the host retains only the root resource pool, even if you used a DRS cluster and decided to graft the host resource pool when you added the host to the cluster. In that case, the hierarchy remains with the cluster. You can create a host-specific resource pool hierarchy.
Note: Ensure that you remove the host from the cluster by first placing it in maintenance mode. If you instead disconnect the host before removing it from the cluster, the host retains the resource pool that reflects the cluster hierarchy.
- **Virtual Machines** – A host must be in maintenance mode before you can remove it from the cluster and for a host to enter maintenance mode all powered-on virtual machines must be migrated off that host. When you request that a host

enter maintenance mode, you are also asked whether you want to migrate all the powered-off virtual machines on that host to other hosts in the cluster.

- **Invalid Clusters** – When you remove a host from a cluster, the resources available for the cluster decrease. If the cluster has enough resources to satisfy the reservations of all virtual machines and resource pools in the cluster, the cluster adjusts resource allocation to reflect the reduced amount of resources. If the cluster does not have enough resources to satisfy the reservations of all resource pools, but there are enough resources to satisfy the reservations for all virtual machines, an alarm is issued and the cluster is marked yellow. DRS continues to run.

Place a Host in Maintenance Mode

You place a host in maintenance mode when you need to service it, for example, to install more memory. A host enters or leaves maintenance mode only as the result of a user request.

Virtual machines that are running on a host entering maintenance mode need to be migrated to another host (either manually or automatically by DRS) or shut down. The host is in a state of **Entering Maintenance Mode** until all running virtual machines are powered down or migrated to different hosts. You cannot power on virtual machines or migrate virtual machines to a host entering maintenance mode.

When no more running virtual machines are on the host, the host's icon changes to include **under maintenance** and the host's Summary panel indicates the new state. While in maintenance mode, the host does not allow you to deploy or power on a virtual machine.

Note: DRS does not recommend (or perform, in fully automated mode) any virtual machine migrations off of a host entering maintenance or standby mode if the vSphere HA failover level would be violated after the host enters the requested mode.

1. Browse to the host in the vSphere Client.
2. Right-click the host and select **Maintenance Mode > Enter Maintenance Mode**.
 - If the host is part of a partially automated or manual DRS cluster, browse to **Cluster > Monitor > DRS > Recommendations** and click **Apply Recommendations**.
 - If the host is part of an automated DRS cluster, virtual machines are migrated to different hosts when the host enters maintenance mode.
3. If applicable, click **Yes**.

The host is in maintenance mode until you select **Maintenance Mode > Exit Maintenance Mode**.

Remove a Host from a Cluster

You can remove hosts from a cluster.

1. Browse to the host in the vSphere Client.
2. Right-click the host and select **Maintenance Mode > Enter Maintenance Mode**.
When the host is in maintenance mode, move it to a different inventory location, either the top-level data center or to a different cluster.
3. Right-click the host and select **Move To...**
4. Select a new location for the host and click **OK**.

When you move the host, its resources are removed from the cluster. If you grafted the host's resource pool hierarchy onto the cluster, that hierarchy remains with the cluster.

After you remove a host from a cluster, you can perform the following tasks.

- Remove the host from vCenter Server.

- Run the host as a standalone host under vCenter Server.
- Move the host into another cluster.

Using Standby Mode

When a host machine is placed in standby mode, it is powered off.

Normally, hosts are placed in standby mode by the vSphere DPM feature to optimize power usage. You can also place a host in standby mode manually. However, DRS might undo (or recommend undoing) your change the next time it runs. To force a host to remain off, place it in maintenance mode and power it off.

Configuring VM Best Effort Restart Policy

This policy controls the evacuation behavior of a VM during maintenance mode.

When you enter maintenance mode, you must make sure that a host is ready for maintenance. This is often to upgrade the host. Whether you are performing an upgrade or not, you must evacuate all running VMs to put the host into maintenance mode. To do this you must determine how to deal with powered-on VMs. If you are performing an upgrade, you need to verify if the host can enter maintenance mode. If there are some VMs you cannot vMotion, you can power-off these VMs before you enter maintenance mode.

You can use compute policies to automate situations like this to evacuate VMs by powering-off the VMs when entering maintenance mode. Then you can tag all the desired VMs, and the system will power these VMs off when entering maintenance mode.

`BestEffortRestart` policy

During enter maintenance mode, `BestEffortRestart` policy attempts to shut down the VMs, if that fails the VMs are powered-off. Instead of waiting for the host to exit MM, the policy finds the best host for the VM while the original host is still in maintenance mode. `BestEffortRestart` creates tasks for VMs to be powered-on on the best hosts during enter maintenance mode so that these VMs are in powered-on state as soon as possible. If a VM cannot be powered-on for whatever reason, there is a remediation cycle that runs every 3 minutes. Exiting maintenance mode is no longer prerequisite for this VM to be running.

If you are moving a host out of a cluster and want VMs to remain in the cluster and not go away with the host you can use `BestEffortRestart`. In this case, you want VMs to remain in the cluster when the host is getting moved out of the cluster. So, during enter maintenance mode not only we have to power-off VMs but also find the best host for them excluding the current host and power them on there.

You can benefit from using `BestEffortRestart` for vGPU VMs and other passthrough devices. vGPU VMs have large memory framebuffers which are costly to migrate and may exceed the 100 second default vMotion timeout. These vMotions are likely to time out leaving vGPU VMs in an undesirable state. Instead, you can power-off vGPU VMs during enter maintenance mode. Ideally pass-through VMs, including vGPU VMs, should be powered-on on a different host during enter maintenance mode so that they don't have to wait until exitMM completes.

DRS must be activated for this policy to work properly. If DRS is deactivated either at the host or VM level, a VM is powered-off during enter maintenance mode.

If `BestEffortRestart` policy is deleted, depending on the timing of when the policy is deleted, there could be multiple outcomes.

- If a policy is removed before the evacuation action for a VM, this VM is treated the same as any other VM. There is no connection to `BestEffortRestart` policy.
- If the policy is deleted after the evacuation action a VM can be powered-off. The VM will not be remediated because this VM is no longer associated with `BestEffortRestart` policy.
- If a tag associated with `BestEffortRestart` policy is deleted before the evacuation action for a VM, this VM is treated the same as any other VM. There is no connection to `BestEffortRestart` policy.

The VMs should be in the same state that they were before maintenance mode. A VM will eventually be powered-on if it was powered-on before maintenance mode.

When `BestEffortRestart` policy is successful all `BestEffortRestart` policy VMs are powered-off and powered-on on the best hosts.

DRS Cluster Validity

The vSphere Client indicates whether a DRS cluster is valid, overcommitted (yellow), or invalid (red).

DRS clusters become overcommitted or invalid for several reasons.

- A cluster might become overcommitted if a host fails.
- A cluster becomes invalid if vCenter Server is unavailable and you power on virtual machines using the vSphere Client.
- A cluster becomes invalid if the user reduces the reservation on a parent resource pool while a virtual machine is in the process of failing over.
- If changes are made to hosts or virtual machines using the vSphere Client while vCenter Server is unavailable, those changes take effect. When vCenter Server becomes available again, you might find that clusters have turned red or yellow because cluster requirements are no longer met.


When considering cluster validity scenarios, you should understand these terms.

Reservation	A fixed, guaranteed allocation for the resource pool input by the user.
Reservation Used	The sum of the reservation or reservation used (whichever is larger) for each child resource pool, added recursively.
Unreserved	This nonnegative number differs according to resource pool type. • Nonexpandable resource pools: Reservation minus reservation used. • Expandable resource pools: (Reservation minus reservation used) plus any unreserved resources that can be borrowed from its ancestor resource pools.

Valid DRS Clusters

A valid cluster has enough resources to meet all reservations and to support all running virtual machines.


The following figure shows an example of a valid cluster with fixed resource pools and how its CPU and memory resources are computed.

Figure 134: Valid Cluster with Fixed Resource Pools

The cluster has the following characteristics:

- A cluster with total resources of 12GHz.
- Three resource pools, each of type **Fixed** (**Expandable Reservation** is not selected).
- The total reservation of the three resource pools combined is 11GHz (4+4+3 GHz). The total is shown in the **Reserved Capacity** field for the cluster.
- RP1 was created with a reservation of 4GHz. Two virtual machines. (VM1 and VM7) of 2GHz each are powered on (**Reservation Used**: 4GHz). No resources are left for powering on additional virtual machines. VM6 is shown as not powered on. It consumes none of the reservation.
- RP2 was created with a reservation of 4GHz. Two virtual machines of 1GHz and 2GHz are powered on (**Reservation Used**: 3GHz). 1GHz remains unreserved.
- RP3 was created with a reservation of 3GHz. One virtual machine with 3GHz is powered on. No resources for powering on additional virtual machines are available.

The following figure shows an example of a valid cluster with some resource pools (RP1 and RP3) using reservation type **Expandable**.

Figure 135: Valid Cluster with Expandable Resource Pools


A valid cluster can be configured as follows:

- A cluster with total resources of 16GHz.
- RP1 and RP3 are of type **Expandable**, RP2 is of type **Fixed**.
- The total reservation used of the three resource pools combined is 16GHz (6GHz for RP1, 5GHz for RP2, and 5GHz for RP3). 16GHz shows up as the **Reserved Capacity** for the cluster at top level.
- RP1 was created with a reservation of 4GHz. Three virtual machines of 2GHz each are powered on. Two of those virtual machines (for example, VM1 and VM7) can use RP1's reservations, the third virtual machine (VM6) can use reservations from the cluster's resource pool. (If the type of this resource pool were **Fixed**, you could not power on the additional virtual machine.)
- RP2 was created with a reservation of 5GHz. Two virtual machines of 1GHz and 2GHz are powered on (**Reservation Used**: 3GHz). 2GHz remains unreserved.
- RP3 was created with a reservation of 5GHz. Two virtual machines of 3GHz and 2GHz are powered on. Even though this resource pool is of type **Expandable**, no additional 2GHz virtual machine can be powered on because the parent's extra resources are already used by RP1.

Overcommitted DRS Clusters

A cluster becomes overcommitted (yellow) when the tree of resource pools and virtual machines is internally consistent but the cluster does not have the capacity to support all resources reserved by the child resource pools.

There will always be enough resources to support all running virtual machines because, when a host becomes unavailable, all its virtual machines become unavailable. A cluster typically turns yellow when cluster capacity is suddenly reduced, for example, when a host in the cluster becomes unavailable. VMware recommends that you leave adequate additional cluster resources to avoid your cluster turning yellow.

Figure 136: Yellow Cluster

In this example:

- A cluster with total resources of 12GHz coming from three hosts of 4GHz each.
 - Three resource pools reserving a total of 12GHz.
 - The total reservation used by the three resource pools combined is 12GHz (4+5+3 GHz). That shows up as the **Reserved Capacity** in the cluster.
 - One of the 4GHz hosts becomes unavailable, so total resources reduce to 8GHz.
 - At the same time, VM4 (1GHz) and VM3 (3GHz), which were running on the host that failed, are no longer running.
 - The cluster is now running virtual machines that require a total of 6GHz. The cluster still has 8GHz available, which is sufficient to meet virtual machine requirements.
- The resource pool reservations of 12GHz can no longer be met, so the cluster is marked as yellow.

Invalid DRS Clusters

A cluster enabled for DRS becomes invalid (red) when the tree is no longer internally consistent, that is, resource constraints are not observed.

The total amount of resources in the cluster does not affect whether the cluster is red. A cluster can be red, even if enough resources exist at the root level, if there is an inconsistency at a child level.


You can resolve a red DRS cluster problem either by powering off one or more virtual machines, moving virtual machines to parts of the tree that have sufficient resources, or editing the resource pool settings in the red part. Adding resources typically helps only when you are in the yellow state.

A cluster can also turn red if you reconfigure a resource pool while a virtual machine is failing over. A virtual machine that is failing over is disconnected and does not count toward the reservation used by the parent resource pool. You

might reduce the reservation of the parent resource pool before the failover completes. After the failover is complete, the virtual machine resources are again charged to the parent resource pool. If the pool's usage becomes larger than the new reservation, the cluster turns red.

If a user is able to start a virtual machine (in an unsupported way) with a reservation of 3GHz under resource pool 2, the cluster would become red, as shown in the following figure.

Figure 137: Red Cluster



Managing Power Resources

The vSphere Distributed Power Management (DPM) feature allows a DRS cluster to reduce its power consumption by powering hosts on and off based on cluster resource utilization.

vSphere DPM monitors the cumulative demand of all virtual machines in the cluster for memory and CPU resources and compares this to the total available resource capacity of all hosts in the cluster. If sufficient excess capacity is found, vSphere DPM places one or more hosts in standby mode and powers them off after migrating their virtual machines to other hosts. Conversely, when capacity is deemed to be inadequate, DRS brings hosts out of standby mode (powers them on) and uses vMotion to migrate virtual machines to them. When making these calculations, vSphere DPM considers not only current demand, but it also honors any user-specified virtual machine resource reservations.

If you enable **Forecasted Metrics** when you create a DRS cluster, DPM will issue proposals in advance depending on the rolling forecast window you select.

Note: ESXi hosts cannot automatically be brought out of standby mode unless they are running in a cluster managed by vCenter Server.

vSphere DPM can use one of three power management protocols to bring a host out of standby mode: Intelligent Platform Management Interface (IPMI), Hewlett-Packard Integrated Lights-Out (iLO), or Wake-On-LAN (WOL). Each protocol requires its own hardware support and configuration. If a host does not support any of these protocols it cannot be put

into standby mode by vSphere DPM. If a host supports multiple protocols, they are used in the following order: IPMI, iLO, WOL.

Note: Do not disconnect a host in standby mode or move it out of the DRS cluster without first powering it on, otherwise vCenter Server is not able to power the host back on.

Configure IPMI or iLO Settings for vSphere DPM

IPMI is a hardware-level specification and Hewlett-Packard iLO is an embedded server management technology. Each of them describes and provides an interface for remotely monitoring and controlling computers.

Both IPMI and iLO require a hardware Baseboard Management Controller (BMC) to provide a gateway for accessing hardware control functions, and allow the interface to be accessed from a remote system using serial or LAN connections. The BMC is powered-on even when the host itself is powered-off. If properly enabled, the BMC can respond to remote power-on commands.

If you plan to use IPMI or iLO as a wake protocol, you must configure the BMC. BMC configuration steps vary according to model. See your vendor's documentation for more information. With IPMI, you must also ensure that the BMC LAN channel is configured to be always available and to allow operator-privileged commands. On some IPMI systems, when you enable "IPMI over LAN" you must configure this in the BIOS and specify a particular IPMI account.

vSphere DPM using only IPMI supports MD5- and plaintext-based authentication, but MD2-based authentication is not supported. vCenter Server uses MD5 if a host's BMC reports that it is supported and enabled for the Operator role. Otherwise, plaintext-based authentication is used if the BMC reports it is supported and enabled. If neither MD5 nor plaintext authentication is enabled, IPMI cannot be used with the host and vCenter Server attempts to use Wake-on-LAN.

You must perform the following procedure on each host.

1. Browse to the host in the vSphere Client.
2. Click the **Configure** tab.
3. Under **System**, click **Power Management**.
4. Click **Edit**.
5. Enter the following information.
 - User name and password for a BMC account. (The user name must have the ability to remotely power the host on.)
 - IP address of the NIC associated with the BMC, as distinct from the IP address of the host. The IP address should be static or a DHCP address with infinite lease.
 - MAC address of the NIC associated with the BMC.
6. Click **OK**.

Test Wake-on-LAN for vSphere DPM

The use of Wake-on-LAN (WOL) for the vSphere DPM feature is fully supported, if you configure and successfully test it according to the VMware guidelines. You must perform these steps before activating vSphere DPM for a cluster for the first time or on any host that is being added to a cluster that is using vSphere DPM.

Before testing WOL, ensure that your cluster meets the prerequisites.

- Your cluster must contain at least two hosts that are version ESX 3.5 (or ESX 3i version 3.5) or later.
- Each host's vMotion networking link must be working correctly. The vMotion network should also be a single IP subnet, not multiple subnets separated by routers.
- The vMotion NIC on each host must support WOL. To check for WOL support, first determine the name of the physical network adapter corresponding to the VMkernel port by selecting the host in the inventory panel of the vSphere Client, selecting the **Configuration** tab, and clicking **Networking**. After you have this information, click on **Network Adapters**

and find the entry corresponding to the network adapter. The **Wake On LAN Supported** column for the relevant adapter should show Yes.

- To display the WOL-compatibility status for each NIC on a host, select the host in the inventory panel of the vSphere Client, select the **Configuration** tab, and click **Network Adapters**. The NIC must show Yes in the **Wake On LAN Supported** column.
- The switch port that each WOL-supporting vMotion NIC is plugged into should be set to auto negotiate the link speed, and not set to a fixed speed (for example, 1000 Mb/s). Many NICs support WOL only if they can switch to 100 Mb/s or less when the host is powered off.

After you verify these prerequisites, test each ESXi host that is going to use WOL to support vSphere DPM. When you test these hosts, ensure that the vSphere DPM feature is deactivated for the cluster.

CAUTION: Ensure that any host being added to a vSphere DPM cluster that uses WOL as a wake protocol is tested and deactivated from using power management if it fails the testing. If this is not done, vSphere DPM might power off hosts that it subsequently cannot power back up.

1. Browse to the host in the vSphere Client.
2. Right-click the host and select **Power > Enter Standby Mode**
This action powers down the host.
3. Right-click the host and select **Power > Power On** to attempt to bring it out of standby mode.
4. Observe whether or not the host successfully powers back on.
5. For any host that fails to exit standby mode successfully, perform the following steps.
 - a) Select the host in the vSphere Client and select the **Configure** tab.
 - b) Under **Hardware > Power Management**, click **Edit** to adjust the power management policy.
 After you do this, vSphere DPM does not consider that host a candidate for being powered off.

Activating vSphere DPM for a DRS Cluster

After you have performed configuration or testing steps required by the wake protocol you are using on each host, you can enable vSphere DPM.

Configure the power management automation level, threshold, and host-level overrides. These settings are configured under **Power Management** in the cluster's Settings dialog box.

You can also create scheduled tasks to enable and deactivate DPM for a cluster using the **Schedule Task: Change Cluster Power Settings** wizard.

Note: If a host in your DRS cluster has USB devices connected, deactivate DPM for that host. Otherwise, DPM might turn off the host and sever the connection between the device and the virtual machine that was using it.

Automation Level

Whether the host power state and migration recommendations generated by vSphere DPM are run automatically or not depends upon the power management automation level selected for the feature.

The automation level is configured under **Power Management** in the cluster's Settings dialog box.

Note: The power management automation level is not the same as the DRS automation level.

Table 363: Power Management Automation Level

Option	Description
Off	The feature is deactivated and no recommendations are made.
Manual	Host power operation and related virtual machine migration recommendations are made, but not automatically run.
Automatic	Host power operations are automatically run if related virtual machine migrations can all be run automatically.

vSphere DPM Threshold

The power state (host power on or off) recommendations generated by the vSphere DPM feature are assigned priorities that range from priority-one recommendations to priority-five recommendations.

These priority ratings are based on the amount of over- or under-utilization found in the DRS cluster and the improvement that is expected from the intended host power state change. A priority-one recommendation is mandatory, while a priority-five recommendation brings only slight improvement.

The threshold is configured under **Power Management** in the cluster's Settings dialog box. Each level you move the vSphere DPM Threshold slider to the right allows the inclusion of one more lower level of priority in the set of recommendations that are executed automatically or appear as recommendations to be manually executed. At the Conservative setting, vSphere DPM only generates priority-one recommendations, the next level to the right only priority-two and higher, and so on, down to the Aggressive level which generates priority-five recommendations and higher (that is, all recommendations.)

Note: The DRS threshold and the vSphere DPM threshold are essentially independent. You can differentiate the aggressiveness of the migration and host-power-state recommendations they respectively provide.

Host-Level Overrides

When you activate vSphere DPM in a DRS cluster, by default all hosts in the cluster inherit its vSphere DPM automation level.

You can override this default for an individual host by selecting the Host Options page of the cluster's Settings dialog box and clicking its **Power Management** setting. You can change this setting to the following options:

- Disabled
- Manual
- Automatic

Note: Do not change a host's Power Management setting if it has been set to Disabled due to failed exit standby mode testing.

After activating and running vSphere DPM, you can verify that it is functioning properly by viewing each host's **Last Time Exited Standby** information displayed on the Host Options page in the cluster Settings dialog box and on the **Hosts** tab for each cluster. This field shows a timestamp and whether vCenter Server Succeeded or Failed the last time it attempted to bring the host out of standby mode. If no such attempt has been made, the field displays Never.

Note: Times for the **Last Time Exited Standby** text box are derived from the vCenter Server event log. If this log is cleared, the times are reset to Never.

Monitoring vSphere DPM

You can use event-based alarms in vCenter Server to monitor vSphere DPM.

The most serious potential error you face when using vSphere DPM is the failure of a host to exit standby mode when its capacity is needed by the DRS cluster. You can monitor for instances when this error occurs by using the preconfigured **Exit Standby Error** alarm in vCenter Server. If vSphere DPM cannot bring a host out of standby mode (vCenter Server event `DrsExitStandbyModeFailedEvent`), you can configure this alarm to send an alert email to the administrator or to send notification using an SNMP trap. By default, this alarm is cleared after vCenter Server is able to successfully connect to that host.

To monitor vSphere DPM activity, you can also create alarms for the following vCenter Server events.

Table 364: vCenter Server Events

Event Type	Event Name
Entering Standby mode (about to power off host)	<code>DrsEnteringStandbyModeEvent</code>
Successfully entered Standby mode (host power off succeeded)	<code>DrsEnteredStandbyModeEvent</code>
Exiting Standby mode (about to power on the host)	<code>DrsExitingStandbyModeEvent</code>
Successfully exited Standby mode (power on succeeded)	<code>DrsExitedStandbyModeEvent</code>

For more information about creating and editing alarms, see the *vSphere Monitoring and Performance* documentation.

If you use monitoring software other than vCenter Server, and that software triggers alarms when physical hosts are powered off unexpectedly, you might have a situation where false alarms are generated when vSphere DPM places a host into standby mode. If you do not want to receive such alarms, work with your vendor to deploy a version of the monitoring software that is integrated with vCenter Server. You could also use vCenter Server itself as your monitoring solution, because starting with vSphere 4.x, it is inherently aware of vSphere DPM and does not trigger these false alarms.

VM Static Power

vSphere 8.0 U1 introduces a new power metric to estimate VM static power.

You can view this power metric in the **Advanced** pane of the vCenter Performance Monitor page in vCenter. This power metric is a semi-fixed value that is a function of the VM's configuration, specifically the number of vCPUs and memory size as well as the configuration of the host it is running on. VM static power estimates the minimal amount of required power based on the requirements of a physical machine with the same configuration, running on similar hardware as the ESXi host. A VM may report different values of static power on two different hosts, if one host is more power efficient than the other. This is useful for dynamic workload placement strategies to optimize for overall power and may help estimate energy savings due to virtualization of workloads.

Using Affinity Rules with vSphere DRS

You can control the placement of virtual machines on hosts within a cluster by using affinity rules.

Note: This topic describes using affinity rules with DRS. You can also use affinity rules without DRS.

You can create two types of rules.

- Used to specify affinity or anti-affinity between a group of virtual machines and a group of hosts. An affinity rule specifies that the members of a selected virtual machine DRS group can or must run on the members of a specific host DRS group. An anti-affinity rule specifies that the members of a selected virtual machine DRS group cannot run on the members of a specific host DRS group.
See [VM-Host Affinity Rules with vSphere DRS](#) for information about creating and using this type of rule.
- Used to specify affinity or anti-affinity between individual virtual machines. A rule specifying affinity causes DRS to try to keep the specified virtual machines together on the same host, for example, for performance reasons. With an anti-

affinity rule, DRS tries to keep the specified virtual machines apart, for example, so that when a problem occurs with one host, you do not lose both virtual machines.

See [VM-VM Affinity Rules with vSphere DRS](#) for information about creating and using this type of rule.

When you add or edit an affinity rule, and the cluster's current state is in violation of the rule, the system continues to operate and tries to correct the violation. For manual and partially automated DRS clusters, migration recommendations based on rule fulfillment and load balancing are presented for approval. You are not required to fulfill the rules, but the corresponding recommendations remain until the rules are fulfilled.

To check whether any enabled affinity rules are being violated and cannot be corrected by DRS, select the cluster's **DRS** tab and click **Faults**. Any rule currently being violated has a corresponding fault on this page. Read the fault to determine why DRS is not able to satisfy the particular rule. Rules violations also produce a log event.

Note: VM-VM and VM-Host affinity rules are different from an individual host's CPU affinity rules.

VM-VM Affinity Rules with vSphere DRS

A VM-VM affinity rule with DRS specifies whether selected individual virtual machines should run on the same host or be kept on separate hosts. This type of rule is used to create affinity or anti-affinity between individual virtual machines that you select.

Note: This topic describes using affinity rules with DRS. You can also use affinity rules without DRS.

When an affinity rule is created, DRS tries to keep the specified virtual machines together on the same host. You might want to do this, for example, for performance reasons.

With an anti-affinity rule, DRS tries to keep the specified virtual machines apart. You could use such a rule if you want to guarantee that certain virtual machines are always on different physical hosts. In that case, if a problem occurs with one host, not all virtual machines would be placed at risk.

VM-Host Affinity Rules with vSphere DRS

A VM-Host affinity rule specifies whether or not the members of a selected virtual machine DRS group can run on the members of a specific host DRS group.

Note: This topic describes using affinity rules for DRS. You can also use affinity rules without DRS.

Unlike a VM-VM affinity rule, which specifies affinity (or anti-affinity) between individual virtual machines, a VM-Host affinity rule specifies an affinity relationship between a group of virtual machines and a group of hosts. There are 'required' rules (designated by "must") and 'preferential' rules (designated by "should").

A VM-Host affinity rule includes the following components.

- One virtual machine DRS group.
- One host DRS group.
- A designation of whether the rule is a requirement ("must") or a preference ("should") and whether it is affinity ("run on") or anti-affinity ("not run on").

Because VM-Host affinity rules are cluster-based, the virtual machines and hosts that are included in a rule must all reside in the same cluster. If a virtual machine is removed from the cluster, it loses its DRS group affiliation, even if it is later returned to the cluster.

Create a VM-Host Group

A VM-Host affinity rule establishes an affinity (or anti-affinity) relationship between a virtual machine group with a host group. You must create both of these groups before you can create a rule that links them.

1. Browse to the cluster in the vSphere Client.
2. Click the **Configure** tab.
3. Under **Configuration**, select **VM/Host Groups** and click **Add**.
4. In the **Create VM/Host Group** dialog box, type a name for the group.
5. Select **Host Group** from the **Type** drop down box and click **Add**.
6. Click the check box next to a host to add it. Continue this process until all desired hosts have been added.
7. Click **OK**.

Using this host group, you can create a VM-Host affinity rule that establishes an affinity (or anti-affinity) relationship with an appropriate virtual machine group.

Create a Virtual Machine Group

Affinity rules establish an affinity (or anti-affinity) relationship between groups. You must create groups before you can create a rule that links them.

1. Browse to the cluster in the vSphere Client.
2. Click the **Configure** tab.
3. Under **Configuration**, select **VM/Host Groups** and click **Add**.
4. In the **Create VM/Host Group** dialog box, type a name for the group.
5. Select **VM Group** from the **Type** drop down box and click **Add**.
6. Click the check box next to a virtual machine to add it. Continue this process until all desired virtual machines have been added.
7. Click **OK**.

Create a VM-VM Affinity Rule

You can create VM-VM affinity rules to specify whether selected individual virtual machines should run on the same host or be kept on separate hosts.

Note: If you use the vSphere HA Specify Failover Hosts admission control policy and designate multiple failover hosts, VM-VM affinity rules are not supported.

1. Browse to the cluster in the vSphere Client.
2. Click the **Configure** tab.
3. Under **Configuration**, click **VM/Host Rules**.
4. Click **Add**.
5. In the **Create VM/Host Rule** dialog box, type a name for the rule.
6. From the **Type** drop-down menu, select either **Keep Virtual Machines Together** or **Separate Virtual Machines**.
7. Click **Add**.
8. Select at least two virtual machines to which the rule will apply and click **OK**.
9. Click **OK**.

VM-VM Affinity Rule Conflicts

You can create and use multiple VM-VM affinity rules, however, this might lead to situations where the rules conflict with one another.

If two VM-VM affinity rules are in conflict, you cannot activate both. For example, if one rule keeps two virtual machines together and another rule keeps the same two virtual machines apart, you cannot activate both rules. Select one of the rules to apply and deactivate or remove the conflicting rule.

When two VM-VM affinity rules conflict, the older one takes precedence and the newer rule is deactivated. DRS only tries to satisfy activated rules and deactivated rules are ignored. DRS gives higher precedence to preventing violations of anti-affinity rules than violations of affinity rules.

Create a VM-Host Affinity Rule

You can create VM-Host affinity rules to specify whether or not the members of a selected virtual machine group can run on the members of a specific host group.

Create the virtual machine and host groups to which the VM-Host affinity rule applies.

1. Browse to the cluster in the vSphere Client.
2. Click the **Configure** tab.
3. Under **Configuration**, click **VM/Host Rules**.
4. Click **Add**.
5. In the **Create VM/Host Rule** dialog box, type a name for the rule.
6. From the **Type** drop down menu, select **Virtual Machines to Hosts**.
7. Select the virtual machine group and the host group to which the rule applies.
8. Select a specification for the rule.
 - **Must run on hosts in group.** Virtual machines in VM Group 1 must run on hosts in Host Group A.
 - **Should run on hosts in group.** Virtual machines in VM Group 1 should, but are not required, to run on hosts in Host Group A.
 - **Must not run on hosts in group.** Virtual machines in VM Group 1 must never run on host in Host Group A.
 - **Should not run on hosts in group.** Virtual machines in VM Group 1 should not, but might, run on hosts in Host Group A.
9. Click **OK**.

Using VM-Host Affinity Rules

You use a VM-Host affinity rule to specify an affinity relationship between a group of virtual machines and a group of hosts. When using VM-Host affinity rules, you should be aware of when they could be most useful, how conflicts between rules are resolved, and the importance of caution when setting required affinity rules.

If you create more than one VM-Host affinity rule, the rules are not ranked, but are applied equally. Be aware that this has implications for how the rules interact. For example, a virtual machine that belongs to two groups, each of which belongs to a different required rule, can run only on hosts that belong to both of the host groups represented in the rules.

When you create a VM-Host affinity rule, its ability to function in relation to other rules is not checked. So it is possible for you to create a rule that conflicts with the other rules you are using. When two VM-Host affinity rules conflict, the older one takes precedence and the newer rule is deactivated. DRS only tries to satisfy activated rules and deactivated rules are ignored.

DRS, vSphere HA, and vSphere DPM never take any action that results in the violation of required affinity rules (those where the virtual machine group 'must run on' or 'must not run on' the host group). Accordingly, you should exercise caution when using this type of rule because of its potential to adversely affect the functioning of the cluster. If improperly used, required VM-Host affinity rules can fragment the cluster and inhibit the proper functioning of DRS, vSphere HA, and vSphere DPM.

A number of cluster functions are not performed if doing so would violate a required affinity rule.

- DRS does not evacuate virtual machines to place a host in maintenance mode.
- DRS does not place virtual machines for power-on or load balance virtual machines.
- vSphere HA does not perform failovers.
- vSphere DPM does not optimize power management by placing hosts into standby mode.

To avoid these situations, exercise caution when creating more than one required affinity rule or consider using VM-Host affinity rules that are preferential only (those where the virtual machine group 'should run on' or 'should not run on' the host group). Ensure that the number of hosts in the cluster with which each virtual machine is affinity is large enough that losing a host does not result in a lack of hosts on which the virtual machine can run. Preferential rules can be violated to allow the proper functioning of DRS, vSphere HA, and vSphere DPM.

Note: You can create an event-based alarm that is triggered when a virtual machine violates a VM-Host affinity rule. Add a new alarm for the virtual machine and select **VM is violating VM-Host Affinity Rule** as the event trigger. For more information about creating and editing alarms, see the vSphere Monitoring and Performance documentation.

Using Affinity Rules without vSphere DRS

You can use affinity rules without using vSphere DRS.

VM affinity rules allow you to control the placement of VMs within a vSphere cluster. VM-VM affinity means the VMs must be run on the same host within a cluster. Workloads are comprised of multiple VMs that have dependencies on each other. Using **Keep Virtual Machines Together** rules ensure that all selected VMs run on the same host.

VM-VM anti-affinity means the VMs must be run on different hosts within a cluster. You can use a **Separate Virtual Machines** rule when you have multiple resource intensive VMs to ensure that these VMs are always running on different hosts in the cluster to avoid resource contention.

VM-Host affinity rules control the VM to Host relationship and has 4 sub-rules. Use cases for VM-Host rules would include making sure workloads run (or do not run) on specific hosts in a mixed-hardware cluster. For licensing reasons, you may need to ensure that an application only runs on a specific host or specific subset of hosts.

The VM-Host rules leverage VM-Host groups. A VM group is a user-defined group of one or more VMs. A host group is a user-defined group of one or more hosts.

Rule	Description	Persistence
Must run on hosts in group	The VM group must be run on the defined host group.	A hard rule that cannot be violated.
Must not run on hosts in group	The VM group must not be run on the defined host group.	A hard rule that cannot be violated.
Should run on hosts in group	The VM group should be run on the defined host group.	A soft rule that can be violated.
Should not run on hosts in group	The VM group should not be run on the defined host group.	A soft rule that can be violated.

VM-Host groups and VM-Host rules can be created, deleted, and edited when vSphere DRS is not activated. The interface to manage VM-Host groups and VM-Host rules exists outside of the vSphere DRS configuration.

In the event of a vSphere HA failover of VMs, affinity rules are respected. This is true for rule types **Keep Virtual Machines Together** and **Separate Virtual Machines**, as well as hard VM-Host rule types. vSphere HA is unable to failover VMs if a hard rule would be violated. When using soft VM-Host rules, vSphere HA operations succeed even if the rule is violated. Soft rules allow for violation of their definitions.

For example, a 2-node cluster and a **Separate Virtual Machines** rule defined to place VM-1 on ESX-01 and VM-2 on ESX-02. ESX-01 experiences a failure and vSphere HA is initiated. vSphere HA is unable to power on VM-1 on ESX-02 as it would violate the rule to keep each virtual machine separate.

VM power on operations do not allow affinity rules to be violated. This is true for rule types **Keep Virtual Machines Together** and **Separate Virtual Machines**, as well as hard VM-Host rules. When using soft VM-Host rules, power on operations succeed even if the rule is violated. Soft rules allow for violation of their definitions.

For example, if a rule is defined to **Keep Virtual Machines Together** on VM-1 and VM-2 and VM-1 is currently already powered on, an attempt to power on VM-2 will fail if it is registered to a different host. Since vSphere DRS is not activated, the user must manually migrate the VMs to the same host to be able to power on the VM.

When vSphere DRS is not activated and therefore not automatically migrating VMs, you can initiate migrations using vSphere vMotion on VMs with affinity rules defined. Both rule types **Keep Virtual Machines Together** and **Separate**

Virtual Machines can be violated by migrating using vSphere vMotion. For example, if a rule is defined to **Separate Virtual Machines** on VM-1 and VM-2, an attempt to migrate the VMs to the same host will not be blocked. The operation will succeed even though it would violate the rule.

Any hard VM-Host rule cannot be violated by a user-initiated migration. For example, if a rule is defined that VM-1 and VM-2 must run on hosts ESX-01 and ESX-02, an attempt to migrate the VMs to hosts outside of the host group is blocked.

Soft VM-Host rules can be violated by a user-initiated migration. Soft rules allow for violation of their definitions. For example, if a rule is defined that VM-1 and VM-2 should not run on hosts ESX-03 and ESX-04, an attempt to migrate the VMs to either ESX-03 or ESX-04 is not blocked.

When vSphere DRS is not activated, there is no automatic remediation of VM placement. This includes violations because of manual migrations and violations that occur if new rules are defined that would violate current VM placement. Without vSphere DRS, you must correct VM placement to comply with rules.

Creating a Datastore Cluster in vSphere

A datastore cluster is a collection of datastores with shared resources and a shared management interface. Datastore clusters are to datastores what clusters are to hosts. When you create a datastore cluster, you can use vSphere Storage DRS to manage storage resources.

Note: Datastore clusters are referred to as storage pods in the vSphere API.

When you add a datastore to a datastore cluster, the datastore's resources become part of the datastore cluster's resources. As with clusters of hosts, you use datastore clusters to aggregate storage resources, which enables you to support resource allocation policies at the datastore cluster level. The following resource management capabilities are also available per datastore cluster.

Space utilization load balancing

You can set a threshold for space use. When space use on a datastore exceeds the threshold, Storage DRS generates recommendations or performs Storage vMotion migrations to balance space use across the datastore cluster.

I/O latency load balancing

You can set an I/O latency threshold for bottleneck avoidance. When I/O latency on a datastore exceeds the threshold, Storage DRS generates recommendations or performs Storage vMotion migrations to help alleviate high I/O load.

Anti-affinity rules

You can create anti-affinity rules for virtual machine disks. For example, the virtual disks of a certain virtual machine must be kept on different datastores. By default, all virtual disks for a virtual machine are placed on the same datastore.

Initial Placement and Ongoing Balancing

Storage DRS provides initial placement and ongoing balancing recommendations to datastores in a Storage DRS-enabled datastore cluster.

Initial placement occurs when Storage DRS selects a datastore within a datastore cluster on which to place a virtual machine disk. This happens when the virtual machine is being created or cloned, when a virtual machine disk is being migrated to another datastore cluster, or when you add a disk to an existing virtual machine.

Initial placement recommendations are made in accordance with space constraints and with respect to the goals of space and I/O load balancing. These goals aim to minimize the risk of over-provisioning one datastore, storage I/O bottlenecks, and performance impact on virtual machines.

Storage DRS is invoked at the configured frequency (by default, every eight hours) or when one or more datastores in a datastore cluster exceeds the user-configurable space utilization thresholds. When Storage DRS is invoked, it checks each datastore's space utilization and I/O latency values against the threshold. For I/O latency, Storage DRS uses the 90th percentile I/O latency measured over the course of a day to compare against the threshold.

Storage Migration Recommendations

vCenter Server displays migration recommendations on the Storage DRS Recommendations page for datastore clusters that have manual automation mode.

The system provides as many recommendations as necessary to enforce Storage DRS rules and to balance the space and I/O resources of the datastore cluster. Each recommendation includes the virtual machine name, the virtual disk name, the name of the datastore cluster, the source datastore, the destination datastore, and a reason for the recommendation.

- Balance datastore space use
- Balance datastore I/O load

Storage DRS makes mandatory recommendations for migration in the following situations:

- The datastore is out of space.
- Anti-affinity or affinity rules are being violated.
- The datastore is entering maintenance mode and must be evacuated.

In addition, optional recommendations are made when a datastore is close to running out of space or when adjustments should be made for space and I/O load balancing.

Storage DRS considers moving virtual machines that are powered off or powered on for space balancing. Storage DRS includes powered-off virtual machines with snapshots in these considerations.

Create a Datastore Cluster

You can manage datastore cluster resources using Storage DRS.

1. Browse to data centers in the vSphere Client.
2. Right-click the data center and select **Storage, New Datastore Cluster**.
3. To complete the **New Datastore Cluster** wizard, follow the prompts.
4. Click **Finish**.

Activate and Deactivate Storage DRS

Storage DRS allows you to manage the aggregated resources of a datastore cluster. When Storage DRS is activated, it provides recommendations for virtual machine disk placement and migration to balance space and I/O resources across the datastores in the datastore cluster.

When you activate Storage DRS, you activate the following functions.

- Space load balancing among datastores within a datastore cluster.
- I/O load balancing among datastores within a datastore cluster.
- Initial placement for virtual disks based on space and I/O workload.

The Enable Storage DRS check box in the Datastore Cluster Settings dialog box activates or deactivates all of these components at once. If necessary, you can deactivate I/O-related functions of Storage DRS independently of space balancing functions.

When you deactivate Storage DRS on a datastore cluster, Storage DRS settings are preserved. When you activate Storage DRS, the settings for the datastore cluster are restored to the point where Storage DRS was deactivated.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Services**.
3. Select **Storage DRS** and click **Edit**.
4. Select **Turn ON vSphere DRS** and click **OK**.
5. Optional: To deactivate only I/O-related functions of Storage DRS, leaving space-related controls activated, perform the following steps.
 - a) Under **Storage DRS** select **Edit**.
 - b) Deselect the **Enable I/O metric for Storage DRS** check box and click **OK**.

Set the Automation Level for Datastore Clusters

The automation level for a datastore cluster specifies whether or not placement and migration recommendations from Storage DRS are applied automatically.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Services**.
3. Select **DRS** and click **Edit**.
4. Expand DRS Automation and select an automation level.
Manual is the default automation level.

Option	Description
No Automation (Manual Mode)	Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.
Fully Automated	Placement and migration recommendations run automatically.

5. Click **OK**.

Setting the Aggressiveness Level for Storage DRS

The aggressiveness of Storage DRS is determined by specifying thresholds for space used and I/O latency.

Storage DRS collects resource usage information for the datastores in a datastore cluster. vCenter Server uses this information to generate recommendations for placement of virtual disks on datastores.

When you set a low aggressiveness level for a datastore cluster, Storage DRS recommends Storage vMotion migrations only when absolutely necessary, for example, if when I/O load, space utilization, or their imbalance is high. When you set a high aggressiveness level for a datastore cluster, Storage DRS recommends migrations whenever the datastore cluster can benefit from space or I/O load balancing.

In the vSphere Client, you can use the following thresholds to set the aggressiveness level for Storage DRS:

Space Utilization

Storage DRS generates recommendations or performs migrations when the percentage of space utilization on the datastore is greater than the threshold you set in the vSphere Client.

I/O Latency

Storage DRS generates recommendations or performs migrations when the 90th percentile I/O latency measured over a day for the datastore is greater than the threshold.

You can also set advanced options to further configure the aggressiveness level of Storage DRS.

Space utilization difference

This threshold ensures that there is some minimum difference between the space utilization of the source and the destination. For example, if the space used on datastore A is 82% and datastore B is 79%, the difference is 3. If the threshold is 5, Storage DRS will not make migration recommendations from datastore A to datastore B.

I/O load balancing invocation interval

After this interval, Storage DRS runs to balance I/O load.

I/O imbalance threshold

Lowering this value makes I/O load balancing less aggressive. Storage DRS computes an I/O fairness metric between 0 and 1, which 1 being the fairest distribution. I/O load balancing runs only if the computed metric is less than $1 - (\text{I/O imbalance threshold} / 100)$.

Set Storage DRS Runtime Rules

Set Storage DRS triggers and configure advanced options for the datastore cluster.

- Optional: Select or deselect the **Enable I/O metric for SDRS recommendations** check box to activate or deactivate I/O metric inclusion.

When you deactivate this option, vCenter Server does not consider I/O metrics when making Storage DRS recommendations. When you deactivate this option, you deactivate the following elements of Storage DRS:

- I/O load balancing among datastores within a datastore cluster.
- Initial placement for virtual disks based on I/O workload. Initial placement is based on space only.

- Optional: Set Storage DRS thresholds.

You set the aggressiveness level of Storage DRS by specifying thresholds for used space and I/O latency.

- Use the Utilized Space slider to indicate the maximum percentage of consumed space allowed before Storage DRS is triggered. Storage DRS makes recommendations and performs migrations when space use on the datastores is higher than the threshold.
- Use the I/O Latency slider to indicate the maximum I/O latency allowed before Storage DRS is triggered. Storage DRS makes recommendations and performs migrations when latency is higher than the threshold.

Note: The Storage DRS I/O Latency threshold for the datastore cluster should be lower than or equal to the Storage I/O Control congestion threshold.

- Optional: Configure advanced options.

- No recommendations until utilization difference between source and destination is: Use the slider to specify the space utilization difference threshold. Utilization is $\text{usage} * 100 / \text{capacity}$.
This threshold ensures that there is some minimum difference between the space utilization of the source and the destination. For example, if the space used on datastore A is 82% and datastore B is 79%, the difference is 3. If the threshold is 5, Storage DRS will not make migration recommendations from datastore A to datastore B.
- Check imbalances every: Specify how often Storage DRS should assess space and I/O load balancing.
- I/O imbalance threshold: Use the slider to indicate the aggressiveness of I/O load balancing. Lowering this value makes I/O load balancing less aggressive. Storage DRS computes an I/O fairness metric between 0 and 1, which 1 being the fairest distribution. I/O load balancing runs only if the computed metric is less than $1 - (\text{I/O imbalance threshold} / 100)$.

- Click **OK**.

Datastore Cluster Requirements

Datastores and hosts that are associated with a datastore cluster must meet certain requirements to use datastore cluster features successfully.

Follow these guidelines when you create a datastore cluster.

- Datastore clusters must contain similar or interchangeable datastores.
A datastore cluster can contain a mix of datastores with different sizes and I/O capacities, and can be from different arrays and vendors. However, the following types of datastores cannot coexist in a datastore cluster.
 - NFS and VMFS datastores cannot be combined in the same datastore cluster.
 - Replicated datastores cannot be combined with non-replicated datastores in the same Storage-DRS-enabled datastore cluster.
- All hosts attached to the datastores in a datastore cluster must be ESXi 5.0 and later. If datastores in the datastore cluster are connected to ESX/ESXi 4.x and earlier hosts, Storage DRS does not run.
- Datastores shared across multiple data centers cannot be included in a datastore cluster.
- As a best practice, do not include datastores that have hardware acceleration enabled in the same datastore cluster as datastores that do not have hardware acceleration enabled. Datastores in a datastore cluster must be homogeneous to guarantee hardware acceleration-supported behavior.

Adding and Removing Datastores from a Datastore Cluster

You add and remove datastores to and from an existing datastore cluster.

You can add to a datastore cluster any datastore that is mounted on a host in the vSphere Client inventory, with the following exceptions:

- All hosts attached to the datastore must be ESXi 5.0 and later.
- The datastore cannot be in more than one data center in the same instance of the vSphere Client.

When you remove a datastore from a datastore cluster, the datastore remains in the vSphere Client inventory and is not unmounted from the host.

Use Datastore Clusters to Manage Storage Resources in vSphere

After you create a datastore cluster, you can customize it and use it to manage storage I/O and space utilization resources.

Using Storage DRS Maintenance Mode

You place a datastore in maintenance mode when you need to take it out of use to service it. A datastore enters or leaves maintenance mode only as the result of a user request.

Maintenance mode is available to datastores within a Storage DRS-enabled datastore cluster. Standalone datastores cannot be placed in maintenance mode.

Virtual disks that are located on a datastore that is entering maintenance mode must be migrated to another datastore, either manually or using Storage DRS. When you attempt to put a datastore in maintenance mode, the **Placement Recommendations** tab displays a list of migration recommendations, datastores within the same datastore cluster where virtual disks can be migrated. On the **Faults** tab, vCenter Server displays a list of the disks that cannot be migrated and the reasons why. If Storage DRS affinity or anti-affinity rules prevent disks from being migrated, you can choose to enable the Ignore Affinity Rules for Maintenance option.

The datastore is in a state of Entering Maintenance Mode until all virtual disks have been migrated.

Place a Datastore in Maintenance Mode

If you need to take a datastore out of service, you can place the datastore in Storage DRS maintenance mode.

Storage DRS is enabled on the datastore cluster that contains the datastore that is entering maintenance mode.

No CD-ROM image files are stored on the datastore.

There are at least two datastores in the datastore cluster.

1. Browse to the datastore in the vSphere Client.
2. Right-click the datastore and select **Maintenance Mode > Enter Maintenance Mode**.
A list of recommendations appears for datastore maintenance mode migration.
3. Optional: On the Placement Recommendations tab, deselect any recommendations you do not want to apply.
Note: The datastore cannot enter maintenance mode without evacuating all disks. If you deselect recommendations, you must manually move the affected virtual machines.
4. If necessary, click **Apply Recommendations**.
vCenter Server uses Storage vMotion to migrate the virtual disks from the source datastore to the destination datastore and the datastore enters maintenance mode.

The datastore icon might not be immediately updated to reflect the datastore's current state. To update the icon immediately, click **Refresh**.

Ignore Storage DRS Affinity Rules for Maintenance Mode

Storage DRS affinity or anti-affinity rules might prevent a datastore from entering maintenance mode. You can ignore these rules when you put a datastore in maintenance mode.

When you activate the Ignore Affinity Rules for Maintenance option for a datastore cluster, vCenter Server ignores Storage DRS affinity and anti-affinity rules that prevent a datastore from entering maintenance mode.

Storage DRS rules are ignored only for evacuation recommendations. vCenter Server does not violate the rules when making space and load balancing recommendations or initial placement recommendations.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Services**.
3. Select **DRS** and click **Edit**.
4. Expand **Advanced Options** and click **Add**.
5. In the Option column, type `IgnoreAffinityRulesForMaintenance`.
6. In the Value column, type `1` to activate the option.
Type `0` to deactivate the option.
7. Click **OK**.

The Ignore Affinity Rules for Maintenance Mode option is applied to the datastore cluster.

Applying Storage DRS Recommendations

Storage DRS collects resource usage information for all datastores in a datastore cluster. Storage DRS uses the information to generate recommendations for virtual machine disk placement on datastores in a datastore cluster.

Storage DRS recommendations appear on the **Storage DRS** tab in the vSphere Client datastore view. Recommendations also appear when you attempt to put a datastore into Storage DRS maintenance mode. When you apply Storage DRS recommendations, vCenter Server uses Storage vMotion to migrate virtual machine disks to other datastores in the datastore cluster to balance the resources.

You can apply a subset of the recommendations by selecting the Override Suggested DRS Recommendations check box and selecting each recommendation to apply.

Table 365: Storage DRS Recommendations

Label	Description
Priority	Priority level (1-5) of the recommendation. (Hidden by default.)
Recommendation	Action being recommended by Storage DRS.
Reason	Why the action is needed.
Space Utilization % Before (source) and (destination)	Percentage of space used on the source and destination datastores before migration.
Space Utilization % After (source) and (destination)	Percentage of space used on the source and destination datastores after migration.
I/O Latency Before (source)	Value of I/O latency on the source datastore before migration.
I/O Latency Before (destination)	Value of I/O latency on the destination datastore before migration.

Refresh Storage DRS Recommendations

Storage DRS migration recommendations appear on the **Storage DRS** tab in the vSphere Client. You can refresh these recommendations by running Storage DRS.

At least one datastore cluster must exist in the vSphere Client inventory.

Enable Storage DRS for the datastore cluster. The **Storage DRS** tab appears only if Storage DRS is enabled.

1. In the vSphere Client datastore view, select the datastore cluster and click the **Storage DRS** tab.
2. Select the **Recommendations** view and click the **Run Storage DRS** link in the upper right corner.

The recommendations are updated. The Last Updated timestamp displays the time when Storage DRS recommendations were refreshed.

Change Storage DRS Automation Level for a Virtual Machine

You can override the datastore cluster-wide automation level for individual virtual machines. You can also override default virtual disk affinity rules.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Configuration**.
3. Under **VM Overrides**, select **Add**.
4. Select a virtual machine.
5. Click the Automation level drop-down menu, and select an automation level for the virtual machine.

Option	Description
Default (Manual)	Placement and migration recommendations are displayed, but do not run until you manually apply the recommendation.
Fully Automated	Placement and migration recommendations run automatically.
Disabled	vCenter Server does not migrate the virtual machine or provide migration recommendations for it.

6. Click the **Keep VMDKs together**, drop-down menu to override default VMDK affinity.
See [Override VMDK Affinity Rules](#).

7. Click **OK**.

Set Up Off-Hours Scheduling for Storage DRS

You can create a scheduled task to change Storage DRS settings for a datastore cluster so that migrations for fully automated datastore clusters are more likely to occur during off-peak hours.

Enable Storage DRS.

You can create a scheduled task to change the automation level and aggressiveness level for a datastore cluster. For example, you might configure Storage DRS to run less aggressively during peak hours, when performance is a priority, to minimize the occurrence of storage migrations. During non-peak hours, Storage DRS can run in a more aggressive mode and be invoked more frequently.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Services**.
3. Under **vSphere DRS** click the **Schedule DRS** button.
4. In the Edit Datastore Cluster dialog box, click **SDRS Scheduling**.
5. Expand **DRS Automation**.
 - a) Select an automation level.
 - b) Set the Migration threshold.

Use the Migration slider to select the priority level of vCenter Server recommendations that adjust the cluster's load balance.
 - c) Select whether to enable Virtual Machine Automation.

Override for individual virtual machines can be set from the VM Overrides page.
6. Expand **Power Managment**.
 - a) Select an automation level.
 - b) Set the DPM threshold.

Use the DPM slider to select the power recommendations that vCenter Server will apply.
7. Type a Task name.
8. Type a description of the Task you have created.
9. Under Configured Scheduler, click **Change** and select the time for the task to run and click **OK**.
10. Type an email address to send a notification email to when the task is complete.
11. Click **OK**.

The scheduled task runs at the specified time.

Storage DRS Anti-Affinity Rules

You can create Storage DRS anti-affinity rules to control which virtual disks should not be placed on the same datastore within a datastore cluster. By default, a virtual machine's virtual disks are kept together on the same datastore.

When you create an anti-affinity rule, it applies to the relevant virtual disks in the datastore cluster. Anti-affinity rules are enforced during initial placement and Storage DRS-recommendation migrations, but are not enforced when a migration is initiated by a user.

Note: Anti-affinity rules do not apply to CD-ROM ISO image files that are stored on a datastore in a datastore cluster, nor do they apply to swapfiles that are stored in user-defined locations.

VM Anti-Affinity Rules

Specify which virtual machines should never be kept on the same datastore. See [Create VM Anti-Affinity Rules](#).

VMDK Anti-Affinity Rules

Specify which virtual disks associated with a particular virtual machine must be kept on different datastores. See [Create VMDK Anti-Affinity Rules](#).

If you move a virtual disk out of the datastore cluster, the affinity or anti-affinity rule no longer applies to that disk.

When you move virtual disk files into a datastore cluster that has existing affinity and anti-affinity rules, the following behavior applies:

- Datastore Cluster B has an intra-VM affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, any rule that applied to the virtual disk for a given virtual machine in Datastore Cluster A no longer applies. The virtual disk is now subject to the intra-VM affinity rule in Datastore Cluster B.
- Datastore Cluster B has an VM anti-affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, any rule that applied to the virtual disk for a given virtual machine in Datastore Cluster A no longer applies. The virtual disk is now subject to the VM anti-affinity rule in Datastore Cluster B.
- Datastore Cluster B has an VMDK anti-affinity rule. When you move a virtual disk out of Datastore Cluster A and into Datastore Cluster B, the VMDK anti-affinity rule does not apply to the virtual disk for a given virtual machine because the rule is limited to only specified virtual disks in Datastore Cluster B.

Note: Storage DRS rules might prevent a datastore from entering maintenance mode. You can choose to ignore Storage DRS rules for maintenance mode by enabling the Ignore Affinity Rules for Maintenance option.

Create VM Anti-Affinity Rules

You can create an anti-affinity rule to indicate that all virtual disks of certain virtual machines must be kept on different datastores. The rule applies to individual datastore clusters.

Virtual machines that participate in an VM anti-affinity rule in a datastore cluster must be associated with an intra-VM affinity rule in the datastore cluster. The virtual machines must also comply with the intra-VM affinity rule.

If a virtual machine is subject to an VM anti-affinity rule, the following behavior applies:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.
- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

No VM anti-affinity rules are defined by default.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Configuration**.
3. Select **VM/Host Rules**.
4. Click **Add**.
5. Type a name for the rule.
6. From the Type menu, select **VM anti-affinity**.
7. Click **Add**.
8. Click **Select Virtual Machine**.
9. Select at least two virtual machines and click **OK**.
10. Click **OK** to save the rule.

Create VMDK Anti-Affinity Rules

You can create a VMDK anti-affinity rule for a virtual machine that indicates which of its virtual disks must be kept on different datastores.

VMDK anti-affinity rules apply to the virtual machine for which the rule is defined, not to all virtual machines. The rule is expressed as a list of virtual disks that are to be separated from one another.

If you attempt to set an VMDK anti-affinity rule and an intra-VM affinity rule for a virtual machine, vCenter Server rejects the most recently defined rule.

If a virtual machine is subject to a VMDK anti-affinity rule, the following behavior applies:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.
- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

No VMDK anti-affinity rules are defined by default.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Configuration**.
3. Select **VM/Host Rules**.
4. Click **Add**.
5. Type a name for the rule.
6. From the Type menu, select **VMDK anti-affinity**.
7. Click **Add**.
8. Click **Select Virtual Machine**.
9. Select a virtual machine and click **OK**.
10. Select at least two virtual disks to which the rule applies and click **OK**.
11. Click **OK** to save the rule.

Override VMDK Affinity Rules

VMDK affinity rules indicate that all virtual disks in a datastore cluster that are associated with a particular virtual machine are located on the same datastore in the datastore cluster. The rules apply to individual datastore clusters.

VMDK affinity rules are activated by default for all virtual machines that are in a datastore cluster. You can override the default setting for the datastore cluster or for individual virtual machines.

Virtual machines that are subject to VMDK affinity rules have the following behavior:

- Storage DRS places the virtual machine's virtual disks according to the rule.
- Storage DRS migrates the virtual disks using vMotion according to the rule, even if the migration is for a mandatory reason such as putting a datastore in maintenance mode.
- If the virtual machine's virtual disk violates the rule, Storage DRS makes migration recommendations to correct the error or reports the violation as a fault if it cannot make a recommendation that will correct the error.

When you add a datastore to a datastore cluster that is activated for Storage DRS, the VMDK affinity rule is deactivated for any virtual machine that has virtual disks on that datastore if it also has virtual disks on other datastores.

1. Browse to the datastore cluster in the vSphere Client.
2. Click the **Configure** tab and click **Configuration**.
3. Select **VM Overrides**.
4. Click **Add**.
5. Use the **+** button to select virtual machines.
6. Click the **Keep VMDKs together** drop-down menu and select **No**.
7. Click **OK**.

Clear Storage DRS Statistics

To diagnose problems with Storage DRS, you can clear Storage DRS statistics before you manually run Storage DRS.

Activate Storage DRS for the datastore cluster.

Important: When you activate the option to clear Storage DRS statistics, statistics are cleared every time Storage DRS runs until you deactivate the option. Always deactivate the option after you diagnose the Storage DRS problem.

1. Activate the `ClearIoStatsOnSdrsRun` option.
 - a) Browse to the datastore cluster in the vSphere Client.
 - b) Click the **Configuration** tab and click **Services**.
 - c) Select **vSphere DRS** and click **Edit**.
 - d) Expand **Advanced Options** and click **Add**.
 - e) In the Option column, type `ClearIoStatsOnSdrsRun`.
 - f) In the corresponding Value text box, type 1.
 - g) Click **OK**.
2. Run Storage DRS on the datastore cluster.

The current Storage DRS statistics for all datastores and virtual disks in all datastore clusters in the vSphere Client inventory are cleared, but no new statistics are collected.
3. Change the `ClearIoStatsOnSdrsRun` flag value to 0 to deactivate it.
4. Run Storage DRS again.

Storage DRS runs normally. Allow several hours for the new setting to take effect.

Storage vMotion Compatibility with Datastore Clusters

A datastore cluster has certain vSphere Storage vMotion® requirements.

- The host must be running a version of ESXi that supports Storage vMotion.
- The host must have write access to both the source datastore and the destination datastore.
- The host must have enough free memory resources to accommodate Storage vMotion.
- The destination datastore must have sufficient disk space.
- The destination datastore must not be in maintenance mode or entering maintenance mode.

Using NUMA Systems with ESXi

ESXi supports memory access optimization for Intel and AMD Opteron processors in server architectures that support NUMA (non-uniform memory access).

After you understand how ESXi NUMA scheduling is performed and how the VMware NUMA algorithms work, you can specify NUMA controls to optimize the performance of your virtual machines.

What is NUMA?

NUMA systems are advanced server platforms with more than one system bus. They can harness large numbers of processors in a single system image with superior price to performance ratios.

A multi-gigahertz CPU needs to be supplied with a large amount of memory bandwidth to use its processing power effectively. Even a single CPU running a memory-intensive workload, such as a scientific computing application, can be constrained by memory bandwidth.

This problem is amplified on symmetric multiprocessing (SMP) systems, where many processors must compete for bandwidth on the same system bus. Some high-end systems often try to solve this problem by building a high-speed data bus. However, such a solution is expensive and limited in scalability.

NUMA is an alternative approach that links several small, cost-effective nodes using a high-performance connection. Each node contains processors and memory, much like a small SMP system. However, an advanced memory controller allows a node to use memory on all other nodes, creating a single system image. When a processor accesses memory that does

not lie within its own node (remote memory), the data must be transferred over the NUMA connection, which is slower than accessing local memory. Memory access times are not uniform and depend on the location of the memory and the node from which it is accessed, as the technology's name implies.

Challenges for Operating Systems

Because a NUMA architecture provides a single system image, it can often run an operating system with no special optimizations.

The high latency of remote memory accesses can leave the processors under-utilized, constantly waiting for data to be transferred to the local node, and the NUMA connection can become a bottleneck for applications with high-memory bandwidth demands.

Furthermore, performance on such a system can be highly variable. It varies, for example, if an application has memory located locally on one benchmarking run, but a subsequent run happens to place all of that memory on a remote node. This phenomenon can make capacity planning difficult.

Some high-end UNIX systems provide support for NUMA optimizations in their compilers and programming libraries. This support requires software developers to tune and recompile their programs for optimal performance. Optimizations for one system are not guaranteed to work well on the next generation of the same system. Other systems have allowed an administrator to explicitly decide on the node on which an application should run. While this might be acceptable for certain applications that demand 100 percent of their memory to be local, it creates an administrative burden and can lead to imbalance between nodes when workloads change.

Ideally, the system software provides transparent NUMA support, so that applications can benefit immediately without modifications. The system should maximize the use of local memory and schedule programs intelligently without requiring constant administrator intervention. Finally, it must respond well to changing conditions without compromising fairness or performance.

How ESXi NUMA Scheduling Works

ESXi uses a sophisticated NUMA scheduler to dynamically balance processor load and memory locality or processor load balance.

1. Each virtual machine managed by the NUMA scheduler is assigned a home node. A home node is one of the system's NUMA nodes containing processors and local memory, as indicated by the System Resource Allocation Table (SRAT).
2. When memory is allocated to a virtual machine, the ESXi host preferentially allocates it from the home node. The virtual CPUs of the virtual machine are constrained to run on the home node to maximize memory locality.
3. The NUMA scheduler can dynamically change a virtual machine's home node to respond to changes in system load. The scheduler might migrate a virtual machine to a new home node to reduce processor load imbalance. Because this might cause more of its memory to be remote, the scheduler might migrate the virtual machine's memory dynamically to its new home node to improve memory locality. The NUMA scheduler might also swap virtual machines between nodes when this improves overall memory locality.

Some virtual machines are not managed by the ESXi NUMA scheduler. For example, if you manually set the processor or memory affinity for a virtual machine, the NUMA scheduler might not be able to manage this virtual machine. Virtual machines that are not managed by the NUMA scheduler still run correctly. However, they don't benefit from ESXi NUMA optimizations.

The NUMA scheduling and memory placement policies in ESXi can manage all virtual machines transparently, so that administrators do not need to address the complexity of balancing virtual machines between nodes explicitly.

The optimizations work seamlessly regardless of the type of guest operating system. ESXi provides NUMA support even to virtual machines that do not support NUMA hardware, such as Windows NT 4.0. As a result, you can take advantage of new hardware even with legacy operating systems.

A virtual machine that has more virtual processors than the number of physical processor cores available on a single hardware node can be managed automatically. The NUMA scheduler accommodates such a virtual machine by having it span NUMA nodes. That is, it is split up as multiple NUMA clients, each of which is assigned to a node and then managed by the scheduler as a normal, non-spanning client. This can improve the performance of certain memory-intensive workloads with high locality. For information on configuring the behavior of this feature, see [Advanced Virtual Machine Attributes](#).

ESXi includes support for exposing virtual NUMA topology to guest operating systems. For more information about virtual NUMA control, see [Using Virtual NUMA](#).

VMware NUMA Optimization Algorithms and Settings

This section describes the algorithms and settings used by ESXi to maximize application performance while still maintaining resource guarantees.

Home Nodes and Initial Placement

When a virtual machine is powered on, ESXi assigns it a home node. A virtual machine runs only on processors within its home node, and its newly allocated memory comes from the home node as well.

Unless a virtual machine's home node changes, it uses only local memory, avoiding the performance penalties associated with remote memory accesses to other NUMA nodes.

When a virtual machine is powered on, it is assigned an initial home node so that the overall CPU and memory load among NUMA nodes remains balanced. Because internode latencies in a large NUMA system can vary greatly, ESXi determines these internode latencies at boot time and uses this information when initially placing virtual machines that are wider than a single NUMA node. These wide virtual machines are placed on NUMA nodes that are close to each other for lowest memory access latencies.

Initial placement-only approaches are usually sufficient for systems that run only a single workload, such as a benchmarking configuration that remains unchanged as long as the system is running. However, this approach is unable to guarantee good performance and fairness for a datacenter-class system that supports changing workloads. Therefore, in addition to initial placement, ESXi does dynamic migration of virtual CPUs and memory between NUMA nodes for improving CPU balance and increasing memory locality.

Dynamic Load Balancing and Page Migration

ESXi combines the traditional initial placement approach with a dynamic rebalancing algorithm. Periodically (every two seconds by default), the system examines the loads of the various nodes and determines if it should rebalance the load by moving a virtual machine from one node to another.

This calculation takes into account the resource settings for virtual machines and resource pools to improve performance without violating fairness or resource entitlements.

The rebalancer selects an appropriate virtual machine and changes its home node to the least loaded node. When it can, the rebalancer moves a virtual machine that already has some memory located on the destination node. From that point on (unless it is moved again), the virtual machine allocates memory on its new home node and it runs only on processors within the new home node.

Rebalancing is an effective solution to maintain fairness and ensure that all nodes are fully used. The rebalancer might need to move a virtual machine to a node on which it has allocated little or no memory. In this case, the virtual machine incurs a performance penalty associated with a large number of remote memory accesses. ESXi can eliminate this penalty by transparently migrating memory from the virtual machine's original node to its new home node:

1. The system selects a page (4KB of contiguous memory) on the original node and copies its data to a page in the destination node.

2. The system uses the virtual machine monitor layer and the processor's memory management hardware to seamlessly remap the virtual machine's view of memory, so that it uses the page on the destination node for all further references, eliminating the penalty of remote memory access.

When a virtual machine moves to a new node, the ESXi host immediately begins to migrate its memory in this fashion. It manages the rate to avoid overtaxing the system, particularly when the virtual machine has little remote memory remaining or when the destination node has little free memory available. The memory migration algorithm also ensures that the ESXi host does not move memory needlessly if a virtual machine is moved to a new node for only a short period.

When initial placement, dynamic rebalancing, and intelligent memory migration work in conjunction, they ensure good memory performance on NUMA systems, even in the presence of changing workloads. When a major workload change occurs, for instance when new virtual machines are started, the system takes time to readjust, migrating virtual machines and memory to new locations. After a short period, typically seconds or minutes, the system completes its readjustments and reaches a steady state.

Transparent Page Sharing Optimized for NUMA

Many ESXi workloads present opportunities for sharing memory across virtual machines.

You might have multiple virtual machines running instances of the same guest operating system, have the same applications or components loaded, or contain common data. In such cases, ESXi systems use a proprietary transparent page-sharing technique to eliminate redundant copies of memory pages. With memory sharing, a workload running in virtual machines often consumes less memory than it might when running on physical machines. As a result, higher levels of overcommitment can be supported efficiently.

Transparent page sharing for ESXi systems has also been optimized for use on NUMA systems. On NUMA systems, pages are shared per-node, so each NUMA node has its own local copy of heavily shared pages. When virtual machines use shared pages, they do not need access to remote memory.

Note: This default behavior is the same in all previous versions of ESX and ESXi.

Resource Management in NUMA Architectures

You can perform resource management with different types of NUMA architecture.

With the proliferation of highly multicore systems, NUMA architectures are becoming more popular as these architectures allow better performance scaling of memory intensive workloads. All modern Intel and AMD systems have NUMA support built into the processors. Additionally, there are traditional NUMA systems like the IBM Enterprise X-Architecture that extend Intel and AMD processors with NUMA behavior with specialized chipset support.

Typically, you can use BIOS settings to activate and deactivate NUMA behavior. For example, in AMD Opteron-based HP ProLiant servers, NUMA can be deactivated by activating node interleaving in the BIOS. If NUMA is activated, the BIOS builds a system resource allocation table (SRAT) which ESXi uses to generate the NUMA information used in optimizations. For scheduling fairness, NUMA optimizations are not activated for systems with too few cores per NUMA node or too few cores overall. You can modify the `numa.rebalancecorestotal` and `numa.rebalancecoresnode` options to change this behavior.

Using Virtual NUMA

vSphere includes support for exposing virtual NUMA topology to guest operating systems, which can improve performance by facilitating guest operating system and application NUMA optimizations.

Virtual NUMA topology is available to virtual machines and is activated by default when the number of virtual CPUs is greater than eight. You can also manually influence virtual NUMA topology using advanced configuration options.

The first time a virtual NUMA activated virtual machine is powered on, its virtual NUMA topology is based on the NUMA topology of the underlying physical host. Once a virtual machines virtual NUMA topology is initialized, it does not change unless the number of vCPUs in that virtual machine is changed.

The virtual NUMA topology does not consider the memory configured to a virtual machine. The virtual NUMA topology is not influenced by the number of virtual sockets and number of cores per socket for a virtual machine.

If the virtual NUMA topology needs to be overridden, see [Virtual NUMA Controls](#).

Note: Enabling CPU HotAdd will deactivate virtual NUMA. See <https://kb.vmware.com/kb/2040375>.

Virtual Topology in ESXi 8.0

ESXi 8.0 includes an enhanced virtual topology feature.

The virtual topology of a VM enables optimization within GOS for placement and load balancing. Selecting an accurate virtual topology that aligns with the underlying physical topology of the host where the VM is running is crucial for application performance.

ESXi 8.0 automatically selects optimal `coresPerSocket` for a VM and optimal virtual L3 size. It also includes a new virtual motherboard layout to expose NUMA for virtual devices and vNUMA topology when CPU hotplug is enabled.

Note: Enhanced virtual topology is only available on ESXi 8.0. VMs must have hardware version 20 or later to use this feature.

1. To manually configure VM topology, first browse to the VM.
2. Select **VM Options**. Under **CPU Topology** you can adjust **Cores per socket** and **NUMA nodes**.
In order to enable a hot-added CPU to a new NUMA node, add `numa.allowHotadd` under advanced configuration options. Then you can manually add a NUMA configuration.
Note: By default, a VM with CPU hotplug enabled will enforce single NUMA node topology. Any hot-added CPUs go to the single NUMA node.
3. Under **Device Assignment** you can also assign a device to a virtual NUMA node, or leave it unassigned.

This newly configured topology is displayed as **Manual** in the existing VM topology section. Without manual configuration, this tab is displayed as **assigned at power on**.

Virtual NUMA Controls

For virtual machines with disproportionately large memory consumption, you can use advanced options to override the default virtual CPU settings.

You can add these advanced options to the virtual machine configuration file.

Table 366: Advanced Options for Virtual NUMA Controls

Option	Description	Default Value
<code>cpuid.coresPerSocket</code>	Determines the number of virtual cores per virtual CPU socket. This option does not affect the virtual NUMA topology unless <code>numa.vcpu.followcorespersocket</code> is configured. Note: ESXi 8.0 automatically selects optimal <code>coresPerSocket</code> for a VM, the default value is displayed as 0.	1
<code>numa.vcpu.maxPerVirtualNode</code>	Determines the number of virtual NUMA nodes by splitting the total vCPU count evenly with this value as its divisor.	8

Option	Description	Default Value
<code>numa.autosize.once</code>	When you create a virtual machine template with these settings, the settings remain the same every time you then power on the virtual machine with the default value TRUE. If the value is set to FALSE, the virtual NUMA topology is updated every time it is powered on. The virtual NUMA topology is reevaluated when the configured number of virtual CPUs on the virtual machine is modified at any time.	FALSE
<code>numa.vcpu.min</code>	The minimum number of virtual CPUs in a virtual machine that are required to generate a virtual NUMA topology. A virtual machine is always UMA when its size is smaller than <code>numa.vcpu.min</code>	9
<code>numa.vcpu.followcorespersocket</code>	If set to 1, reverts to the old behavior of virtual NUMA node sizing being tied to <code>cpuid.coresPerSocket</code> .	0
<code>numa.allowHotadd</code>	In order to activate the capability of hot-adding CPU to a new NUMA node, add <code>numa.allowHotadd</code> under advanced configuration options. Then you can manually add a NUMA configuration while CPU hotadd is activated.	FALSE
<code>numa.vcpu.coresPerNode</code>	VMX parameter to configure size of virtual NUMA node, which is interpreted from the UI reconfiguration. This parameter is only effective for HWv20. Defaults to 0, which indicates ESXi automatically selects vNUMA size. Note: If this option conflicts with <code>numa.vcpu.maxPerVirtualNode</code> , the VM cannot power on.	0
<code>vcpu.hotadd</code>	When this option is TRUE, virtual NUMA is deactivated. The VM always sees one virtual NUMA node when VM is capable of CPU hotplug.	
<code>llc.multiLLCPerSocket</code>	VM exposes truthful vLLC on AMD Epyc when this option is TRUE. The exposed vLLC size can be different than virtual socket size.	FALSE
<code>llc.size.vcpu</code>	Manually configured number of vcpus for vLLC on AMD Epyc. The value is ignored if it is incompatible with other settings of the VM.	
<code>chipset.motherboardLayout</code>	Type of virtual motherboard used by this VM. It can only have one of the following two values: <code>acpi</code> : New motherboard layout starting from HWv 20. <code>i440bx</code> : Legacy motherboard layout	
<code>cpuid.coresPerSocket.cookie</code>	This is a vmx entry generated by ESXi to store the automatically generated <code>coresPerSocket</code> value. It is there to ensure vMotion consistency. Do not change or remove it manually.	

Specifying NUMA Controls

If you have applications that use a lot of memory or have a small number of virtual machines, you might want to optimize performance by specifying virtual machine CPU and memory placement explicitly.

Specifying controls is useful if a virtual machine runs a memory-intensive workload, such as an in-memory database or a scientific computing application with a large data set. You might also want to optimize NUMA placements manually if the

system workload is known to be simple and unchanging. For example, an eight-processor system running eight virtual machines with similar workloads is easy to optimize explicitly.

Note: In most situations, the ESXi host's automatic NUMA optimizations result in good performance.

ESXi provides three sets of controls for NUMA placement, so that administrators can control memory and processor placement of a virtual machine.

You can specify the following options.

NUMA Node Affinity	When you set this option, NUMA can schedule a virtual machine only on the nodes specified in the affinity.
CPU Affinity	When you set this option, a virtual machine uses only the processors specified in the affinity.
Memory Affinity	When you set this option, the server allocates memory only on the specified nodes.

A virtual machine is still managed by NUMA when you specify NUMA node affinity, but its virtual CPUs can be scheduled only on the nodes specified in the NUMA node affinity. Likewise, memory can be obtained only from the nodes specified in the NUMA node affinity. When you specify CPU or memory affinities, a virtual machine ceases to be managed by NUMA. NUMA management of these virtual machines is effective when you remove the CPU and memory affinity constraints.

Manual NUMA placement might interfere with ESXi resource management algorithms, which distribute processor resources fairly across a system. For example, if you manually place 10 virtual machines with processor-intensive workloads on one node, and manually place only 2 virtual machines on another node, it is impossible for the system to give all 12 virtual machines equal shares of systems resources.

Associate Virtual Machines with Specific Processors

You might be able to improve the performance of the applications on a virtual machine by pinning its virtual CPUs to fixed processors. This allows you to prevent the virtual CPUs from migrating across NUMA nodes.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and click **Edit Settings**.
3. Select the **Virtual Hardware** tab, and expand **CPU**.
4. Under Scheduling Affinity, set the CPU affinity to the preferred processors.

Note: You must manually select all processors in the NUMA node. CPU affinity is specified on a per-processor, not on a per-node, basis.

Associate Memory Allocations with Specific NUMA Nodes Using Memory Affinity

You can specify that all future memory allocations on a virtual machine use pages associated with specific NUMA nodes (also known as manual memory affinity).

Note: Specify nodes to be used for future memory allocations only if you have also specified CPU affinity. If you make manual changes only to the memory affinity settings, automatic NUMA rebalancing does not work properly.

1. Browse to the virtual machine in the vSphere Client.
2. Click the **Configure** tab.
3. Click **Settings**, and click **VM Hardware**.
4. Click **Edit**.
5. Select the **Virtual Hardware** tab, and expand **Memory**.
6. Under NUMA Memory Affinity, set memory affinity.

Binding a Virtual Machine to a Single NUMA Node

The following example illustrates manually binding the last four physical CPUs to a single NUMA node for a two-way virtual machine on an eight-way server.

The CPUs (for example, 4, 5, 6, and 7) are the physical CPU numbers.

1. In the vSphere Client, right-click the virtual machine and select **Edit Settings**.
2. Click the **Advanced Parameters** tab.
3. In the vSphere Client, turn on CPU affinity for processors 4, 5, 6, and 7.

Then, you want this virtual machine to run only on node 1.

1. In the vSphere Client inventory panel, select the virtual machine and select **Edit Settings**.
2. Click the **Advanced Parameters** tab.
3. In the vSphere Client, set memory affinity for the NUMA node to 1.

Completing these two tasks ensures that the virtual machine runs only on NUMA node 1 and, when possible, allocates memory from the same node.

Associate Virtual Machines with Specified NUMA Nodes

When you associate a NUMA node with a virtual machine to specify NUMA node affinity, you constrain the set of NUMA nodes on which ESXi can schedule a virtual machine's virtual CPU and memory.

Note: When you constrain NUMA node affinities, you might interfere with the ability of the ESXi NUMA scheduler to rebalance virtual machines across NUMA nodes for fairness. Specify NUMA node affinity only after you consider the rebalancing issues.

1. Browse to the VM in the vSphere Client.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click the **Advanced Parameters** tab.
4. Optional: To add a parameter, enter an attribute name and value for the parameter and click **Add**.
5. Optional: To change a parameter, enter a new value in the **Value** text box for that parameter and click the check mark to confirm.
6.
 - To specify NUMA node for the virtual machine, in the Name column, enter `numa.nodeAffinity`.
 - To specify NUMA node for a specific Virtual NUMA node on the virtual machine, in the Name column, enter `sched.nodeX.affinity`, where X is the Virtual NUMA node number. For example, `sched.node0.affinity` specifies Virtual NUMA node 0 on the virtual machine.
7. In the Value column, enter the NUMA nodes where the virtual machine or the virtual NUMA node can be scheduled. Use a comma-separated list for multiple nodes. For example, enter `0,1` to constrain the virtual machine resource scheduling to NUMA nodes 0 and 1.
8. Click **OK**.

Advanced Attributes in vSphere

You can set advanced attributes for hosts or individual virtual machines to help you customize resource management.

In most cases, adjusting the basic resource allocation settings (reservation, limit, shares) or accepting default settings results in appropriate resource allocation. However, you can use advanced attributes to customize resource management for a host or a specific virtual machine.

How Do You Set Advanced Host Attributes in the vSphere Client

Learn how to set advanced attributes for an ESXi host.

CAUTION: Changing advanced options is considered unsupported. Typically, the default settings produce the optimum result. Change the advanced options only when you get specific instructions from VMware technical support or a knowledge base article.

1. Browse to the host in the vSphere Client.
2. Click the **Configure** tab.
3. Under **System**, click **Advanced System Settings**.
4. Click the **Edit** button.
5. Find the appropriate item and change the value.
6. Click **OK**.

Advanced Memory Attributes

You can use the advanced memory attributes to customize memory resource usage.

Table 367: Advanced Memory Attributes

Attribute	Description	Default
Mem.ShareForceSalting	Mem.ShareForceSalting 0: Inter-virtual machine Transparent Page Sharing (TPS) behavior is still retained. The value of VMX option <code>sched.mem.pshare.salt</code> is ignored even if present. Mem.ShareForceSalting 1: By default the salt value is taken from <code>sched.mem.pshare.salt</code> . If not specified, it falls back to old TPS (inter-VM) behavior by considering salt values for the virtual machine as 0. Mem.ShareForceSalting 2: By default the salt value is taken from <code>sched.mem.pshare.salt</code> if present, or <code>vc.uuid</code> . If it does not exist, then the page sharing algorithm generates random and unique value for salting per virtual machine, which is not configurable by users.	2
Mem.SamplePeriod	Specifies the periodic time interval, measured in seconds of the virtual machine's execution time, over which memory activity is monitored to estimate working set sizes.	60
Mem.BalancePeriod	Specifies the periodic time interval, in seconds, for automatic memory reallocations. Significant changes in the amount of free memory also trigger reallocations.	15
Mem.IdleTax	Specifies the idle memory tax rate, as a percentage. This tax effectively charges virtual machines more for idle memory than for memory they are actively using. A tax rate of 0 percent defines an allocation policy that ignores working sets and allocates memory strictly based on shares. A high tax rate results in an allocation policy that allows idle memory to be reallocated away from virtual machines that are unproductively hoarding it.	75
Mem.ShareScanGHz	Specifies the maximum amount of memory pages to scan (per second) for page sharing opportunities for each GHz of available host CPU resource. For example, defaults to 4 MB/sec per 1 GHz.	4
Mem.ShareScanTime	Specifies the time, in minutes, within which an entire virtual machine is scanned for page sharing opportunities. Defaults to 60 minutes.	60

Attribute	Description	Default
Mem.CtlMaxPercent	Limits the maximum amount of memory reclaimed from any virtual machine using the memory balloon driver (<code>vmmemctl</code>), based on a percentage of its configured memory size. Specify 0 to deactivate reclamation for all virtual machines.	65
Mem.AllocGuestLargePage	Activates backing of guest large pages with host large pages. Reduces TLB misses and improves performance in server workloads that use guest large pages. 0=deactivate.	1
Mem.AllocUsePSharePool and Mem.AllocUseGuestPool	Reduces memory fragmentation by improving the probability of backing guest large pages with host large pages. If host memory is fragmented, the availability of host large pages is reduced. 0 = deactivate.	15
Mem.MemZipEnable	Activates memory compression for the host. 0 = deactivate.	1
Mem.MemZipMaxPct	Specifies the maximum size of the compression cache in terms of the maximum percentage of each virtual machine's memory that can be stored as compressed memory.	10
LPage.LPageDefragEnable	Activates large page defragmentation. 0 = deactivate.	1
LPage.LPageDefragRateVM	Maximum number of large page defragmentation attempts per second per virtual machine. Accepted values range from 1 to 1024.	32
LPage.LPageDefragRateTotal	Maximum number of large page defragmentation attempts per second. Accepted values range from 1 to 10240.	256
LPage.LPageAlwaysTryForNPT	Try to allocate large pages for nested page tables (called 'RVI' by AMD or 'EPT' by Intel). If you activate this option, all guest memory is backed with large pages in machines that use nested page tables (for example, AMD Barcelona). If NPT is not available, only some portion of guest memory is backed with large pages. 0= deactivate.	1

Advanced NUMA Attributes

You can use the advanced NUMA attributes to customize NUMA usage.

Table 368: Advanced NUMA Attributes

Attribute	Description	Default
Numa.RebalancePeriod	Controls the frequency of rebalance periods, specified in milliseconds. More frequent rebalancing can increase CPU overheads, particularly on machines with a large number of running virtual machines. More frequent rebalancing can also improve fairness.	2000
Numa.MigImbalanceThreshold	The NUMA rebalancer computes the CPU imbalance between nodes, accounting for the difference between each virtual machine's CPU time entitlement and its actual consumption. This option controls the minimum load imbalance between nodes needed to trigger a virtual machine migration, in percent.	10
Numa.RebalanceEnable	Activate NUMA rebalancing and scheduling. Set this option to 0 to deactivate all NUMA rebalancing and initial placement of virtual machines, effectively deactivating the NUMA scheduling system.	1

Attribute	Description	Default
Numa.RebalanceCoresTotal	Specifies the minimum number of total processor cores on the host required to activate the NUMA rebalancer.	4
Numa.RebalanceCoresNode	Specifies the minimum number of processor cores per node required to activate the NUMA rebalancer. This option and Numa.RebalanceCoresTotal are useful when deactivating NUMA rebalancing on small NUMA configurations (for example, two-way Opteron hosts), where the small number of total or per-node processors can compromise scheduling fairness when you activate NUMA rebalancing.	2
Numa.AutoMemAffinity	Automatically set memory affinity for virtual machines that have CPU affinity set.	1
Numa.PageMigEnable	Automatically migrate pages between NUMA nodes to improve memory locality. Page migration rates set manually are still in effect.	1

Set Advanced Virtual Machine Attributes

You can set advanced attributes for a virtual machine.

Power off the VM before setting advanced attributes.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and select **Edit Settings**.
3. Click the **Advanced Parameters** tab.
4. Optional: To add a parameter, enter an attribute name and value for the parameter and click **Add**.
5. Optional: To change a parameter, enter a new value in the **Value** text box for that parameter and click the check mark to confirm.
6. Click **OK**.

Advanced Virtual Machine Attributes

You can use the advanced virtual machine attributes to customize virtual machine configuration.

Table 369: Advanced Virtual Machine Attributes

Attribute	Description	Default
sched.mem.maxmemctl	Maximum amount of memory reclaimed from the selected virtual machine by ballooning, in megabytes (MB). If the ESXi host needs to reclaim additional memory, it is forced to swap. Swapping is less desirable than ballooning.	-1 (Unlimited)
sched.mem.pshare.enable	Enables memory sharing for a selected virtual machine.	True

Attribute	Description	Default
	This boolean value defaults to True. If you set it to False for a virtual machine, this turns off memory sharing.	
sched.mem.pshare.salt	A salt value is a configurable VMX option for each virtual machine. If this option is not present in the virtual machine's VMX file, then the value of <code>vc.uuid vmx</code> option is taken as the default value. Since the <code>vc.uuid</code> is unique to each virtual machine, by default transparent page sharing happens only among the pages belonging to a particular virtual machine (intra-VM). If a group of virtual machines are considered trustworthy, it is possible to share pages among them by setting a common salt value for all of those virtual machines (inter-VM).	user configurable
sched.swap.persist	Specifies whether the virtual machine's swap files should persist or be deleted when the virtual machine is powered off. By default, the system creates the swap file for a virtual machine when the virtual machine is powered on, and deletes the swap file when the virtual machine is powered off.	False
sched.swap.dir	Directory location of the virtual machine's swap file. Defaults to the virtual machine's working directory, that is, the directory that contains its configuration file. This directory must remain on a host that is accessible to the virtual machine. If you move the virtual machine (or any clones created from it), you might need to reset this attribute.	Equals <code>workingDir</code>

Advanced Virtual NUMA Attributes

You can use the advanced virtual NUMA attributes to customize virtual NUMA usage.

Table 370: Advanced NUMA Attributes

Attribute	Description	Default
cpuid.coresPerSocket	Determines the number of virtual cores per virtual CPU socket. If the value is greater than 1, also determines the size of virtual NUMA nodes if a virtual machine has a virtual NUMA topology. You can set this option if you know the exact virtual NUMA topology for each physical host.	1
numa.autosize	When you set this option, the virtual NUMA topology has the same number of virtual CPUs per virtual node as there are cores on each physical node.	FALSE

Attribute	Description	Default
numa.autosize.once	When you create a virtual machine template with these settings, the settings are guaranteed to remain the same every time you subsequently power on the virtual machine. The virtual NUMA topology will be reevaluated if the configured number of virtual CPUs on the virtual machine is modified.	TRUE
numa.vcpu.maxPerVirtualNode	If <code>cpuid.coresPerSocket</code> is too restrictive as a power of two, you can set <code>numa.vcpu.maxPerVirtualNode</code> directly. In this case, do not set <code>cpuid.coresPerSocket</code> .	8
numa.vcpu.min	Minimum number of virtual CPUs in a virtual machine that are required in order to generate a virtual NUMA topology.	9
numa.vcpu.maxPerMachineNode	Maximum number of virtual CPUs that belong to the same virtual machine that can be scheduled on a NUMA node at the same time. Use this attribute to ensure maximum bandwidth, by forcing different NUMA clients on different NUMA nodes.	Number of cores per node on the physical host where a virtual machine is running.
numa.vcpu.maxPerClient	Maximum number of virtual CPUs in a NUMA client. A client is a group of virtual CPUs that are NUMA-managed as a single entity. By default, each virtual NUMA node is a NUMA client, but if a virtual NUMA node is larger than a physical NUMA node, a single virtual NUMA node can be backed by multiple NUMA clients.	Equals <code>numa.vcpu.maxPerMachineNode</code>
numa.nodeAffinity	Constrains the set of NUMA nodes on which a virtual machine's virtual CPU and memory can be scheduled. Note: When you constrain NUMA node affinities, you might interfere with the ability of the NUMA scheduler to rebalance virtual machines across NUMA nodes for fairness. Specify NUMA node affinity only after you consider the rebalancing issues.	
numa.mem.interleave	Specifies whether the memory allocated to a virtual machine is statically interleaved across all the NUMA nodes on which its constituent NUMA clients are running and there is no virtual NUMA topology exposed.	True

Latency Sensitivity

You can adjust the latency sensitivity of a virtual machine to optimize the scheduling delay for latency sensitive applications.

ESXi is optimized to deliver high throughput. You can optimize your virtual machine to meet the low latency requirement of latency sensitive applications. Examples of latency sensitive applications are VOIP or media player applications, or applications that require frequent access to the mouse or keyboard devices.

Adjust Latency Sensitivity

You can adjust the latency sensitivity of a virtual machine.

ESXi requires full CPU reservation to power on a VM with hardware version 14 when Latency Sensitivity is set to **high**.

1. Browse to the virtual machine in the vSphere Client.
 - a) To find a virtual machine, select a data center, folder, cluster, resource pool, or host.
 - b) Click the **VMs** tab.
2. Right-click the virtual machine and click **Edit Settings**.
3. Click **VM Options** and click **Advanced**.
4. Select a setting from the **Latency Sensitivity** drop-down menu.
5. Click **OK**.

Virtual Hyperthreading support for VMs

Virtual Hyperthreading (vHT) is supported for VMs.

In vSphere 8.0, vHT is supported for VMs. vHT is deactivated by default and can be activated under the latency sensitivity settings for each VM. vHT supports a maximum HT size of 2.

vHT is an extension of latency sensitivity high feature. Applications which benefit from hyperthreading awareness can see performance improvements with latency sensitivity high and vHT activated. The performance gain might come from sufficient resource reservation and that the VM has exclusive physical CPUs.

Without vHT activated on ESXi, each virtual CPU (vCPU) is equivalent to a single non-hyperthreaded core available to the guest operating system. With vHT activated, each guest vCPU is treated as a single hyperthread of a virtual core (vCore).

Virtual hyperthreads of the same vCore occupy the same physical core. As a result, vCPUs of the VM can share the same core as opposed to using multiple cores on VMs with latency sensitivity high that have vHT deactivated.

ESXi hosts and VM running older hardware versions cannot use this feature.

vHT Full CPU Reservation

You can use a formula to calculate the full CPU reservation for vHT.

For a low latency VM without vHT, each vCPU of the VM has exclusively affinity to a thread of the physical core. For hosts with Hyperthreading activated, the partner hyperthread has exclusive affinity to an idle world. Each vCPU of a low latency VM is allocated a dedicated physical core.

The CPU reservation for low latency VM is computed as:

```
Low latency VM (without vHT)
CPU minimum reservation = numVcpus * cpuFrequency
```

However, when vHT is activated for a VM, each hypervisor of the physical core is shared between multiple vCPUs of the VM where each hypervisor has exclusive affinity to a vCPU of the VM. This means a core with numSMT physical hyperthreads is shared by numSMT virtual threads. In this case, the CPU reservation requirement would be computed as:

```
Low latency VM (with vHT)
CPU minimum reservation = (numVcpus / numSMT) * cpuFrequency
```

Table 371: Example of booting a low latency VM with 20 vCPUs on a host with CPU frequency as 2 GHz

	numSMT = 1(without vHT)	numSMT = 2(with vHT)
numVcpus	20	20
number of Physical cores	20	10 (each core shared by 2 vCPUs)

	numSMT = 1(without vHT)	numSMT = 2(with vHT)
Required minimum CPU reservation	20 * 2.0 GHz = 40 GHz	(20/2) * 2.0 GHz = 20 GHz

Activate vHT for VMs

vHT is supported with ESXi 8.0, but it is deactivated by default. You can activate vHT under the latency sensitivity settings for each VM.

When vHT is activated, CPU and Memory must be set to the full reservation. A warning is displayed if you set the reservation lower.

1. In vSphere, select the VM.
2. Select **Actions** and click **Edit Settings**.
3. Under Latency Sensitivity, click the drop-down menu and select **High with Hyperthreading**.
4. Click **OK**.

vHT is activated.

Reliable Memory

ESXi supports reliable memory.

Some systems have reliable memory, which is a part of memory that is less likely to have hardware memory errors than other parts of the memory in the system. If the hardware exposes information about the different levels of reliability, ESXi might be able to achieve higher system reliability.

View Reliable Memory

You can view whether or not the license permits reliable memory.

1. Browse to the host in the vSphere Client.
2. Click the **Configure** tab and click **System**.
3. Select **Licensing**.
4. Under **Licensed Features** verify Reliable Memory is displayed.

You can look up how much memory is considered reliable by using the ESXCLI `hardware memory get` command.

Backing Guest vRAM with 1GB Pages

vSphere ESXi provides a limited support for backing guest vRAM with 1GB pages.

In order to use 1GB pages for backing guest memory you must apply the option `sched.mem.lpage.enable1GPage = "TRUE"` for the VM. You can set this under Advanced options when you select **Edit Settings**. You can only enable 1GB pages on a VM that is powered off.

A VM with 1GB pages enabled must have full memory reservation. Otherwise, the VM will not be able to power on. All of the vRAM for VMs with 1GB pages enabled is preallocated on power-on. Since these VMs have full memory reservation they are not affected by memory reclamation and their memory consumption stays at the maximum level for the entire lifetime of the VM.

1GB page vRAM backing is opportunistic and 1GB pages are allocated on a best effort basis. This includes cases where host CPUs do not have 1GB capabilities. To maximize the chances of having guest vRAM backed with 1GB pages, we recommended to start VMs requiring 1GB pages on a freshly booted host because over time the host RAM is fragmented.

A VM with 1GB pages enabled can be migrated to a different host. However, the 1GB page size might not be allocated on the destination host in the same amount as it was on the source host. You might also see part of vRAM backed with a 1GB page on the source host is no longer backed with a 1GB page on the destination host.

The opportunistic nature of 1GB pages extends to vSphere services such as HA and DRS that might not preserve 1GB page vRAM backing. These services are not aware of 1GB capabilities of destination host and do not take 1GB memory backing into account while making placement decisions.

Fault Definitions

DRS faults indicate the reasons that prevent the generation of DRS actions (or the recommendation of those actions in manual mode).

The DRS faults are defined within this section.

Note: In this chapter, "Memory" can refer to physical RAM or Persistent Memory.

Virtual Machine is Pinned

This fault occurs when DRS cannot move a virtual machine because DRS is deactivated on it. That is, the virtual machine is "pinned" on its registered host.

Virtual Machine not Compatible with any Host

This fault occurs when DRS cannot find a host that can run the virtual machine.

This might occur, for example, if no host can satisfy the virtual machine's CPU or memory resource needs or if no host currently has network or storage access needed by the virtual machine.

To address this problem, provide a host that can meet the virtual machine's requirements.

VM/VM DRS Rule Violated when Moving to another Host

This fault occurs when more than one virtual machines running on the same host and share affinity rules with each other cannot be moved to another host.

This might occur because not all the virtual machines can vMotion off the current host. For example, one of the virtual machines in the group is DRS-deactivated.

To prevent this, check for reasons why some virtual machines in the group cannot vMotion.

Host Incompatible with Virtual Machine

This fault occurs when DRS considers migrating a virtual machine to a host, but finds that the host is incompatible with the given virtual machine.

This might occur because the target host does not have access to the network or storage connection needed by the virtual machine. Another reason this fault occurs is if the target host has a CPU that differs sufficiently from the current host so that using vMotion amongst the hosts is not supported.

To avoid this, create clusters such that all hosts are configured consistently and vMotion is compatible amongst the hosts.

Another reason the host is incompatible with the virtual machine is that there is a required VM/Host DRS rule in place that instructs DRS to never place this virtual machine on this host.

Host Has Virtual Machine That Violates VM/VM DRS Rules

This fault occurs when the virtual machine, when powered on or moved by starting vMotion, might violate a VM/VM DRS rule.

The virtual machine can still be manually powered on or moved with vMotion, but vCenter Server cannot automatically do so.

Host has Insufficient Capacity for Virtual Machine

This fault occurs when the host does not have enough CPU or memory capacity for running the virtual machine.

Host in Incorrect State

This fault occurs when the host is entering maintenance or standby state when needed for DRS action to occur.

To address this fault, cancel the request for the host to enter standby or maintenance mode.

Host Has Insufficient Number of Physical CPUs for Virtual Machine

This fault occurs when the host hardware does not enough CPUs (hyperthreads) to support the number of virtual CPUs in the virtual machine.

Host has Insufficient Capacity for Each Virtual Machine CPU

This fault occurs when the host does not have enough CPU capacity for running the virtual machine.

The Virtual Machine Is in vMotion

This fault occurs when DRS cannot move a virtual machine because it is in vMotion.

No Active Host in Cluster

This fault occurs when the cluster in which the virtual machine is being moved does not contain any hosts that are connected and in a non-maintenance state.

This can occur, for example, when all the hosts are disconnected or in maintenance mode.

Insufficient Resources

This fault occurs when an attempted operation conflicts with a resource configuration policy.

This fault may occur, for example, if a power-on operation reserves more memory than is allocated to a resource pool.

Retry the operation after adjusting the resources to allow more memory.

Insufficient Resources to Satisfy Configured Failover Level for HA

This fault occurs when the HA configuration of CPU or memory resources reserved for failover is violated or is insufficient for the DRS operation.

This fault is reported when:

- The host is requested to enter maintenance or standby mode.
- The virtual machine violates failover when it attempts to power on.

No Compatible Hard Affinity Host

No host is available for the virtual machine that satisfies its mandatory VM/Host DRS affinity or anti-affinity rules.

No Compatible Soft Affinity Host

No host is available for the virtual machine that satisfied its preferred VM/Host DRS affinity or anti-affinity rules.

Soft Rule Violation Correction Disallowed

DRS migration threshold is set at mandatory-only.

This does not allow the generation of DRS actions to correct non-mandatory VM/Host DRS affinity rules.

Soft Rule Violation Correction Impact

Correcting the non-mandatory VM/Host DRS affinity rule does not occur because it impacts performance.

DRS Troubleshooting Information

This information describes vSphere® Distributed Resource Scheduler (DRS) problems for particular categories: cluster, host, and virtual machine problems.

Note: In this chapter, "Memory" can refer to physical RAM or Persistent Memory.

Cluster Problems

Cluster problems can prevent DRS from performing optimally or from reporting faults.

Load Imbalance on Cluster

A cluster has a load imbalance of resources.

A cluster might become unbalanced because of uneven resource demands from virtual machines and unequal capacities of hosts.

The following are possible reasons why the cluster has a load imbalance:

- The migration threshold is too high.
A higher threshold makes the cluster a more likely candidate for load imbalance.
- VM/VM or VM/Host DRS rules prevent virtual machines from being moved.
- DRS is deactivated for one or more virtual machines.
- A device is mounted to one or more virtual machines preventing DRS from moving the virtual machine in order to balance the load.
- Virtual machines are not compatible with the hosts to which DRS would move them. That is, at least one of the hosts in the cluster is incompatible for the virtual machines that would be migrated. For example, if host A's CPU is not vMotion-compatible with host B's CPU, then host A becomes incompatible for powered-on virtual machines running on host B.
- It would be more detrimental for the virtual machine's performance to move it than for it to run where it is currently located. This may occur when loads are unstable or the migration cost is high compared to the benefit gained from moving the virtual machine.
- vMotion is not activated or set up for the hosts in the cluster.

Address the problem that is causing the load imbalance.

Cluster is Yellow

The cluster is yellow due to a shortage of resources.

If the cluster does not have enough resources to satisfy the reservations of all resource pools and virtual machines, but does have enough resources to satisfy the reservations of all running virtual machines, DRS continues to run and the cluster is yellow.

A cluster can become yellow if the host resources are removed from the cluster (for example, if a host fails).

Add host resources to the cluster or reduce the resource pool reservations.

Cluster is Red Because of Inconsistent Resource Pool

A DRS cluster becomes red when it is invalid. It may become red because the resource pool tree is not internally consistent.

If the cluster resource pool tree is not internally consistent (for example, the sum of the children's reservations is greater than the parent pool's nonexpandable reservation), the cluster does not have enough resources to satisfy the reservations of all running virtual machines making the cluster red.

This can occur if vCenter Server is unavailable or if resource pool settings are changed while a virtual machine is in a failover state.

Revert the associated changes or otherwise revise the resource pool settings.

Cluster Is Red Because Failover Capacity Is Violated

A DRS cluster becomes red when it is invalid. It may become red because failover capacity is violated.

The cluster attempts to fail over virtual machines if there is host failure, but is not guaranteed to have enough resources available to fail over all virtual machines covered by the failover requirements.

If a cluster enabled for HA loses so many resources that it can no longer fulfill its failover requirements, a message appears and the cluster's status changes to red.

Review the list of configuration issues in the yellow box at the top of the cluster Summary page and address the issue that is causing the problem.

No Hosts are Powered Off When Total Cluster Load is Low

Hosts are not powered off when the total cluster load is low.

Hosts are not powered off when the total cluster load is low because extra capacity is needed for HA failover reservations.

Hosts might not be powered off for the following reasons:

- The `MinPoweredOn{Cpu|Memory}Capacity` advanced options settings need to be met.
- Virtual machines cannot be consolidated onto fewer hosts due to their resource reservations, VM/Host DRS rules, VM/VM DRS rules, not being DRS-activated, or not being compatible with the hosts having available capacity.
- Loads are unstable.
- DRS migration threshold is at the highest setting and only allows mandatory moves.
- vMotion is unable to run because it is not configured.
- DPM is deactivated on the hosts that might be powered off.
- Hosts are not compatible for virtual machines to be moved to another host.
- Host does not have Wake On LAN, IPMI, or iLO technology. Either one is required for DPM to enter a host in standby.

Address the issue that prevents hosts from being powered off when the total cluster load is low.

Hosts Are Powered-off When Total Cluster Load Is High

Hosts are powered off when total cluster load is high.

DRS determined that virtual machines can be run on fewer hosts without degrading the host or virtual machine performance. DRS is also constrained from moving the virtual machines running on the highly utilized hosts to the hosts scheduled for power-off.

The total cluster load is too high.

Reduce the cluster load.

DRS Seldom or Never Performs vMotion Migrations

DRS seldom or never performs vMotion migrations.

DRS does not perform vMotion migrations.

DRS never performs vMotion migrations when one or more of the following issues is present on the cluster.

- DRS is deactivated on the cluster.
- The hosts do not have shared storage.
- The hosts in the cluster do not contain a vMotion network.
- DRS is manual and no one has approved the migration.

DRS seldom performs vMotion when one or more of the following issues is present on the cluster:

- Loads are unstable, or vMotion takes a long time, or both. A move is not appropriate.
- DRS seldom or never migrates virtual machines.
- DRS migration threshold is set too high.

DRS moves virtual machines for the following reasons:

- Evacuation of host that a user requested enter maintenance or standby mode.
- VM/Host DRS rules or VM/VM DRS rules.
- Reservation violations.
- Load imbalance.
- Power management.

Address the issues that are causing DRS to avoid performing vMotion migrations.

Host Problems

Host problems might cause DRS to not perform as expected.

DRS Recommends Host Be Powered on to Increase Capacity When Total Cluster Load Is Low

The host must be powered on to help provide more capacity for the cluster or help hosts that are overloaded.

DRS recommends that the host is powered on to increase capacity when the total cluster load is low.

The recommendation might be made because:

- The cluster is a DRS-HA cluster. Additional powered-on hosts are needed to provide more failover capability.
- Some hosts are overloaded and virtual machines on currently powered-on hosts can be moved to hosts in standby mode to balance the load.
- The capacity is necessary to meet the `MinPoweredOn{Cpu|Memory}Capacity` advanced options.

Power on the host.

Total Cluster Load Is High

The total cluster load is high.

When the total cluster load is high, DRS does not power-on the host.

The following are possible reasons why DRS does not power-on the host:

- VM/VM DRS rules or VM/Host DRS rules prevent the virtual machine from being moved to this host.

- Virtual machines are pinned to their current hosts, so DRS cannot move these virtual machines to hosts in standby mode to balance the load.
- DRS or DPM is in manual mode and the recommendations were not applied.
- No virtual machines on highly used hosts are moved to that host.
- DPM is deactivated on the host because of a user setting or host previously failing to exit standby.

Address that issue that prevents DRS from powering on the host.

Total Cluster Load Is Low

The total cluster load is low.

When the total cluster load is low, DRS does not power off the host.

The following are possible reasons why DRS does not power off the host:

- Distributed Power Management (DPM) detected better candidates to power off.
- vSphere HA needs extra capacity for failover.
- The load is not low enough to trigger the host to power off.
- DPM projects that the load will increase.
- DPM is not enabled for the host.
- DPM threshold is set too high.
- While DPM is enabled for the host, no suitable power-on mechanism is present for the host.
- DRS cannot evacuate the host.
- The DRS migration threshold is at the highest setting and only performs mandatory moves.

Address the issue that is preventing DRS from powering off the host.

DRS Does Not Evacuate a Host Requested to Enter Maintenance or Standby Mode

DRS does not evacuate a host requested to enter maintenance mode or standby mode.

When you attempt to put a host into maintenance or standby mode, DRS does not evacuate the host as expected.

vSphere HA is activated and evacuating this host might violate HA failover capacity.

There is no solution. If appropriate, deactivate vSphere HA before you attempt to put the host into maintenance mode or standby mode.

DRS Does Not Move Any Virtual Machines onto a Host

DRS does not move any virtual machines onto a host.

DRS does not recommend migration of virtual machine to a host that has been added to a DRS-activated cluster.

After a host has been added to a DRS-activated cluster, the virtual machines deployed to the host become part of the cluster. DRS can recommend migration of some virtual machines to this host just added to the cluster. If that does not occur, there may be problems with vMotion, host compatibility, or affinity rules. The following are possible reasons:

- vMotion is not configured or activated on this host.
- Virtual machines on other hosts are not compatible with this host.
- The host does not have sufficient resources for any virtual machine.
- Moving any virtual machines to this host would violate a VM/VM DRS rule or VM/Host DRS rule.
- This host is reserved for HA failover capacity.
- A device is mounted to the virtual machine.
- The vMotion threshold is too high.

- DRS is deactivated for the virtual machines, hence the virtual machine could not be moved onto the destination host.

Address the issue that prevents DRS from moving virtual machines onto a host.

DRS Does Not Move Any Virtual Machines from a Host

DRS does not move any virtual machines from a host.

Virtual machines are not moved from this host.

This may be because of problems with vMotion, DRS, or host compatibility. The following are the possible reasons:

- vMotion is not configured or activated on this host.
- DRS is deactivated for the virtual machines on this host.
- Virtual machines on this host are not compatible with any other hosts.
- No other hosts have sufficient resources for any virtual machines on this host.
- Moving any virtual machines from this host would violate a VM/VM DRS rule or VM/Host DRS rule.
- DRS is deactivated for one or more virtual machines on the host.
- A device is mounted to the virtual machine.

Address the issues that are preventing DRS from moving virtual machines from the host.

VM migrations initiated by DRS due to incompatibility between the VMs and their host

In a DRS activated cluster, you may observe large amounts of VM migrations initiated by DRS targeting only some hosts in the cluster.

DRS scans hosts in the cluster every minute for possible recommendations as a part of its load balancing workflow. The outcome relies on the results of compatibility checks of powered-on VMs against the hosts in the cluster. If there are compatibility constraints that narrow a VM's compatible set of hosts to only certain hosts in the cluster, then DRS will try to satisfy these constraints by migrating that VM to one of the compatible hosts.

The constraints typically come from two sources, static user configuration and runtime state change in the cluster. This article focuses on the runtime state change that could be unexpected to the users.

The runtime state change that can impact the compatibility between powered-on VMs and hosts can be in one of the following areas.

vSphere High Availability

In a HA activated cluster, all hosts are expected to have a healthy HA status. If a host does not have healthy HA status at some time, it will generate a compatibility failure during VM compat-check. Some examples of such situations are: HA agent unreachable, HA agent isolated, HA agent partitioned. For more information, see, [Troubleshooting vSphere HA Host States](#).

These state changes typically accompany the following event in vCenter Server.

"vSphere HA agent on a host has an error"

Storage Accessibility

If a VM cannot access its configuration file (VMX file), virtual disks (VMDK) or swap file from its current host, it will fail the compat-check leading to the current host being incompatible. If a different host in the cluster still has accessibility to these files DRS could try to migrate the VM to that host. The outcome of such migration varies depending on the accessibility of VM's VMX file from its current host. If the VM only loses accessibility to its VMDK but not the VMX file, the migration could succeed. If the VM loses accessibility to its VMX file, the migration can fail.

Network Accessibility

For environments with NSX-T, the NSX component status could go down on some hosts, or all hosts at different times. In versions of vCenter Server prior to 7.0u2, this could lead to compatibility check failures for VMs and the affected hosts.

Resolution

1. If a user is planning to perform an operation that could potentially lead to the runtime state change discussed above, temporarily setting DRS to manual mode can avoid undesired migrations. Alternatively, a user can also temporarily set the DRS advanced option, VmsPerLBIteration, to 0. It asks DRS to scan no VM during its load balance workflow, hence no migration recommendation.
2. Since vCenter 7.0 Update 1, DRS introduced an advanced option to tolerate a powered-on VM's incompatibility with its current host for a predefined time period, CompatCheckTransientFailureTimeSeconds. Users can configure this option to avoid undesired migrations due to transient incompatibility.

Note: The default value of this option is 600 (10 minutes), which means DRS will not move a VM out just because of incompatibility with its current host if the incompatibility lasts for less than 10 mins. The maximum value of this option is 3600 (60 minutes).

Since vCenter 7.0 Update 3 and 8.0 Update 1, this option can also be set to -1, which disallows DRS to move VMs out due to its incompatibility with its current host.

Since vCenter 8.0 Update 3, the default value of this option is set to -1.

To set the DRS advanced option perform the following steps from the vSphere client:

1. Right-click the DRS cluster and click `Settings > vSphere DRS > Edit > Advanced Options > Add`
2. In the Option column enter the option name.
3. Click in the Value column to enter the desired value, then click OK for this setting to take effect.

Virtual Machine Problems

Virtual machine problems might cause DRS to not perform as expected.

Insufficient CPU or Memory Resources

The virtual machine does not receive enough CPU or memory resources.

In some cases, the virtual machine's demand is greater than its resource entitlement. When this occurs, the virtual machine doesn't receive enough CPU or memory resources.

The following sections describe the factors that influence the entitlement for a virtual machine.

Cluster is Yellow or Red

If the cluster is yellow or red, the capacity is insufficient to meet the resource reservations configured for all virtual machines and resource pools in the cluster. The particular virtual machine might be one that is not receiving its reservation. Check the status of the cluster (red or yellow) and resolve the situation.

Resource Limit is Too Restrictive

The virtual machine, its parent resource pool, or its resource pool ancestors might have a configured resource limit that is too restrictive. Check whether demand is equal to or greater than any configured limits.

Cluster is Overloaded

The cluster on which the virtual machine is running might have insufficient resources. Also, the virtual machine's share value is such that other virtual machines are granted proportionally more of the resources. To determine the demand is larger than the capacity, check the cluster statistics.

Host is Overloaded

To determine if the host's resources are oversubscribed, check the host statistics. If they are oversubscribed, consider why DRS is not moving any of the virtual machines now running on the host to other hosts. This condition might exist for the following reasons:

- The VM/VM DRS rules and VM/Host DRS rules require the current virtual machine-to-host mapping. If such rules are configured in the cluster, consider deactivating one or more of them. Then run DRS and check whether the situation is corrected.
- DRS cannot move this virtual machine or enough of the other virtual machines to other hosts to free up capacity. DRS will not move a virtual machine for any of the following reasons:
 - DRS is deactivated for the virtual machine.
 - A host device is mounted to the virtual machine.
 - Either of its resource reservations is so large that the virtual machine cannot run on any other host in the cluster.
 - The virtual machine is not compatible with any other host in the cluster.

Check whether any of these conditions exist for the virtual machine. If none exist, the conditions might exist for other virtual machines in the cluster. If this is the case, DRS cannot balance the cluster to address the virtual machine's demand.

- Decrease the DRS migration threshold setting and check whether the situation is resolved.
- Increase the virtual machine's reservation.

Address the problem that is causing the virtual machine to not receive enough CPU or memory resources.

VM/VM DRS Rule or VM/Host DRS Rule Violated

DRS rules specify which host a virtual machine must or must not reside on, or which virtual machines must be or must not be on the same host.

A VM/VM DRS rule or a VM/Host DRS rule is violated.

VM/VM DRS rules specify that selected virtual machines should be placed on the same host (affinity) or that virtual machines be placed on different hosts (anti-affinity). VM/Host DRS rules specify that selected virtual machines should be placed on specified hosts (affinity) or that selected virtual machines should not be placed on specified hosts (anti-affinity).

When a VM/VM DRS rule or VM/Host DRS rule is violated, it might be because DRS cannot move some or all of the virtual machines in the rule. The reservation of the virtual machine or other virtual machines in the affinity rule, or their parent resource pools, might prevent DRS from locating all virtual machines on the same host.

- Check the DRS faults panel for faults associated with affinity rules.
- Compute the sum of the reservations of all the virtual machines in the affinity rule. If that value is greater than the available capacity on any host, the rule cannot be satisfied.
- Compute the sum of the reservations of their parent resource pools. If that value is greater than the available capacity of any host, the rule cannot be satisfied if the resources are obtained from a single host.

Virtual Machine Power On Operation Fails

An error message appears stating that the virtual machine fails to power on.

The virtual machine fails to power on.

The virtual machine might fail to power on because of insufficient resources or because there are no compatible hosts for the virtual machine.

If the cluster does not have sufficient resources to power on a single virtual machine or any of the virtual machines in a group power-on attempt, check the resources required by the virtual machine against those available in the cluster or its parent resource pool. If necessary, reduce the reservations of the virtual machine to be powered-on, reduce the reservations of its sibling virtual machines, or increase the resources available in the cluster or its parent resource pool.

DRS Does Not Move the Virtual Machine

DRS does not move the virtual machine when it is initially powered on despite insufficient resources on the host.

When you power on a virtual machine, DRS does not migrate it as expected when there are not enough resources on the host where the virtual machine is registered.

The following are possible reasons why DRS does not move the virtual machine.

- DRS is deactivated on the virtual machine.
- The virtual machine has a device mounted.
- The virtual machine is not compatible with any other hosts.
- No other hosts have a sufficient number of physical CPUs or capacity for each CPU for the virtual machine.
- No other hosts have sufficient CPU or memory resources to satisfy the reservations and required memory of this virtual machine.
- Moving the virtual machine will violate an affinity or anti-affinity rule.
- The DRS automation level of the virtual machine is manual and the user does not approve the migration recommendation.
- DRS will not move fault tolerance-activated virtual machines.
- Check if a VM Override definition exists for the VM which may impede an automatic migration.

Address the issue that prevents DRS from moving the virtual machine.

vSphere Availability

vSphere Availability describes solutions that provide business continuity, including how to establish vSphere® High Availability (HA) and vSphere Fault Tolerance.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

This information is for anyone who wants to provide business continuity through the vSphere HA and Fault Tolerance solutions. The information in this book is for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

You Can Minimize Downtime with vSphere

Downtime, whether planned or unplanned, brings considerable costs. However, solutions that ensure higher levels of availability have traditionally been costly, hard to implement, and difficult to manage.

VMware software makes it simpler and less expensive to provide higher levels of availability for important applications. With vSphere, you can increase the baseline level of availability provided for all applications and provide higher levels of availability more easily and cost effectively. With vSphere, you can:

- Provide high availability independent of hardware, operating system, and applications.
- Reduce the planned downtime for common maintenance operations.
- Provide automatic recovery in cases of failure.

vSphere makes it possible to reduce planned downtime, prevent unplanned downtime, and recover rapidly from outages.

Reduce Planned Downtime with vSphere

Planned downtime typically accounts for over 80% of data center downtime. Hardware maintenance, server migration, and firmware updates all require downtime for physical servers. To minimize the impact of this downtime, organizations are forced to delay maintenance until inconvenient and difficult-to-schedule downtime windows.

vSphere makes it possible for organizations to dramatically reduce planned downtime. Because workloads in a vSphere environment can be dynamically moved to different physical servers without downtime or service interruption, server maintenance can be performed without requiring application and service downtime. With vSphere, organizations can:

- Eliminate downtime for common maintenance operations.
- Eliminate planned maintenance windows.
- Perform maintenance at any time without disrupting users and services.

The vSphere vMotion® and Storage vMotion functionality in vSphere makes it possible for organizations to reduce planned downtime because workloads in a VMware environment can be dynamically moved to different physical servers or to different underlying storage without service interruption. Administrators can perform faster and completely transparent maintenance operations, without being forced to schedule inconvenient maintenance windows.

Prevent Unplanned Downtime with vSphere

While an ESXi host provides a robust platform for running applications, an organization must also protect itself from unplanned downtime caused from hardware or application failures. vSphere builds important capabilities into data center infrastructure that can help you prevent unplanned downtime.

These vSphere capabilities are part of virtual infrastructure and are transparent to the operating system and applications running in virtual machines. These features can be configured and utilized by all the virtual machines on a physical system, reducing the cost and complexity of providing higher availability. Key availability capabilities are built into vSphere:

- Shared storage. Eliminate single points of failure by storing virtual machine files on shared storage, such as Fibre Channel or iSCSI SAN, or NAS. The use of SAN mirroring and replication features can be used to keep updated copies of virtual disk at disaster recovery sites.
- Network interface teaming. Provide tolerance of individual network card failures.
- Storage multipathing. Tolerate storage path failures.

In addition to these capabilities, the vSphere HA and Fault Tolerance features can minimize or eliminate unplanned downtime by providing rapid recovery from outages and continuous availability, respectively.

vSphere HA Provides Rapid Recovery from Outages

vSphere HA leverages multiple ESXi hosts configured as a cluster to provide rapid recovery from outages and cost-effective high availability for applications running in virtual machines.

vSphere HA protects application availability in the following ways:

- It protects against a server failure by restarting the virtual machines on other hosts within the cluster.
- It protects against application failure by continuously monitoring a virtual machine and resetting it in the event that a failure is detected.
- It protects against datastore accessibility failures by restarting affected virtual machines on other hosts which still have access to their datastores.
- It protects virtual machines against network isolation by restarting them if their host becomes isolated on the management or vSAN network. This protection is provided even if the network has become partitioned.

Unlike other clustering solutions, vSphere HA provides the infrastructure to protect all workloads with the infrastructure:

- You do not need to install special software within the application or virtual machine. All workloads are protected by vSphere HA. After vSphere HA is configured, no actions are required to protect new virtual machines. They are automatically protected.
- You can combine vSphere HA with vSphere Distributed Resource Scheduler (DRS) to protect against failures and to provide load balancing across the hosts within a cluster.

vSphere HA has several advantages over traditional failover solutions:

Minimal setup

After a vSphere HA cluster is set up, all virtual machines in the cluster get failover support without additional configuration.

Reduced hardware cost and setup

The virtual machine acts as a portable container for the applications and it can be moved among hosts. Administrators avoid duplicate configurations on multiple machines. When you use vSphere HA, you must have sufficient resources to fail over the number of hosts you want to protect with vSphere HA. However, the VMware vCenter Server® system automatically manages resources and configures clusters.

Increased application availability

Any application running inside a virtual machine has access to increased availability. Because the virtual machine can recover from hardware failure, all applications that start at boot have increased availability without increased computing needs, even if the application is not itself a clustered application. By monitoring and responding to VMware Tools heartbeats and restarting nonresponsive virtual machines, it protects against guest operating system crashes.

DRS and vMotion integration

If a host fails and virtual machines are restarted on other hosts, DRS can provide migration recommendations or migrate virtual machines for balanced resource allocation. If one or both of the

source and destination hosts of a migration fail, vSphere HA can help recover from that failure.

vSphere Fault Tolerance Provides Continuous Availability

vSphere HA provides a base level of protection for your virtual machines by restarting virtual machines in the event of a host failure. vSphere Fault Tolerance provides a higher level of availability, allowing users to protect any virtual machine from a host failure with no loss of data, transactions, or connections.

Fault Tolerance provides continuous availability by ensuring that the states of the Primary and Secondary VMs are identical at any point in the instruction execution of the virtual machine.

If either the host running the Primary VM or the host running the Secondary VM fails, an immediate and transparent failover occurs. The functioning ESXi host seamlessly becomes the Primary VM host without losing network connections or in-progress transactions. With transparent failover, there is no data loss and network connections are maintained. After a transparent failover occurs, a new Secondary VM is respawned and redundancy is re-established. The entire process is transparent and fully automated and occurs even if vCenter Server is unavailable.

Protecting vCenter Server with vCenter High Availability

vCenter High Availability (vCenter HA) protects not only against host and hardware failures but also against vCenter Server application failures. Using automated failover from active to passive, vCenter HA supports high availability with minimal downtime.

You configure vCenter HA from the vSphere Client. The configuration wizard provides these options.

Option	Description
Automatic	The automatic option clones the Active node to the Passive node and witness node, and configures the nodes for you. If your environment meets the following requirements, you can use this option. The vCenter Server that becomes the Active node is managing its own ESXi host and its own virtual machine. This configuration is sometimes called a self-managed vCenter Server.
Manual	The manual option offers more flexibility. You can use this option provided that your environment meets hardware and software requirements. If you select this option, you are responsible for cloning the Active node to the Passive node and the Witness node. You must also perform some networking configuration.

Protecting vCenter Server with VMware Service Lifecycle Manager

Availability of vCenter Server is provided by VMware Service Lifecycle Manager.

If a vCenter service fails, VMware Service Lifecycle Manager restarts it. VMware Service Lifecycle Manager monitors the health of services and it takes preconfigured remediation action when it detects a failure. Service does not restart if multiple attempts to remediate fail.

Creating and Using vSphere HA Clusters

vSphere HA clusters enable a collection of ESXi hosts to work together so that, as a group, they provide higher levels of availability for virtual machines than each ESXi host can provide individually. When you plan the creation and usage

of a new vSphere HA cluster, the options you select affect the way that cluster responds to failures of hosts or virtual machines.

Before you create a vSphere HA cluster, you should know how vSphere HA identifies host failures and isolation and how it responds to these situations. You also should know how admission control works so that you can choose the policy that fits your failover needs. After you establish a cluster, you can customize its behavior with advanced options and optimize its performance by following recommended best practices.

Note: You might get an error message when you try to use vSphere HA. For information about error messages related to vSphere HA, see the VMware knowledge base article at <http://kb.vmware.com/kb/1033634>.

How vSphere HA Works

vSphere HA provides high availability for virtual machines by pooling the virtual machines and the hosts they reside on into a cluster. Hosts in the cluster are monitored and in the event of a failure, the virtual machines on a failed host are restarted on alternate hosts.

When you create a vSphere HA cluster, a single host is automatically elected as the primary host. The primary host communicates with vCenter Server and monitors the state of all protected virtual machines and of the secondary hosts. Different types of host failures are possible, and the primary host must detect and appropriately deal with the failure. The primary host must distinguish between a failed host and one that is in a network partition or that has become network isolated. The primary host uses network and datastore heartbeating to determine the type of failure.

Primary and Secondary Hosts

When you add a host to a vSphere HA cluster, an agent is uploaded to the host and configured to communicate with other agents in the cluster. Each host in the cluster functions as a primary host or a secondary host.

When vSphere HA is enabled for a cluster, all active hosts (that are not in standby, maintenance mode or not disconnected) participate in an election to choose the cluster's primary host. The host that mounts the greatest number of datastores has an advantage in the election. Only one primary host typically exists per cluster and all other hosts are secondary hosts. If the primary host fails, is shut down or put in standby mode, or is removed from the cluster a new election is held.

The primary host in a cluster has several responsibilities:

- Monitoring the state of secondary hosts. If a secondary host fails or becomes unreachable, the primary host identifies which virtual machines must be restarted.
- Monitoring the power state of all protected virtual machines. If one virtual machine fails, the primary host ensures that it is restarted. Using a local placement engine, the primary host also determines where the restart takes place.
- Managing the lists of cluster hosts and protected virtual machines.
- Acting as the vCenter Server management interface to the cluster and reporting the cluster health state.

The secondary hosts primarily contribute to the cluster by running virtual machines locally, monitoring their runtime states, and reporting state updates to the primary host. A primary host can also run and monitor virtual machines. Both secondary hosts and primary hosts implement the VM and Application Monitoring features.

One of the functions performed by the primary host is to orchestrate restarts of protected virtual machines. A virtual machine is protected by a primary host after vCenter Server observes that the virtual machine's power state has changed from powered off to powered on in response to a user action. The primary host persists the list of protected virtual machines in the cluster's datastores. A newly elected primary host uses this information to determine which virtual machines to protect.

Note: If you disconnect a host from a cluster, the virtual machines registered to that host are unprotected by vSphere HA.

Host Failure Types

The primary host of a VMware vSphere® High Availability cluster is responsible for detecting the failure of secondary hosts. Depending on the type of failure detected, the virtual machines running on the hosts might need to be failed over.

In a vSphere HA cluster, three types of host failure are detected:

- Failure. A host stops functioning.
- Isolation. A host becomes network isolated.
- Partition. A host loses network connectivity with the primary host.

The primary host monitors the liveness of the secondary hosts in the cluster. This communication happens through the exchange of network heartbeats every second. When the primary host stops receiving these heartbeats from a secondary host, it checks for host liveness before declaring the host failed. The liveness check that the primary host performs is to determine whether the secondary host is exchanging heartbeats with one of the datastores. See [Datastore Heartbeating](#) . Also, the primary host checks whether the host responds to ICMP pings sent to its management IP addresses.

If a primary host cannot communicate directly with the agent on a secondary host, the secondary host does not respond to ICMP pings. If the agent is not issuing heartbeats, it is viewed as failed. The host's virtual machines are restarted on alternate hosts. If such a secondary host is exchanging heartbeats with a datastore, the primary host assumes that the secondary host is in a network partition or is network isolated. So, the primary host continues to monitor the host and its virtual machines. See [Network Partitions](#) .

Host network isolation occurs when a host is still running, but it can no longer observe traffic from vSphere HA agents on the management network. If a host stops observing this traffic, it attempts to ping the cluster isolation addresses. If this pinging also fails, the host declares that it is isolated from the network.

The primary host monitors the virtual machines that are running on an isolated host. If the primary host observes that the VMs power off, and the primary host is responsible for the VMs, it restarts them.

Note: If you ensure that the network infrastructure is sufficiently redundant and that at least one network path is always available, host network isolation is less likely to occur.

Proactive HA Failures

A Proactive HA failure occurs when a host component fails, which results in a loss of redundancy or a noncatastrophic failure. However, the functional behavior of the VMs residing on the host is not yet affected. For example, if a power supply on the host fails, but other power supplies are available, that is a Proactive HA failure.

If a Proactive HA failure occurs, you can automate the remediation action taken in the vSphere Availability section of the vSphere Client. The VMs on the affected host can be evacuated to other hosts and the host is either placed in Quarantine mode or Maintenance mode.

Note: Your cluster must use vSphere DRS for the Proactive HA failure monitoring to work.

Determining Responses to Host Issues

If a host fails and its virtual machines must be restarted, you can control the order in which the virtual machines are restarted with the VM restart priority setting. You can also configure how vSphere HA responds if hosts lose management network connectivity with other hosts by using the host isolation response setting. Other factors are also considered when vSphere HA restarts a virtual machine after a failure.

The following settings apply to all virtual machines in the cluster in the case of a host failure or isolation. You can also configure exceptions for specific virtual machines. See [Customize an Individual Virtual Machine](#) .

Host Isolation Response

Host isolation response determines what happens when a host in a vSphere HA cluster loses its management network connections, but continues to run. You can use the isolation response to have vSphere HA power off virtual machines that are running on an isolated host and restart them on a non-isolated host. Host isolation responses require that Host Monitoring Status is activated. If Host Monitoring Status is deactivated, host isolation responses are also suspended. A host determines that it is isolated when it is unable to communicate with the agents running on the other hosts, and it is unable to ping its isolation addresses. The host then executes its isolation response. The responses are Power off and restart VMs or Shutdown and restart VMs. You can customize this property for individual virtual machines.

Note: If a virtual machine has a restart priority setting of Disabled, no host isolation response is made.

To use the Shutdown and restart VMs setting, you must install VMware Tools in the guest operating system of the virtual machine. Shutting down the virtual machine provides the advantage of preserving its state. Shutting down is better than powering off the virtual machine, which does not flush most recent changes to disk or commit transactions. Virtual machines that are in the process of shutting down take longer to fail over while the shutdown completes. Virtual Machines that have not shut down in 300 seconds, or the time specified in the advanced option `das.isolationshutdowntimeout`, are powered off.

After you create a vSphere HA cluster, you can override the default cluster settings for Restart Priority and Isolation Response for specific virtual machines. Such overrides are useful for virtual machines that are used for special tasks. For example, virtual machines that provide infrastructure services like DNS or DHCP might need to be powered on before other virtual machines in the cluster.

A virtual machine "split-brain" condition can occur when a host becomes isolated or partitioned from a primary host and the primary host cannot communicate with it using heartbeat datastores. In this situation, the primary host cannot determine that the host is alive and so declares it dead. The primary host then attempts to restart the virtual machines that are running on the isolated or partitioned host. This attempt succeeds if the virtual machines remain running on the isolated/partitioned host and that host lost access to the virtual machines' datastores when it became isolated or partitioned. A split-brain condition then exists because there are two instances of the virtual machine. However, only one instance is able to read or write the virtual machine's virtual disks. VM Component Protection can be used to prevent this split-brain condition. When you activate VMCP with the aggressive setting, it monitors the datastore accessibility of powered-on virtual machines, and shuts down those that lose access to their datastores.

To recover from this situation, ESXi generates a question on the virtual machine that has lost the disk locks for when the host comes out of isolation and cannot reacquire the disk locks. vSphere HA automatically answers this question, allowing the virtual machine instance that has lost the disk locks to power off, leaving just the instance that has the disk locks.

Virtual Machine Dependencies

You can create dependencies between groups of virtual machines. To do so, you must first create the VM groups in the vSphere Client by going to the **Configure** tab for the cluster and selecting **VM/Host Groups**. Once the groups have been created, you can create restart dependency rules between the groups by browsing to **VM/Host Rules** and in the Type drop-down menu, select **Virtual Machines to Virtual Machines**. These rules can specify that certain VM groups cannot be restarted until other, specified VM groups have been Ready first.

Factors Considered for Virtual Machine Restarts

After a failure, the cluster's primary host attempts to restart affected virtual machines by identifying a host that can power them on. When choosing such a host, the primary host considers a number of factors.

File accessibility

Before a virtual machine can be started, its files must be accessible from one of the active cluster hosts that the primary can communicate with over the network

Virtual machine and host compatibility

If there are accessible hosts, the virtual machine must be compatible with at least one of them. The compatibility set for a virtual machine includes the effect of any required VM-Host affinity

Resource reservations

rules. For example, if a rule only permits a virtual machine to run on two hosts, it is considered for placement on those two hosts. Of the hosts that the virtual machine can run on, at least one must have sufficient unreserved capacity to meet the memory overhead of the virtual machine and any resource reservations. Four types of reservations are considered: CPU, Memory, vNIC, and Virtual flash. Also, sufficient network ports must be available to power on the virtual machine.

Host limits

In addition to resource reservations, a virtual machine can only be placed on a host if doing so does not violate the maximum number of allowed virtual machines or the number of in-use vCPUs.

Feature constraints

If the advanced option has been set that requires vSphere HA to enforce VM to VM anti-affinity rules, vSphere HA does not violate this rule. Also, vSphere HA does not violate any configured per host limits for fault tolerant virtual machines.

If no hosts satisfy the preceding considerations, the primary host issues an event stating that there are not enough resources for vSphere HA to start the VM and tries again when the cluster conditions have changed. For example, if the virtual machine is not accessible, the primary host tries again after a change in file accessibility.

VM and Application Monitoring

VM Monitoring restarts individual virtual machines if their VMware Tools heartbeats are not received within a set time. Similarly, Application Monitoring can restart a virtual machine if the heartbeats for an application it is running are not received. You can enable these features and configure the sensitivity with which vSphere HA monitors non-responsiveness.

When you enable VM Monitoring, the VM Monitoring service (using VMware Tools) evaluates whether each virtual machine in the cluster is running by checking for regular heartbeats and I/O activity from the VMware Tools process running inside the guest. If no heartbeats or I/O activity are received, this is most likely because the guest operating system has failed or VMware Tools is not being allocated any time to complete tasks. In such a case, the VM Monitoring service determines that the virtual machine has failed and the virtual machine is rebooted to restore service. For information on how to enable VM monitoring, see [Activate VM Monitoring](#).

Occasionally, virtual machines or applications that are still functioning properly stop sending heartbeats. To avoid unnecessary resets, the VM Monitoring service also monitors a virtual machine's I/O activity. If no heartbeats are received within the failure interval, the I/O stats interval (a cluster-level attribute) is checked. The I/O stats interval determines if any disk or network activity has occurred for the virtual machine during the previous two minutes (120 seconds). If not, the virtual machine is reset. This default value (120 seconds) can be changed using the advanced option `das.iostatsinterval`.

To enable Application Monitoring, you must first obtain the appropriate SDK (or be using an application that supports VMware Application Monitoring) and use it to set up customized heartbeats for the applications you want to monitor. After you have done this, Application Monitoring works much the same way that VM Monitoring does. If the heartbeats for an application are not received for a specified time, its virtual machine is restarted.

You can configure the level of monitoring sensitivity. Highly sensitive monitoring results in a more rapid conclusion that a failure has occurred. While unlikely, highly sensitive monitoring might lead to falsely identifying failures when the virtual machine or application in question is actually still working, but heartbeats have not been received due to factors such as resource constraints. Low sensitivity monitoring results in longer interruptions in service between actual failures and virtual machines being reset. Select an option that is an effective compromise for your needs.

You can also specify custom values for the monitoring sensitivity interval by selecting the **Custom** check box.

Table 372: VM Monitoring Settings

Preset Setting	Failure interval (seconds)	Minimum uptime (seconds)	Maximum per-VM resets	Maximum resets time window (hours)
High	30	120	3	1
Medium	60	240	3	24
Low	120	480	3	168

- **Failure interval:** If no VM heartbeat is received in the specified interval, vSphere HA restarts the VM.
- **Minimum uptime:** After a VM starts, vSphere HA waits until the specified interval is reached, before monitoring for VM heartbeats, storage and network traffic.
- **Maximum per-VM resets:** The maximum number of resets, that vSphere HA can restart a VM a within the specified interval for **Maximum resets time window**.
- **Maximum reset time window:** The maximum interval for the number of specified **Maximum per-VM resets**.

Note: To set a custom value for the I/O stats interval, you can use the `das.iostatsInterval` advanced configuration option. For more information, see [KB 2033250](#).

After failures are detected, vSphere HA resets virtual machines. The reset ensures that services remain available. To avoid resetting virtual machines repeatedly for nontransient errors, by default, virtual machines will be reset only three times during a certain configurable time interval. After virtual machines have been reset three times, vSphere HA makes no further attempts to reset the virtual machines after subsequent failures until after the specified time has elapsed. You can configure the number of resets using the **Maximum per-VM resets** custom setting.

Note: The reset statistics are cleared when a virtual machine is powered off then back on, or when it is migrated using vMotion to another host. This causes the guest operating system to reboot, but is not the same as a 'restart' in which the power state of the virtual machine is changed.

VM Component Protection

If VM Component Protection (VMCP) is activated, vSphere HA can detect datastore accessibility failures and provide automated recovery for affected virtual machines.

VMCP provides protection against datastore accessibility failures that can affect a virtual machine running on a host in a vSphere HA cluster. When a datastore accessibility failure occurs, the affected host can no longer access the storage path for a specific datastore. You can determine the response that vSphere HA will make to such a failure, ranging from the creation of event alarms to virtual machine restarts on other hosts.

Note: When you use the VM Component Protection feature, your ESXi hosts must be version 6.0 or higher.

Types of Failure

There are two types of datastore accessibility failure:

PDL	PDL (Permanent Device Loss) is an unrecoverable loss of accessibility that occurs when a storage device reports the datastore is no longer accessible by the host. This condition cannot be reverted without powering off virtual machines.
APD	APD (All Paths Down) represents a transient or unknown accessibility loss or any other unidentified delay in I/O processing. This type of accessibility issue is recoverable.

Configuring VMCP

VM Component Protection is configured in the vSphere Client. Go to the **Configure** tab and click **vSphere Availability** and **Edit**. Under **Failures and Responses** you can select **Datastore with PDL** or **Datastore with APD**. The storage protection levels you can choose and the virtual machine remediation actions available differ depending on the type of database accessibility failure.

PDL Failures

Under **Datastore with PDL**, you can select **Issue events** or **Power off and restart VMs**.

APD Failures

The response to APD events is more complex and accordingly the configuration is more fine-grained. You can select **Issue events**, **Power off and restart VMs--conservative restart policy**, or **Power off and restart VMs--aggressive restart policy**.

Note: If either the Host Monitoring or VM Restart Priority settings are deactivated, VMCP cannot perform virtual machine restarts. Storage health can still be monitored and events can be issued, however.

Network Partitions

When a management network failure occurs for a vSphere HA cluster, a subset of the cluster's hosts might be unable to communicate over the management network with the other hosts. Multiple partitions can occur in a cluster.

A partitioned cluster leads to degraded virtual machine protection and cluster management functionality. Correct the partitioned cluster as soon as possible.

- Virtual machine protection. vCenter Server allows a virtual machine to be powered on, but it can be protected only if it is running in the same partition as the primary host that is responsible for it. The primary host must be communicating with vCenter Server. A primary host is responsible for a virtual machine if it has exclusively locked a system-defined file on the datastore that contains the virtual machine's configuration file.
- Cluster management. vCenter Server can communicate with the primary host, but only a subset of the secondary hosts. As a result, changes in configuration that affect vSphere HA might not take effect until after the partition is resolved. This failure could result in one of the partitions operating under the old configuration, while another uses the new settings.

Datastore Heartbeating

When the primary host in a VMware vSphere® High Availability cluster cannot communicate with a secondary host over the management network, the primary host uses datastore heartbeating to determine whether the secondary host has failed, is in a network partition, or is network isolated. If the secondary host has stopped datastore heartbeating, it is considered to have failed and its virtual machines are restarted elsewhere.

VMware vCenter Server® selects a preferred set of datastores for heartbeating. This selection is made to maximize the number of hosts that have access to a heartbeating datastore and minimize the likelihood that the datastores are backed by the same LUN or NFS server.

You can use the advanced option `das.heartbeatdsperhost` to change the number of heartbeat datastores selected by vCenter Server for each host. The default is two and the maximum valid value is five.

vSphere HA creates a directory at the root of each datastore that is used for both datastore heartbeating and for persisting the set of protected virtual machines. The name of the directory is `.vSphere-HA`. Do not delete or modify the files stored in this directory, because this can have an impact on operations. Because more than one cluster might use a datastore, subdirectories for this directory are created for each cluster. Root owns these directories and files and only root can read and write to them. The disk space used by vSphere HA depends on several factors including which VMFS version is in use and the number of hosts that use the datastore for heartbeating. With `vmfs3`, the maximum usage is 2GB and the typical usage is 3MB. With `vmfs5` and `vmfs6`, the maximum and typical usage is 3MB. vSphere HA use of the datastores adds negligible overhead and has no performance impact on other datastore operations.

vSphere HA limits the number of virtual machines that can have configuration files on a single datastore. See *Configuration Maximums* for updated limits. If you place more than this number of virtual machines on a datastore and power them on, vSphere HA protects virtual machines only up to the limit.

Note: A vSAN datastore cannot be used for datastore heartbeating. Therefore, if no other shared storage is accessible to all hosts in the cluster, there can be no heartbeat datastores in use. However, if you have storage that is accessible by an alternate network path independent of the vSAN network, you can use it to set up a heartbeat datastore.

vSphere HA Security

vSphere HA is enhanced by several security features.

Select firewall ports opened	vSphere HA uses TCP and UDP port 8182 for agent-to-agent communication. The firewall ports open and close automatically to ensure they are open only when needed.
Configuration files protected using file system permissions	vSphere HA stores configuration information on the local storage or on ramdisk if there is no local datastore. These files are protected using file system permissions and they are accessible only to the root user. Hosts without local storage are only supported if they are managed by Auto Deploy.
Detailed logging	The location where vSphere HA places log files depends on the version of host. - For ESXi hosts, vSphere HA writes to syslog only by default, so logs are placed where syslog is configured to put them. The log file names for vSphere HA are prepended with <code>fdm</code>, fault domain manager, which is a service of vSphere HA. - For legacy ESXi hosts, vSphere HA writes to <code>/var/log/vmware/fdm</code> on local disk, as well as syslog if it is configured.
Secure vSphere HA logins	vSphere HA logs onto the vSphere HA agents using a user account, vpuser , created by vCenter Server. This account is the same account used by vCenter Server to manage the host. vCenter Server creates a random password for this account and changes the password periodically. The time period is set by the vCenter Server <code>VirtualCenter.VimPasswordExpirationInDays</code> setting. Users with administrative privileges on the root folder of the host can log in to the agent.
Secure communication	All communication between vCenter Server and the vSphere HA agent is done over SSL. Agent-to-agent communication also uses SSL except for election messages, which occur over UDP. Election messages are verified over SSL so that a rogue agent can prevent only the host on which the agent is running from being elected as a primary host. In this case, a configuration issue for the cluster is issued so the user is aware of the problem.
Host SSL certificate verification required	vSphere HA requires that each host have a verified SSL certificate. Each host generates a self-signed certificate when it is booted for the first time. This certificate can then be regenerated or replaced with one issued by an authority. If the certificate is replaced, vSphere HA needs to be reconfigured on the host. If a host becomes disconnected from vCenter Server after its certificate is updated and the ESXi or ESX Host agent is restarted, then vSphere HA is automatically reconfigured when the host is reconnected to vCenter Server. If the disconnection does not occur because vCenter Server host SSL certificate verification is deactivated at the time, verify the new certificate and reconfigure vSphere HA on the host.

vSphere HA Admission Control

vSphere HA uses admission control to ensure that sufficient resources are reserved for virtual machine recovery when a host fails.

Admission control imposes constraints on resource usage. Any action that might violate these constraints is not permitted. Actions that might be disallowed include the following examples:

- Powering on a virtual machine
- Migrating a virtual machine
- Increasing the CPU or memory reservation of a virtual machine

The basis for vSphere HA admission control is how many host failures your cluster is allowed to tolerate and still guarantee failover. The host failover capacity can be set in three ways:

- Cluster resource percentage
- Slot policy
- Dedicated failover hosts

Note: vSphere HA admission control can be deactivated. However, without it you have no assurance that the expected number of virtual machines can be restarted after a failure. Do not permanently deactivate admission control.

Note: Within a cluster, you must deactivate HA admission control temporarily to let vSphere vMotion proceed. This action prevents downtime for the machines on the hosts that you remediate. Disabling HA admission control before you remediate a two-node cluster causes the cluster to practically lose all its high availability guarantees. The reason is that when one of the two hosts enters maintenance mode, vCenter Server cannot failover virtual machines to that host and HA failovers are never successful.

Note: To use vSphere HA admission control you must have at least 3 hosts in the cluster.

Regardless of the admission control option chosen, a VM resource reduction threshold also exists. You use this setting to specify the percentage of resource degradation to tolerate, but it is not available unless vSphere DRS is activated.

The resource reduction calculation is checked for both CPU and memory. It considers a virtual machine's reserved memory and memory overload to decide whether to permit it to power on, migrate, or have reservation changes. The actual memory consumed by the virtual machine is not considered in the calculation because the memory reservation does not always correlate with the actual memory usage of the virtual machine. If the actual usage is more than reserved memory, insufficient failover capacity is available, resulting in performance degradation on failover.

Setting a performance reduction threshold allows you to specify the occurrence of a configuration issue. For example:

- The default value is 100%, which produces no warnings.
- If you reduce the threshold to 0%, a warning is generated when cluster usage exceeds the available capacity.
- If you reduce the threshold to 20%, the performance reduction that can be tolerated is calculated as $\text{performance reduction} = \text{current utilization} * 20\%$. When the current usage minus the performance reduction exceeds the available capacity, a configuration notice is issued.

Cluster Resources Percentage Admission Control

You can configure vSphere HA to perform admission control by reserving a specific percentage of cluster CPU and memory resources for recovery from host failures.

With this type of admission control, vSphere HA ensures that a specified percentage of aggregate CPU and memory resources are reserved for failover.

With the cluster resources percentage option, vSphere HA enforces admission control as follows:

1. Calculates the total resource requirements for all powered-on virtual machines in the cluster.
2. Calculates the total host resources available for virtual machines.

3. Calculates the Current CPU Failover Capacity and Current Memory Failover Capacity for the cluster.
4. Determines if either the Current CPU Failover Capacity or Current Memory Failover Capacity is less than the corresponding Configured Failover Capacity (provided by the user).
If so, admission control disallows the operation.

vSphere HA uses the actual reservations of the virtual machines. If a virtual machine does not have reservations, meaning that the reservation is 0, a default of 0MB memory and 32MHz CPU is applied.

Note: The cluster resources percentage option for admission control also checks that there are at least two vSphere HA-enabled hosts in the cluster (excluding hosts that are entering maintenance mode). If there is only one vSphere HA-enabled host, an operation is not allowed, even if there is a sufficient percentage of resources available. The reason for this extra check is that vSphere HA cannot perform failover if there is only a single host in the cluster.

Computing the Current Failover Capacity

The total resource requirements for the powered-on virtual machines is comprised of two components, CPU and memory. vSphere HA calculates these values.

- The CPU component by summing the CPU reservations of the powered-on virtual machines. If you have not specified a CPU reservation for a virtual machine, it is assigned a default value of 32MHz (this value can be changed using the `das.vmcputminmhz` advanced option.)
- The memory component by summing the memory reservation (plus memory overhead) of each powered-on virtual machine.


The total host resources available for virtual machines is calculated by adding the hosts' CPU and memory resources. These amounts are those contained in the host's root resource pool, not the total physical resources of the host. Resources being used for virtualization purposes are not included. Only hosts that are connected, not in maintenance mode, and have no vSphere HA errors are considered.

The Current CPU Failover Capacity is computed by subtracting the total CPU resource requirements from the total host CPU resources and dividing the result by the total host CPU resources. The Current Memory Failover Capacity is calculated similarly.

Admission Control Using Cluster Resources Percentage

The way that Current Failover Capacity is calculated and used with this admission control policy is shown with an example. Make the following assumptions about a cluster:

- The cluster is comprised of three hosts, each with a different amount of available CPU and memory resources. The first host (H1) has 9GHz of available CPU resources and 9GB of available memory, while Host 2 (H2) has 9GHz and 6GB and Host 3 (H3) has 6GHz and 6GB.
- There are five powered-on virtual machines in the cluster with differing CPU and memory requirements. VM1 needs 2GHz of CPU resources and 1GB of memory, while VM2 needs 2GHz and 1GB, VM3 needs 1GHz and 2GB, VM4 needs 1GHz and 1GB, and VM5 needs 1GHz and 1GB.
- The Configured Failover Capacity for CPU and Memory are both set to 25%.

Figure 138: Admission Control Example with Percentage of Cluster Resources Reserved Policy

The total resource requirements for the powered-on virtual machines is 7GHz and 6GB. The total host resources available for virtual machines is 24GHz and 21GB. Based on this, the Current CPU Failover Capacity is 70% $((24\text{GHz} - 7\text{GHz})/24\text{GHz})$. Similarly, the Current Memory Failover Capacity is 71% $((21\text{GB} - 6\text{GB})/21\text{GB})$.

Because the cluster's Configured Failover Capacity is set to 25%, 45% of the cluster's total CPU resources and 46% of the cluster's memory resources are still available to power on additional virtual machines.

Slot Policy Admission Control

With the slot policy option, vSphere HA admission control ensures that a specified number of hosts can fail and sufficient resources remain in the cluster to fail over all the virtual machines from those hosts.

Using the slot policy, vSphere HA performs admission control in the following way:

1. Calculates the slot size.
A slot is a logical representation of memory and CPU resources. By default, it is sized to satisfy the requirements for any powered-on virtual machine in the cluster.
2. Determines how many slots each host in the cluster can hold.
3. Determines the Current Failover Capacity of the cluster.
This is the number of hosts that can fail and still leave enough slots to satisfy all of the powered-on virtual machines.
4. Determines whether the Current Failover Capacity is less than the Configured Failover Capacity (provided by the user).
If it is, admission control disallows the operation.

Note: You can set a specific slot size for both CPU and memory in the admission control section of the vSphere HA settings in the vSphere Client.

Slot Size Calculation

Slot size is comprised of two components, CPU and memory.

- vSphere HA calculates the CPU component by obtaining the CPU reservation of each powered-on virtual machine and selecting the largest value. If you have not specified a CPU reservation for a virtual machine, it is assigned a default value of 32MHz. You can change this value by using the `das.vmcputminmhz` advanced option.)
- vSphere HA calculates the memory component by obtaining the memory reservation, plus memory overhead, of each powered-on virtual machine and selecting the largest value. There is no default value for the memory reservation.

If your cluster contains any virtual machines that have much larger reservations than the others, they will distort slot size calculation. To avoid this, you can specify an upper bound for the CPU or memory component of the slot size by using the `das.slotcpuinmhz` or `das.slotmeminmb` advanced options, respectively. See [vSphere HA Advanced Options](#).

You can also determine the risk of resource fragmentation in your cluster by viewing the number of virtual machines that require multiple slots. This can be calculated in the admission control section of the vSphere HA settings in the vSphere Client. Virtual machines might require multiple slots if you have specified a fixed slot size or a maximum slot size using advanced options.

Using Slots to Compute the Current Failover Capacity

After the slot size is calculated, vSphere HA determines each host's CPU and memory resources that are available for virtual machines. These amounts are those contained in the host's root resource pool, not the total physical resources of the host. The resource data for a host that is used by vSphere HA can be found on the host's **Summary** tab on the vSphere Client. If all hosts in your cluster are the same, this data can be obtained by dividing the cluster-level figures by the number of hosts. Resources being used for virtualization purposes are not included. Only hosts that are connected, not in maintenance mode, and that have no vSphere HA errors are considered.


The maximum number of slots that each host can support is then determined. To do this, the host's CPU resource amount is divided by the CPU component of the slot size and the result is rounded down. The same calculation is made for the host's memory resource amount. These two numbers are compared and the smaller number is the number of slots that the host can support.

The Current Failover Capacity is computed by determining how many hosts (starting from the largest) can fail and still leave enough slots to satisfy the requirements of all powered-on virtual machines.

Admission Control Using Slot Policy

The way that slot size is calculated and used with this admission control policy is shown in an example. Make the following assumptions about a cluster:

- The cluster is comprised of three hosts, each with a different amount of available CPU and memory resources. The first host (H1) has 9GHz of available CPU resources and 9GB of available memory, while Host 2 (H2) has 9GHz and 6GB and Host 3 (H3) has 6GHz and 6GB.
- There are five powered-on virtual machines in the cluster with differing CPU and memory requirements. VM1 needs 2GHz of CPU resources and 1GB of memory, while VM2 needs 2GHz and 1GB, VM3 needs 1GHz and 2GB, VM4 needs 1GHz and 1GB, and VM5 needs 1GHz and 1GB.
- The Host Failures Cluster Tolerates is set to one.

Figure 139: Admission Control Example with Host Failures Cluster Tolerates Policy

1. Slot size is calculated by comparing both the CPU and memory requirements of the virtual machines and selecting the largest.
The largest CPU requirement (shared by VM1 and VM2) is 2GHz, while the largest memory requirement (for VM3) is 2GB. Based on this, the slot size is 2GHz CPU and 2GB memory.
2. Maximum number of slots that each host can support is determined.
H1 can support four slots. H2 can support three slots (which is the smaller of 9GHz/2GHz and 6GB/2GB) and H3 can also support three slots.
3. Current Failover Capacity is computed.
The largest host is H1 and if it fails, six slots remain in the cluster, which is sufficient for all five of the powered-on virtual machines. If both H1 and H2 fail, only three slots remain, which is insufficient. Therefore, the Current Failover Capacity is one.

The cluster has one available slot (the six slots on H2 and H3 minus the five used slots).

Dedicated Failover Hosts Admission Control

You can configure vSphere HA to designate specific hosts as the failover hosts.

With dedicated failover hosts admission control, when a host fails, vSphere HA attempts to restart its virtual machines on any of the specified failover hosts. If restarting the virtual machines is not possible, for example the failover hosts have failed or have insufficient resources, then vSphere HA attempts to restart those virtual machines on other hosts in the cluster.

To ensure that spare capacity is available on a failover host, you are prevented from powering on virtual machines or using vMotion to migrate virtual machines to a failover host. Also, DRS does not use a failover host for load balancing.

Note: If you use dedicated failover hosts admission control and designate multiple failover hosts, DRS does not attempt to enforce VM-VM affinity rules for virtual machines that are running on failover hosts.

vSphere HA Interoperability

vSphere HA can interoperate with many other features, such as DRS and vSAN.

Before configuring vSphere HA, you should be aware of the limitations of its interoperability with these other features or products.

Using vSphere HA with vSAN

You can use vSAN as the shared storage for a vSphere HA cluster. If activated, vSAN aggregates the specified local storage disks available on the hosts into a single datastore shared by all hosts.

To use vSphere HA with vSAN, you must be aware of certain considerations and limitations for the interoperability of these two features.

For information about vSAN, see *Administering VMware vSAN*.

Note: You can use vSphere HA with vSAN stretched clusters.

ESXi Host Requirements

You can use vSAN with a vSphere HA cluster only if the following conditions are met:

- All the cluster's ESXi hosts must be version 5.5 or later.
- The cluster must have a minimum of three ESXi hosts.

Networking Differences

vSAN has its own network. If vSAN and vSphere HA are activated for the same cluster, the HA interagent traffic flows over this storage network rather than the management network. vSphere HA uses the management network only if vSAN is deactivated. vCenter Server chooses the appropriate network if vSphere HA is configured on a host.

Note: You can activate vSAN only if vSphere HA is deactivated.

If you change the vSAN network configuration, the vSphere HA agents do not automatically pick up the new network settings. To make changes to the vSAN network, you must take the following steps in the vSphere Client:

1. Deactivate Host Monitoring for the vSphere HA cluster.
2. Make the vSAN network changes.
3. Right-click all hosts in the cluster and select **Reconfigure for vSphere HA**.
4. Re-activate Host Monitoring for the vSphere HA cluster.

The following table shows the differences in vSphere HA networking whether or not vSAN is used.

Table 373: vSphere HA Networking Differences

	vSAN Activated	vSAN Deactivated
Network used by vSphere HA	vSAN storage network	Management network
Heartbeat datastores	Any datastore mounted to > 1 host, but not vSAN datastores	Any datastore mounted to > 1 host
Host declared isolated	Isolation addresses not pingable and vSAN storage network inaccessible	Isolation addresses not pingable and management network inaccessible

Capacity Reservation Settings

When you reserve capacity for your vSphere HA cluster with an admission control policy, you must coordinate this setting with the corresponding vSAN setting that ensures data accessibility on failures. Specifically, the Number of Failures Tolerated setting in the vSAN rule set must not be lower than the capacity that the vSphere HA admission control setting reserved.

For example, if the vSAN rule set allows for only two failures, the vSphere HA admission control policy must reserve capacity that is equivalent to only one or two host failures. If you are using the Percentage of Cluster Resources Reserved policy for a cluster that has eight hosts, you must not reserve more than 25% of the cluster resources. In the same cluster, with the Host Failures Cluster Tolerates policy, the setting must not be higher than two hosts. If vSphere HA reserves less capacity, failover activity might be unpredictable. Reserving too much capacity overly constrains the powering on of virtual machines and intercluster vSphere vMotion migrations.

Using vSphere HA and DRS Together

Using vSphere HA with Distributed Resource Scheduler (DRS) combines automatic failover with load balancing. This combination can result in a more balanced cluster after vSphere HA has moved virtual machines to different hosts.

When vSphere HA performs failover and restarts virtual machines on different hosts, its first priority is the immediate availability of all virtual machines. After the virtual machines have been restarted, those hosts on which they were powered on might be heavily loaded, while other hosts are comparatively lightly loaded. vSphere HA uses the virtual machine's CPU and memory reservation and overhead memory to determine if a host has enough spare capacity to accommodate the virtual machine.

In a cluster using DRS and vSphere HA with admission control turned on, virtual machines might not be evacuated from hosts entering maintenance mode. This behavior occurs because of the resources reserved for restarting virtual machines in the event of a failure. You must manually migrate the virtual machines off of the hosts using vMotion.

In some scenarios, vSphere HA might not be able to fail over virtual machines because of resource constraints. This can occur for several reasons.

- HA admission control is deactivated and Distributed Power Management (DPM) is activated. This can result in DPM consolidating virtual machines onto fewer hosts and placing the empty hosts in standby mode leaving insufficient powered-on capacity to perform a failover.
- VM-Host affinity (required) rules might limit the hosts on which certain virtual machines can be placed.
- There might be sufficient aggregate resources but these can be fragmented across multiple hosts so that they can not be used by virtual machines for failover.

In such cases, vSphere HA can use DRS to try to adjust the cluster (for example, by bringing hosts out of standby mode or migrating virtual machines to defragment the cluster resources) so that HA can perform the failovers.

If DPM is in manual mode, you might need to confirm host power-on recommendations. Similarly, if DRS is in manual mode, you might need to confirm migration recommendations.

If you are using VM-Host affinity rules that are required, be aware that these rules cannot be violated. vSphere HA does not perform a failover if doing so would violate such a rule.

For more information about DRS, see the *vSphere Resource Management* documentation.

Note: vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere Cluster. Starting with vSphere 7.0 Update 1, DRS depends on the availability of vCLS VMs. See *vSphere Cluster Services (vCLS)* in *vSphere Resource Management* for more information.

vSphere HA and DRS Affinity Rules

If you create a DRS affinity rule for your cluster, you can specify how vSphere HA applies that rule during a virtual machine failover.

The two types of rules for which you can specify vSphere HA failover behavior are the following:

- VM anti-affinity rules force specified virtual machines to remain apart during failover actions.
- VM-Host affinity rules place specified virtual machines on a particular host or a member of a defined group of hosts during failover actions.

When you edit a DRS affinity rule, you must use vSphere HA advanced options to enforce the desired failover behavior for vSphere HA.

- **HA must respect VM anti-affinity rules during failover** -- When the advanced option for VM anti-affinity rules is set, vSphere HA does not fail over a virtual machine if doing so violates a rule. Instead, vSphere HA issues an event reporting there are insufficient resources to perform the failover.
- **HA should respect VM to Host affinity rules during failover** --vSphere HA attempts to place VMs with this rule on the specified hosts if at all possible.

For more information, see vSphere HA Advanced Options.

Note: vSphere HA can restart a VM in a DRS-deactivated cluster, overriding a VM-Host affinity rules mapping if the host failure happens soon (by default, within 5 minutes) after setting the rule.

Other vSphere HA Interoperability Issues

To use vSphere HA, you must be aware of the following additional interoperability issues.

VM Component Protection

VM Component Protection (VMCP) has the following interoperability issues and limitations:

- VMCP does not detect or respond to accessibility issues for files located on vSAN datastores. If a virtual machine's configuration and VMDK files are located only on vSAN datastores, they are not protected by VMCP.
- VMCP does not detect or respond to accessibility issues for files located on Virtual Volume datastores. If a virtual machine's configuration and VMDK files are located only on Virtual Volume datastores, they are not protected by VMCP.
- VMCP does not protect against inaccessible Raw Device Mapping (RDM)s.

IPv6

vSphere HA can be used with IPv6 network configurations, which are fully supported if the following considerations are observed:

- The cluster contains only ESXi 6.0 or later hosts.
- The management network for all hosts in the cluster must be configured with the same IP version, either IPv6 or IPv4. vSphere HA clusters cannot contain both types of networking configuration.
- The network isolation addresses used by vSphere HA must match the IP version used by the cluster for its management network.
- IPv6 cannot be used in vSphere HA clusters that also utilize vSAN.

In addition to the previous restrictions, the following types of IPv6 address types are not supported for use with the vSphere HA isolation address or management network: link-local, ORCHID, and link-local with zone indices. Also, the loopback address type cannot be used for the management network.

Note: To upgrade an existing IPv4 deployment to IPv6, you must first deactivate vSphere HA.

Creating a vSphere HA Cluster

vSphere HA operates in the context of a cluster of ESXi (or legacy ESX) hosts. You must create a cluster, populate it with hosts, and configure vSphere HA settings before failover protection can be established.

When you create a vSphere HA cluster, you must configure a number of settings that determine how the feature works. Before you do this, identify your cluster's nodes. These nodes are the ESXi hosts that will provide the resources to support virtual machines and that vSphere HA will use for failover protection. You should then determine how those nodes are to be connected to one another and to the shared storage where your virtual machine data resides. After that networking architecture is in place, you can add the hosts to the cluster and finish configuring vSphere HA.

You can activate and configure vSphere HA before you add host nodes to the cluster. However, until the hosts are added, your cluster is not fully operational and some of the cluster settings are unavailable. For example, the Specify a Failover Host admission control policy is unavailable until there is a host that can be designated as the failover host.

Note: The Virtual Machine Startup and Shutdown (automatic startup) feature is deactivated for all virtual machines residing on hosts that are in (or moved into) a vSphere HA cluster. Automatic startup is not supported when used with vSphere HA.

vSphere HA Checklist

The vSphere HA checklist contains requirements that you must be aware of before creating and using a vSphere HA cluster.

Review this list before you set up a vSphere HA cluster. For more information, follow the appropriate cross reference.

- All hosts must be licensed for vSphere HA.
- A cluster must contain at least two hosts.
- All hosts must be configured with static IP addresses. If you are using DHCP, you must ensure that the address for each host persists across reboots.
- All hosts must have at least one management network in common. The best practice is to have at least two management networks in common. You should use the VMkernel network with the **Management traffic** checkbox enabled. The networks must be accessible to each other and vCenter Server and the hosts must be accessible to each other on the management networks. See [Best Practices for Networking](#).
- To ensure that any virtual machine can run on any host in the cluster, all hosts must have access to the same virtual machine networks and datastores. Similarly, virtual machines must be located on shared, not local, storage otherwise they cannot be failed over in the case of a host failure.

Note: vSphere HA uses datastore heartbeating to distinguish between partitioned, isolated, and failed hosts. So if some datastores are more reliable in your environment, configure vSphere HA to give preference to them.

- For VM Monitoring to work, VMware tools must be installed. See [VM and Application Monitoring](#).
- vSphere HA supports both IPv4 and IPv6. See [Other vSphere HA Interoperability Issues](#) for considerations when using IPv6.
- For VM Component Protection to work, hosts must have the All Paths Down (APD) Timeout feature enabled.
- To use VM Component Protection, clusters must contain ESXi 6.0 hosts or later.
- Only vSphere HA clusters that contain ESXi 6.0 or later hosts can be used to enable VMCP. Clusters that contain hosts from an earlier release cannot enable VMCP, and such hosts cannot be added to a VMCP-enabled cluster.
- If your cluster uses Virtual Volume datastores, when vSphere HA is enabled a configuration Virtual Volume is created on each datastore by vCenter Server. In these containers, vSphere HA stores the files it uses to protect virtual machines. vSphere HA does not function correctly if you delete these containers. Only one container is created per Virtual Volume datastore.

Create a vSphere HA Cluster in the vSphere Client

To enable your cluster for vSphere HA, you must first create an empty cluster. After you plan the resources and networking architecture of your cluster, use the vSphere Client to add hosts to the cluster and specify the cluster's vSphere HA settings.

- Verify that all virtual machines and their configuration files reside on shared storage.
- Verify that the hosts are configured to access the shared storage so that you can power on the virtual machines by using different hosts in the cluster.
- Verify that hosts are configured to have access to the virtual machine network.
- Verify that you are using redundant management network connections for vSphere HA. For information about setting up network redundancy, see [Best Practices for Networking](#).
- Verify that you have configured hosts with at least two datastores to provide redundancy for vSphere HA datastore heartbeating.
- Connect vSphere Client to vCenter Server by using an account with cluster administrator permissions.

A vSphere HA-enabled cluster is a prerequisite for vSphere Fault Tolerance.

1. In the vSphere Client, browse to the data center where you want the cluster to reside and click **New Cluster**.
2. Complete the **New Cluster** wizard.
Do not turn on vSphere HA (or DRS).
3. Click **OK** to close the wizard and create an empty cluster.
4. Based on your plan for the resources and networking architecture of the cluster, use the vSphere Client to add hosts to the cluster.
5. Browse to the cluster and enable vSphere HA.
 - a) Click the **Configure** tab.
 - b) Select **vSphere Availability** and click **Edit**.
 - c) Select **vSphere HA**.
6. Under **Failures and Responses** select **Enable Host Monitoring**.
With Host Monitoring enabled, hosts in the cluster can exchange network heartbeats and vSphere HA can take action when it detects failures. Host Monitoring is required for the vSphere Fault Tolerance recovery process to work properly.
7. Select a setting for **VM Monitoring**.
Select **VM Monitoring Only** to restart individual virtual machines if their heartbeats are not received within a set time. You can also select **VM and Application Monitoring** to enable application monitoring.
8. Click **OK**.

You have a vSphere HA cluster, populated with hosts.

Configure the appropriate vSphere HA settings for your cluster.

- Failures and responses
- Admission Control
- Heartbeat Datastores
- Advanced Options

See [Configuring vSphere Availability Settings](#).

Configuring vSphere Availability Settings

When you create a vSphere HA cluster or configure an existing cluster, you must configure settings that determine how the feature works.

In the vSphere Client, you can configure following the vSphere HA settings:

Failures and responses	Provide settings here for host failure responses, host isolation, VM monitoring, and VM Component Protection.
Admission Control	Activate or deactivate admission control for the vSphere HA cluster and choose a policy for how it is enforced.
Heartbeat Datastores	Specify preferences for the datastores that vSphere HA uses for datastore heartbeating.
Advanced Options	Customize vSphere HA behavior by setting advanced options.

Configuring Responses to Failures

The **Failure and Responses** pane of the vSphere HA settings allows you to configure how your cluster should function when problems are encountered.

In this part of the vSphere Client, you can determine the specific responses the vSphere HA cluster has for host failures and isolation. You can also configure VM Component Protection (VMCP) actions when Permanent Device Loss (PDL) and All Paths Down (APD) situations occur and you can enable VM monitoring.

The following tasks are available:

Respond to Host Failure

You can set specific responses to host failures that occur in your vSphere HA cluster.

This page is editable only if you have activated vSphere HA.

1. In the vSphere Client, browse to the vSphere HA cluster.
2. Click the **Configure** tab.
3. Select **vSphere Availability** and click **Edit**.
4. Click **Failures and Responses** and then expand **Host Failure Response**.
5. Select from the following configuration options.

Option	Description
Failure Response	If you select Disabled , this setting turns off host monitoring and VMs are not restarted when host failures occur. If Restart VMs is selected, VMs are failed over based on their restart priority when a host fails.
Default VM Restart Priority	The restart priority determines the order in which virtual machines are restarted when the host fails. Higher priority virtual machines are started first. If multiple hosts fail, all virtual machines are migrated from the first host in order of priority, then all virtual machines from the second host in order of priority, and so on.
VM Restart Priority Condition	A specific condition must be selected as well as a delay after that condition has been met, before vSphere HA is allowed to continue to the next VM restart priority.

6. Click **OK**.

Your settings for the host failure response take effect.

Respond to Host Isolation

You can set specific responses to host isolation that occurs in your vSphere HA cluster.

This page is editable only if you have activated vSphere HA.

1. In the vSphere Client, browse to the vSphere HA cluster.
2. Click the **Configure** tab.
3. Select **vSphere Availability** and click **Edit**.
4. Click **Failures and Responses** and expand **Response for Host Isolation**.
5. To configure the host isolation response, select **Disabled**, **Shut down and restart VMs**, or **Power off and restart VMs**.
6. Click **OK**.

Your setting for the host isolation response takes effect.

Configure VMCP Responses

Configure the response that VM Component Protection (VMCP) makes when a datastore encounters a PDL or APD failure.

This page is editable only if you have activated vSphere HA.

1. In the vSphere Client, browse to the vSphere HA cluster.
2. Click the **Configure** tab.
3. Select **vSphere Availability** and click **Edit**.
4. Click **Failures and Responses**, and expand either **Datastore with PDL** or **Datastore with APD**.
5. If you clicked **Datastore with PDL**, you can set the VMCP failure response for this type of issue, either **Disabled**, **Issue Events**, or **Power off and restart VMs**.
6. If you clicked **Datastore with APD**, you can set the VMCP failure response for this type of issue, either **Disabled**, **Issue Events**, **Power off and restart VMs--Conservative restart policy**, or **Power off and restart VMs--Aggressive restart policy**. You can also set **Response recovery**, which is the number of minutes that VMCP waits before taking action.
7. Click **OK**.

Your settings for the VMCP failure response take effect.

Activate VM Monitoring

You can turn on VM and Application Monitoring and also set the monitoring sensitivity for your vSphere HA cluster.

This page is editable only if you have activated vSphere HA.

Note: vSphere HA might failover a non-healthy VM on a healthy host if VM monitoring is deactivated. The host could be selected because of a DRS recommendation or based on which host has the lowest available resources.

1. In the vSphere Client, browse to the vSphere HA cluster.
2. Click the **Configure** tab.
3. Select **vSphere Availability** and click **Edit**.
4. Click **Failures and Responses** and expand **VM Monitoring**.
5. Select **VM Monitoring** and **Application Monitoring**.
These settings turn on VMware Tools heartbeats and application heartbeats, respectively.
6. To set the heartbeat monitoring sensitivity, move the slider between **Low** and **High** or select **Custom** to provide custom settings.
7. Click **OK**.

Your monitoring settings take effect.

Configure Proactive HA

You can configure how Proactive HA responds when a provider has notified its health degradation to vCenter, indicating a partial failure of that host.

This page is editable only if you have enabled vSphere DRS.

1. In the vSphere Client, browse to the Proactive HA cluster.
2. Click the **Configure** tab.
3. Select **vSphere Availability** and click **Edit**.
4. Select **Turn on Proactive HA**.
5. Click **Proactive HA Failures and Responses**.
6. Select from the following configuration options.

Option	Description
Automation Level	Determine whether host quarantine or maintenance mode and VM migrations are recommendations or automatic. • Manual. vCenter Server suggests migration recommendations for virtual machines. • Automated. Virtual machines are migrated to healthy hosts and degraded hosts are entered into quarantine or maintenance mode depending on the configured Proactive HA automation level.
Remediation	Determine what happens to partially degraded hosts. • Quarantine mode for all failures. Balances performance and availability, by avoiding the usage of partially degraded hosts provided that virtual machine performance is unaffected. • Quarantine mode for moderate and Maintenance mode for severe failure (Mixed). Balances performance and availability, by avoiding the usage of moderately degraded hosts provided that virtual machine performance is unaffected. Ensures that virtual machines do not run on severely failed hosts. • Maintenance mode for all failures. Ensures that virtual machines do not run on partially failed hosts.

Option	Description
	Host.Config.Quarantine and Host.Config.Maintenance privileges are required to put hosts in Quarantine mode and Maintenance mode, respectively.

To enable Proactive HA providers for this cluster, select the check boxes. Providers appear when their corresponding vSphere Client plugin has been installed and the providers monitor every host in the cluster. To view or edit the failure conditions supported by the provider, click the edit link.

- Click **OK**.

Configure Admission Control

After you create a cluster, you can configure admission control to specify whether virtual machines can be started if they violate availability constraints. The cluster reserves resources so that failover can occur for all running virtual machines on the specified number of hosts.

The Admission Control page appears only if you activated vSphere HA.

- In the vSphere Client, browse to the vSphere HA cluster.
- Click the **Configure** tab.
- Select **vSphere Availability** and click **Edit**.
- Click **Admission Control** to display the configuration options.
- Select a number for the **Host failures cluster tolerates**. This is the maximum number of host failures that the cluster can recover from or guarantees failover for.
- Select an option for **Define host failover capacity by**.

Option	Description
Cluster resource percentage	Specify a percentage of the cluster's CPU and memory resources to reserve as spare capacity to support failovers.
Slot Policy (powered-on VMs)	Select a slot size policy that covers all powered on VMs or is a fixed size. You can also calculate how many VMs require multiple slots.
Dedicated failover hosts	Select hosts to use for failover actions. Failovers can still occur on other hosts in the cluster if a default failover host does not have enough resources.
Disabled	Select this option to deactivate admission control and allow virtual machine power ons that violate availability constraints.

- Set the percentage for the **Performance degradation VMs tolerate**.
This setting determines what percentage of performance degradation the VMs in the cluster are allowed to tolerate during a failure.
- Click **OK**.

Your admission control settings take effect.

Configure Heartbeat Datastores

vSphere HA uses datastore heartbeating to distinguish between hosts that have failed and hosts that reside on a network partition. With datastore heartbeating, vSphere HA can monitor hosts when a management network partition occurs and continue to respond to failures.

You can specify the datastores that you want to be used for datastore heartbeating.

1. In the vSphere Client, browse to the vSphere HA cluster.
2. Click the **Configure** tab.
3. Select **vSphere Availability** and click **Edit**.
4. Click **Heartbeat Datastores** to display the configuration options for datastore heartbeating.
5. To instruct vSphere HA about how to select the datastores and how to treat your preferences, select from the following options.

Table 374:

Datastore Heartbeating Options
Automatically select datastores accessible from the host
Use datastores only from the specified list
Use datastores from the specified list and complement automatically if needed

6. In the Available heartbeat datastores pane, select the datastores that you want to use for heartbeating.
The listed datastores are shared by more than one host in the vSphere HA cluster. When a datastore is selected, the lower pane displays all the hosts in the vSphere HA cluster that can access it.
7. Click **OK**.

Set Advanced Options

To customize vSphere HA behavior, set advanced vSphere HA options.

Verify that you have cluster administrator privileges.

Note: Because these options affect the functioning of vSphere HA, change them with caution.

1. In the vSphere Client, browse to the vSphere HA cluster.
2. Click the **Configure** tab.
3. Select **vSphere Availability** and click **Edit**.
4. Click **Advanced Options**.
5. Click **Add** and type the name of the advanced option in the text box.
You can set the value of the option in the text box in the Value column.
6. Repeat step 5 for each new option that you want to add and click **OK**.

The cluster uses the options that you added or modified.

Once you have set an advanced vSphere HA option, it persists until you do one the following:

- Using the vSphere Client, reset its value to the default value.
- Manually edit or delete the option from the fdm.cfg file on all hosts in the cluster.

vSphere HA Advanced Options

You can set advanced options that affect the behavior of your vSphere HA cluster.

Table 375: vSphere HA Advanced Options

Option	Description
das.isolationaddress[...]	Sets the address to ping to determine if a host is isolated from the network. This address is pinged only when heartbeats are not received from any other host in the cluster. If not specified, the default gateway of the management network is used. This default gateway has to be a reliable address that is available, so that the host can determine if it is isolated from the network. You can specify multiple isolation addresses (up to 10) for the cluster: das.isolationAddressX, where X = 0-9. Typically you should specify one per management network. Specifying too many addresses makes isolation detection take too long.
das.usedefaultisolationaddress	By default, vSphere HA uses the default gateway of the console network as an isolation address. This option specifies whether or not this default is used (true false).
das.isolationshutdowntimeout	The period of time the system waits for a virtual machine to shut down before powering it off. This only applies if the host's isolation response is Shut down VM. Default value is 300 seconds.
das.slotmeminmb	Defines the maximum bound on the memory slot size. If this option is used, the slot size is the smaller of this value or the maximum memory reservation plus memory overhead of any powered-on virtual machine in the cluster.
das.slotcpuinmhz	Defines the maximum bound on the CPU slot size. If this option is used, the slot size is the smaller of this value or the maximum CPU reservation of any powered-on virtual machine in the cluster.
das.vmmemoryminmb	Defines the default memory resource value assigned to a virtual machine if its memory reservation is not specified or zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default is 0 MB.
das.vmcputminmhz	Defines the default CPU resource value assigned to a virtual machine if its CPU reservation is not specified or zero. This is used for the Host Failures Cluster Tolerates admission control policy. If no value is specified, the default is 32MHz.
das.iostatsinterval	Changes the default I/O stats interval for VM Monitoring sensitivity. The default is 120 (seconds). Can be set to any value greater than, or equal to 0. Setting to 0 deactivates the check. Note: Values of less than 50 are not recommended since smaller values can result in vSphere HA unexpectedly resetting a virtual machine.
das.ignoreinsufficienthbdastore	Deactivates configuration issues created if the host does not have sufficient heartbeat datastores for vSphere HA. Default value is false.

Option	Description
das.heartbeatdsperhost	Changes the number of heartbeat datastores required. Valid values can range from 2-5 and the default is 2.
das.config.fdm.isolationPolicyDelaySec	The number of seconds system waits before executing the isolation policy once it is determined that a host is isolated. The minimum value is 30. If set to a value less than 30, the delay will be 30 seconds.
das.respectvmvmantiaffinityrules	Determines if vSphere HA enforces VM-VM anti-affinity rules. The default value is "true" and rules are enforced even if vSphere DRS is not activated. In this case, vSphere HA does not fail over a virtual machine if doing so violates a rule, but it issues an event reporting there are insufficient resources to perform the failover. This option can also be set to "false", whereby the rules are not enforced. See <i>vSphere Resource Management</i> for more information on anti-affinity rules.
das.maxresets	The maximum number of reset attempts made by VMCP. If a reset operation on a virtual machine affected by an APD situation fails, VMCP retries the reset this many times before giving up
das.maxterminates	The maximum number of retries made by VMCP for virtual machine termination.
das.terminateretryintervalsec	If VMCP fails to terminate a virtual machine, this is the number of seconds the system waits before it retries a terminate attempt
das.config.fdm.reportfailoverfailevent	When set to 1, activates generation of a detailed per-VM event when an attempt by vSphere HA to restart a virtual machine is unsuccessful. Default value is 0. In versions earlier than vSphere 6.0, this event is generated by default.
vpxd.das.completemetadataupdateintervalsec	The period of time (seconds) after a VM-Host affinity rule is set during which vSphere HA can restart a VM in a DRS-deactivated cluster, overriding the rule. Default value is 300 seconds.
das.config.fdm.memReservationMB	By default vSphere HA agents run with a configured memory limit of 250 MB. A host might not allow this reservation if it runs out of reservable capacity. You can use this advanced option to lower the memory limit to avoid this issue. Only integers greater than 100, which is the minimum value, can be specified. Conversely, to prevent problems during primary agent elections in a large cluster (containing 6,000 to 8,000 VMs) you should raise this limit to 325 MB. Note: Once this limit is changed, for all hosts in the cluster you must run the Reconfigure HA task. Also, when a new host is added to the cluster or an existing host is rebooted, this task should be performed on those hosts in order to update this memory setting.
das.reregisterrestartdisabledvms	When vSphere HA is deactivated on a specific VM this option ensures that the VM is registered on another host after a failure. This allows you to power-on that VM without needing to re-register it manually. Note: When this option is used, vSphere HA does not power on the VM, but only registers it.

Option	Description
das.respectvmhostsoftaffinityrules	Determines if vSphere HA restarts a respective VM on a host that belongs to the same VM-Host group. If no such host is available or if the value of this option is set to "false", vSphere HA restarts the VM on any available host in the cluster. In vSphere 6.5 or later, the default value is "true". This value might not be visibly defined in the advanced HA options of the cluster. If you want to deactivate the option, you must manually set this option as "false" in the advanced HA options for the cluster.

Note: If you change the value of any of the following advanced options, you must deactivate and then re-activate vSphere HA before your changes take effect.

- das.isolationaddress[...]
- das.usedefaultisolationaddress
- das.isolationshutdowntimeout

Customize an Individual Virtual Machine

Each virtual machine in a vSphere HA cluster is assigned the cluster default settings for VM Restart Priority, Host Isolation Response, VM Component Protection, and VM Monitoring. You can specify specific behavior for each virtual machine by changing these defaults. If the virtual machine leaves the cluster, these settings are lost.

1. In the vSphere Client, browse to the vSphere HA cluster.
2. Click the **Configure** tab.
3. Under Configuration, select **VM Overrides** and click **Add**.
4. Use the **+** button to select virtual machines to which to apply the overrides.
5. Click **OK**.
6. Optional: You can change other settings, such as the **Automation level**, **VM restart priority**, **Response for Host Isolation**, VMCP settings, **VM Monitoring**, or **VM monitoring sensitivity** settings.

Note: You can view the cluster defaults for these settings by first expanding **Relevant Cluster Settings** and then expanding **vSphere HA**.

7. Click **OK**.

The virtual machine's behavior now differs from the cluster defaults for each setting that you changed.

Best Practices for VMware vSphere® High Availability Clusters

To ensure optimal vSphere HA cluster performance, you must follow certain best practices. This section highlights some of the key best practices for a vSphere HA cluster.

You can also refer to the *vSphere High Availability Deployment Best Practices* publication for further discussion.

Best Practices for Networking

Observe the following best practices for the configuration of host NICs and network topology for vSphere HA. Best Practices include recommendations for your ESXi hosts, and for cabling, switches, routers, and firewalls.

Network Configuration and Maintenance

The following network maintenance suggestions can help you avoid the accidental detection of failed hosts and network isolation because of dropped vSphere HA heartbeats.

- When changing the networks that your clustered ESXi hosts are on, suspend the Host Monitoring feature. Changing your network hardware or networking settings can interrupt the heartbeats that vSphere HA uses to detect host failures, which might result in unwanted attempts to fail over virtual machines.
- When you change the networking configuration on the ESXi hosts themselves, for example, adding port groups, or removing vSwitches, suspend Host Monitoring. After you have made the networking configuration changes, you must reconfigure vSphere HA on all hosts in the cluster, which causes the network information to be reinspected. Then re-enable Host Monitoring.

Note: Because networking is a vital component of vSphere HA, if network maintenance must be performed inform the vSphere HA administrator.

Networks Used for vSphere HA Communications

To identify which network operations might disrupt the functioning of vSphere HA, you must know which management networks are being used for heart beating and other vSphere HA communications.

- On legacy ESX hosts in the cluster, vSphere HA communications travel over all networks that are designated as service console networks. VMkernel networks are not used by these hosts for vSphere HA communications. To contain vSphere HA traffic to a subset of the ESX console networks, use the `allowedNetworks` advanced option.
- On ESXi hosts in the cluster, vSphere HA communications, by default, travel over VMkernel networks. With an ESXi host, if you want to use a network other than the one vCenter Server uses to communicate with the host for vSphere HA, you must explicitly enable the **Management traffic** check box.

To keep vSphere HA agent traffic on the networks you have specified, configure hosts so vmkNICs used by vSphere HA do not share subnets with vmkNICs used for other purposes. vSphere HA agents send packets using any pNIC that is associated with a given subnet when there is also at least one vmkNIC configured for vSphere HA management traffic. Therefore, to ensure network flow separation, the vmkNICs used by vSphere HA and by other features must be on different subnets.

Network Isolation Addresses

A network isolation address is an IP address that is pinged to determine whether a host is isolated from the network. This address is pinged only when a host has stopped receiving heartbeats from all other hosts in the cluster. If a host can ping its network isolation address, the host is not network isolated, and the other hosts in the cluster have either failed or are network partitioned. However, if the host cannot ping its isolation address, it is likely that the host has become isolated from the network and no failover action is taken.

By default, the network isolation address is the default gateway for the host. Only one default gateway is specified, regardless of how many management networks have been defined. Use the `das.isolationaddress[...]` advanced option to add isolation addresses for additional networks. See [vSphere HA Advanced Options](#).

Network Path Redundancy

Network path redundancy between cluster nodes is important for vSphere HA reliability. A single management network ends up being a single point of failure and can result in failovers although only the network has failed. If you have only one management network, any failure between the host and the cluster can cause an unnecessary (or false) failover activity if heartbeat datastore connectivity is not retained during the networking failure. Possible failures include NIC failures, network cable failures, network cable removal, and switch resets. Consider these possible sources of failure between hosts and try to minimize them, typically by providing network redundancy.

The first way you can implement network redundancy is at the NIC level with NIC teaming. Using a team of two NICs connected to separate physical switches improves the reliability of a management network. Because servers connected

through two NICs (and through separate switches) have two independent paths for sending and receiving heartbeats, the cluster is more resilient. To configure a NIC team for the management network, configure the vNICs in vSwitch configuration for Active or Standby configuration. The recommended parameter settings for the vNICs are:

- Default load balancing = route based on originating port ID
- Failback = No

After you have added a NIC to a host in your vSphere HA cluster, you must reconfigure vSphere HA on that host.

In most implementations, NIC teaming provides sufficient heartbeat redundancy, but as an alternative you can create a second management network connection attached to a separate virtual switch. Redundant management networking allows the reliable detection of failures and prevents isolation or partition conditions from occurring, because heartbeats can be sent over multiple networks. The original management network connection is used for network and management purposes. When the second management network connection is created, vSphere HA sends heartbeats over both management network connections. If one path fails, vSphere HA still sends and receives heartbeats over the other path.

Note: Configure the fewest possible number of hardware segments between the servers in a cluster. The goal being to limit single points of failure. Also, routes with too many hops can cause networking packet delays for heartbeats, and increase the possible points of failure.

Using IPv6 Network Configurations

Only one IPv6 address can be assigned to a given network interface used by your vSphere HA cluster. Assigning multiple IP addresses increases the number of heartbeat messages sent by the cluster's primary host with no corresponding benefit.

Best Practices for Interoperability

Observe the following best practices for allowing interoperability between vSphere HA and other features.

vSphere HA and Storage vMotion Interoperability in a Mixed Cluster

In clusters where ESXi 5.x hosts and ESX/ ESXi 4.1 or earlier hosts are present and where Storage vMotion is used extensively or Storage DRS is activated, do not deploy vSphere HA. vSphere HA might respond to a host failure by restarting a virtual machine on a host with an ESXi version different from the one on which the virtual machine was running before the failure. A problem can occur if, at the time of failure, the virtual machine was involved in a Storage vMotion action on an ESXi 5.x host, and vSphere HA restarts the virtual machine on a host with a version earlier than ESXi 5.0. While the virtual machine might power-on, any subsequent attempts at snapshot operations might corrupt the vdisk state and leave the virtual machine unusable.

Using Auto Deploy with vSphere HA

You can use vSphere HA and Auto Deploy together to improve the availability of your virtual machines. Auto Deploy provisions hosts when they power-on and you can also configure it to install the vSphere HA agent on hosts during the boot process. See the Auto Deploy documentation included in vSphere Installation and Setup for details.

Upgrading Hosts in a Cluster Using vSAN

If you are upgrading the ESXi hosts in your vSphere HA cluster to version 5.5 or later, and you also plan to use vSAN, follow this process.

1. Upgrade all of the hosts.
2. Deactivate vSphere HA.
3. Activate vSAN.
4. Re-activate vSphere HA.

Best Practices for Cluster Monitoring

Observe the following best practices for monitoring the status and validity of your vSphere HA cluster.

Setting Alarms to Monitor Cluster Changes

When vSphere HA or Fault Tolerance take action to maintain availability, for example, a virtual machine failover, you can be notified about such changes. Configure alarms in vCenter Server to be triggered when these actions occur, and have alerts, such as emails, sent to a specified set of administrators.

Several default vSphere HA alarms are available.

- Insufficient failover resources (a cluster alarm)
- Cannot find primary (a cluster alarm)
- Failover in progress (a cluster alarm)
- Host HA status (a host alarm)
- VM monitoring error (a virtual machine alarm)
- VM monitoring action (a virtual machine alarm)
- Failover failed (a virtual machine alarm)

Note: The default alarms include the feature name, vSphere HA.

Change in behavior for HA VIBs

In vSphere 7.0 or later, it is possible the HA VIBs might be removed in some cases when HA is activated on a Lifecycle Manager (vLCM) cluster. In previous releases, vCenter would not attempt to remove HA VIBs from ESXi hosts.

This situation can occur only on vLCM clusters with vSphere HA activated. When a vLCM **Remediate** operation occurs (either as a user-initiated operation or an API invocation) after vSphere HA is deactivated on the cluster, the vSphere HA VIBs might be removed as a consequence.

Note: This change in behavior is harmless because vCenter pushes the required vSphere HA VIBs when HA is activated again.

Providing Fault Tolerance for Virtual Machines

You can use vSphere Fault Tolerance for your virtual machines to ensure continuity with higher levels of availability and data protection.

Fault Tolerance is built on the ESXi host platform, and it provides availability by having identical virtual machines run on separate hosts.

To obtain the optimal results from Fault Tolerance you must be familiar with how it works, how to enable it for your cluster, virtual machines and the best practices for its usage.

How Fault Tolerance Works

You can use vSphere Fault Tolerance (FT) for most mission critical virtual machines. FT provides continuous availability for such a virtual machine by creating and maintaining another VM that is identical and continuously available to replace it in the event of a failover situation.

The protected virtual machine is called the Primary VM. The duplicate virtual machine, the Secondary VM, is created and runs on another host. The primary VM is continuously replicated to the secondary VM so that the secondary VM can take over at any point, thereby providing Fault Tolerant protection.

The Primary and Secondary VMs continuously monitor the status of one another to ensure that Fault Tolerance is maintained. A transparent failover occurs if the host running the Primary VM fails, or encounters an uncorrectable

hardware error in the memory of the Primary VM, in which case the Secondary VM is immediately activated to replace the Primary VM. A new Secondary VM is started and Fault Tolerance redundancy is reestablished automatically. If the host running the Secondary VM fails, it is also immediately replaced. In either case, users experience no interruption in service and no loss of data.

A fault tolerant virtual machine and its secondary copy are not allowed to run on the same host. This restriction ensures that a host failure cannot result in the loss of both VMs.

Note: You can also use VM-Host affinity rules to dictate which hosts designated virtual machines can run on. If you use these rules, be aware that for any Primary VM that is affected by such a rule, its associated Secondary VM is also affected by that rule. For more information about affinity rules, see the vSphere Resource Management documentation.

Fault Tolerance avoids "split-brain" situations, which can lead to two active copies of a virtual machine after recovery from a failure. Atomic file locking on shared storage is used to coordinate failover so that only one side continues running as the Primary VM and a new Secondary VM is respawned automatically.

vSphere Fault Tolerance can accommodate symmetric multiprocessor (SMP) virtual machines with up to 8 vCPUs.

Fault Tolerance Use Cases

Several typical situations can benefit from the use of vSphere Fault Tolerance.

Fault Tolerance provides a higher level of business continuity than vSphere HA. When a Secondary VM is called upon to replace its Primary VM counterpart, the Secondary VM immediately takes over the Primary VM's role with the entire state of the virtual machine preserved. Applications are already running, and data stored in memory does not need to be reentered or reloaded. Failover provided by vSphere HA restarts the virtual machines affected by a failure.

This higher level of continuity and the added protection of state information and data informs the scenarios when you might want to deploy Fault Tolerance.

- Applications which must always be available, especially applications that have long-lasting client connections that users want to maintain during hardware failure.
- Custom applications that have no other way of doing clustering.
- Cases where high availability might be provided through custom clustering solutions, which are too complicated to configure and maintain.

Another key use case for protecting a virtual machine with Fault Tolerance can be described as On-Demand Fault Tolerance. In this case, a virtual machine is adequately protected with vSphere HA during normal operation. During certain critical periods, you might want to enhance the protection of the virtual machine. For example, you might be running a quarter-end report which, if interrupted, might delay the availability of critical information. With vSphere Fault Tolerance, you can protect this virtual machine before running this report and then turn off or suspend Fault Tolerance after the report has been produced. You can use On-Demand Fault Tolerance to protect the virtual machine during a critical time period and return the resources to normal during non-critical operation.

Fault Tolerance Requirements, Limits, and Licensing

Before using vSphere Fault Tolerance (FT), consider the high-level requirements, limits, and licensing that apply to this feature.

Requirements

The following CPU and networking requirements apply to FT.

CPUs that are used in host machines for fault tolerant VMs must be compatible with vSphere vMotion. Also, CPUs that support Hardware MMU virtualization (Intel EPT or AMD RVI) are required. The following CPUs are supported.

- Intel Sandy Bridge or later. Avoton is not supported.

- AMD Bulldozer or later.

Use a 10-Gbit logging network for FT and verify that the network is low latency. A dedicated FT network is highly recommended.

Note: Fault Tolerance is currently unsupported to be enabled on any VM utilizing an NSX T created portgroup (vlan or overlay segment). Fault Tolerance is also unsupported for NSX T Managers and Edge nodes.

Limits

In a cluster configured to use Fault Tolerance, two limits are enforced independently.

das.maxftvmsperhost

The maximum number of fault tolerant VMs allowed on a host in the cluster. The default value is 4. There is no FT VMs per host maximum, you can use larger numbers if the workload performs well in FT VMs. You can deactivate checking by setting the value to 0.

das.maxftvcpusperhost

The maximum number of vCPUs aggregated across all fault tolerant VMs on a host. The default value is 8. There is no FT vCPU per host maximum, you can use larger numbers if the workload performs well. You can deactivate checking by setting the value to 0.

Licensing

The number of vCPUs supported by a single fault tolerant VM is limited by the level of licensing that you have purchased for vSphere. Fault Tolerance is supported as follows:

- vSphere Standard and Enterprise. Allows up to 2 vCPUs.
- vSphere Enterprise Plus. Allows up to 8 vCPUs.
- VMware vSphere Foundation (VVF). Allows up to 8 vCPUs.

Note: FT is supported in vSphere Standard, vSphere Enterprise and vSphere Enterprise Plus Editions.

Fault Tolerance Interoperability

Before configuring vSphere Fault Tolerance, you must be aware of the features and products Fault Tolerance cannot interoperate with.

vSphere Features Not Supported with Fault Tolerance

When configuring your cluster, you should be aware that not all vSphere features can interoperate with Fault Tolerance.

The following vSphere features are not supported for fault tolerant virtual machines.

Note: Before vSphere 7.0 Update 2, vSphere Virtual Machine Encryption was not supported with FT.

- Snapshots. Snapshots must be removed or committed before Fault Tolerance can be enabled on a virtual machine. In addition, it is not possible to take snapshots of virtual machines on which Fault Tolerance is enabled.

Note: Disk-only snapshots created for vStorage APIs - Data Protection (VADP) backups are supported with Fault Tolerance. However, legacy FT does not support VADP.

- Storage vMotion. You cannot invoke Storage vMotion for virtual machines with Fault Tolerance turned on. To migrate the storage, you should temporarily turn off Fault Tolerance, and perform the storage vMotion action. When this is complete, you can turn Fault Tolerance back on.
- Linked clones. You cannot use Fault Tolerance on a virtual machine that is a linked clone, nor can you create a linked clone from an FT-enabled virtual machine.

- Virtual Volume datastores.
- Storage-based policy management. Storage policies are supported for vSAN storage.
- I/O filters.
- VBS enabled VMs.
- VM Namespace DB and VM DataSets.

Features and Devices Incompatible with Fault Tolerance

Not all third party devices, features, or products can interoperate with Fault Tolerance.

For a virtual machine to be compatible with Fault Tolerance, the Virtual Machine must not use the following features or devices.

Table 376: Features and Devices Incompatible with Fault Tolerance and Corrective Actions

Incompatible Feature or Device	Corrective Action
Physical Raw Disk mapping (RDM).	With legacy FT you can reconfigure virtual machines with physical RDM-backed virtual devices to use virtual RDMs instead.
CD-ROM or floppy virtual devices backed by a physical or remote device.	Remove the CD-ROM or floppy virtual device or reconfigure the backing with an ISO installed on shared storage.
USB and sound devices.	Remove these devices from the virtual machine.
N_Port ID Virtualization (NPIV).	Deactivate the NPIV configuration of the virtual machine.
NIC passthrough.	This feature is not supported by Fault Tolerance so it must be turned off.
Hot-plugging devices.	The hot plug feature is automatically deactivated for fault tolerant virtual machines. To hot plug devices (either adding or removing), you must momentarily turn off Fault Tolerance, perform the hot plug, and then turn on Fault Tolerance. Note: When using Fault Tolerance, changing the settings of a virtual network card while a virtual machine is running is a hot-plug operation, since it requires "unplugging" the network card and then "plugging" it in again. For example, with a virtual network card for a running virtual machine, if you change the network that the virtual NIC is connected to, FT must be turned off first.
Serial or parallel ports	Remove these devices from the virtual machine.
Video devices that have 3D activated.	Fault Tolerance does not support video devices that have 3D activated.
Virtual Machine Communication Interface (VMCI)	Not supported by Fault Tolerance.
2TB+ VMDK	Fault Tolerance is not supported with a 2TB+ VMDK.

Using Fault Tolerance with DRS

You can use vSphere Fault Tolerance with vSphere Distributed Resource Scheduler (DRS).

FT VMs do not require EVC to support DRS. You can use FT with DRS on vSphere 6.5 and 6.0 hosts that are managed by a vSphere 6.7 or higher VC.

Note: vSphere DRS is a critical feature of vSphere which is required to maintain the health of the workloads running inside vSphere Cluster. Starting with vSphere 7.0 Update 1, DRS depends on the availability of vCLS VMs. See *vSphere Cluster Services (vCLS)* in *vSphere Resource Management* for more information.

Preparing Your Cluster and Hosts for Fault Tolerance

To enable vSphere Fault Tolerance for your cluster, you must meet the feature's prerequisites and you must perform certain configuration steps on your hosts. After those steps are accomplished and your cluster has been created, you can also check that your configuration complies with the requirements for enabling Fault Tolerance.

The tasks you should complete before attempting to set up Fault Tolerance for your cluster include the following:

- Ensure that your cluster, hosts, and virtual machines meet the requirements outlined in the Fault Tolerance checklist.
- Configure networking for each host.
- Create the vSphere HA cluster, add hosts, and check compliance.

After your cluster and hosts are prepared for Fault Tolerance, you are ready to turn on Fault Tolerance for your virtual machines. See [Turn On Fault Tolerance](#).

Fault Tolerance Checklist

The following checklist contains cluster, host, and virtual machine requirements that you need to be aware of before using vSphere Fault Tolerance.

Review this list before setting up Fault Tolerance.

Note: The failover of fault tolerant virtual machines is independent of vCenter Server, but you must use vCenter Server to set up your Fault Tolerance clusters.

Cluster Requirements for Fault Tolerance

You must meet the following cluster requirements before you use Fault Tolerance.

- Fault Tolerance logging and VMotion networking configured. See [Configure Networking for Host Machines](#).
- vSphere HA cluster created and enabled. See [Creating a vSphere HA Cluster](#). vSphere HA must be enabled before you can power on fault tolerant virtual machines or add a host to a cluster that already supports fault tolerant virtual machines.

Host Requirements for Fault Tolerance

You must meet the following host requirements before you use Fault Tolerance.

- Hosts must use supported processors.
- Hosts must be licensed for Fault Tolerance.
- Hosts must be certified for Fault Tolerance. See <http://www.vmware.com/resources/compatibility/search.php> and select **Search by Fault Tolerant Compatible Sets** to determine if your hosts are certified.
- The configuration for each host must have Hardware Virtualization (HV) enabled in the BIOS.

Note: VMware recommends that the hosts you use to support FT VMs have their BIOS power management settings turned to "Maximum performance" or "OS-managed performance".

To confirm the compatibility of the hosts in the cluster to support Fault Tolerance, you can also run profile compliance checks as described in [Create Cluster and Check Compliance](#).

Virtual Machine Requirements for Fault Tolerance

You must meet the following virtual machine requirements before you use Fault Tolerance.

- No unsupported devices attached to the virtual machine. See [Fault Tolerance Interoperability](#).

- Incompatible features must not be running with the fault tolerant virtual machines. See [Fault Tolerance Interoperability](#).
- Virtual machine files (except for the VMDK files) must be stored on shared storage. Acceptable shared storage solutions include Fibre Channel, (hardware and software) iSCSI, vSAN, NFS, and NAS.

Other Configuration Recommendations

You should also observe the following guidelines when configuring Fault Tolerance.

- If you are using NFS to access shared storage, use dedicated NAS hardware with at least a 1Gbit NIC to obtain the network performance required for Fault Tolerance to work properly.
- The memory reservation of a fault tolerant virtual machine is set to the VM's memory size when Fault Tolerance is turned on. Ensure that a resource pool containing fault tolerant VMs has memory resources above the memory size of the virtual machines. Without this excess in the resource pool, there might not be any memory available to use as overhead memory.
- To ensure redundancy and maximum Fault Tolerance protection, you should have a minimum of three hosts in the cluster. In a failover situation, this provides a host that can accommodate the new Secondary VM that is created.

Configure Networking for Host Machines

On each host that you want to add to a vSphere HA cluster, you must configure two different networking switches (vMotion and FT logging) so that the host can support vSphere Fault Tolerance.

Multiple gigabit Network Interface Cards (NICs) are required. For each host supporting Fault Tolerance, a minimum of two physical NICs is recommended. For example, you need one dedicated to Fault Tolerance logging and one dedicated to vMotion. Use three or more NICs to ensure availability. See [Fault Tolerance Requirements, Limits, and Licensing](#).

To set up Fault Tolerance for a host, you must complete this procedure for each port group option (vMotion and FT logging) to ensure that sufficient bandwidth is available for Fault Tolerance logging. Select one option, finish this procedure, and repeat the procedure a second time, selecting the other port group option.

1. In the vSphere Client, browse to the host.
2. Click the **Configure** tab and click **Networking**.
3. Select **VMkernel adapters**.
4. Click the **Add Networking** icon.
5. Provide appropriate information for your connection type.
6. Click **Finish**.

After you create both a vMotion and Fault Tolerance logging virtual switch, you can create other virtual switches, as needed. Add the host to the cluster and complete any steps needed to turn on Fault Tolerance.

Note: If you configure networking to support FT but subsequently suspend the Fault Tolerance logging port, pairs of fault tolerant virtual machines that are powered on remain powered on. If a failover situation occurs, when the Primary VM is replaced by its Secondary VM a new Secondary VM is not started, causing the new Primary VM to run in a Not Protected state.

Using Fault Tolerance

After you have taken all of the required steps for activating vSphere Fault Tolerance for your cluster, you can use the feature by turning it on for individual virtual machines.

Before Fault Tolerance can be turned on, validation checks are performed on a virtual machine.

After these checks are passed and you turn on vSphere Fault Tolerance for a virtual machine, new options are added to the Fault Tolerance section of its context menu. These include turning off or deactivating Fault Tolerance, migrating the Secondary VM, testing failover, and testing restart of the Secondary VM.

Validation Checks for Turning On Fault Tolerance

If the option to turn on Fault Tolerance is available, this task still must be validated and can fail if certain requirements are not met.

Several validation checks are performed on a virtual machine before Fault Tolerance can be turned on.

- SSL certificate checking must be activated in the vCenter Server settings.
- The host must be in a vSphere HA cluster or a mixed vSphere HA and DRS cluster.
- The host must have ESXi 6.x or greater installed.
- The virtual machine must not have snapshots.
- The virtual machine must not be a template.
- The virtual machine must not have vSphere HA deactivated.
- The virtual machine must not have a video device with 3D activated.

Checks for Powered-On Virtual Machines

Several additional validation checks are performed for powered-on virtual machines (or those that are in the process of being powered on).

- The BIOS of the hosts where the fault tolerant virtual machines reside must have Hardware Virtualization (HV) activated.
- The host that supports the Primary VM must have a processor that supports Fault Tolerance.
- Your hardware should be certified as compatible with Fault Tolerance. To confirm that it is, use the VMware Compatibility Guide at <http://www.vmware.com/resources/compatibility/search.php> and select **Search by Fault Tolerant Compatible Sets**.
- The configuration of the virtual machine must be valid for use with Fault Tolerance (for example, it must not contain any unsupported devices).

Secondary VM Placement

When your effort to turn on Fault Tolerance for a virtual machine passes the validation checks, the Secondary VM is created. The placement and immediate status of the Secondary VM depends upon whether the Primary VM was powered-on or powered-off when you turned on Fault Tolerance.

If the Primary VM is powered on:

- The entire state of the Primary VM is copied and the Secondary VM is created, placed on a separate compatible host, and powered on if it passes admission control.
- The Fault Tolerance Status displayed for the virtual machine is **Protected**.

If the Primary VM is powered off:

- The Secondary VM is immediately created and registered to a host in the cluster (it might be re-registered to a more appropriate host when it is powered on.)
- The Secondary VM is not powered on until after the Primary VM is powered on.
- The Fault Tolerance Status displayed for the virtual machine is **Not Protected, VM not Running**.
- When you attempt to power on the Primary VM after Fault Tolerance has been turned on, the additional validation checks listed above are performed.

After these checks are passed, the Primary and Secondary VMs are powered on and placed on separate, compatible hosts. The virtual machine's Fault Tolerance Status is tagged as **Protected**.

Turn On Fault Tolerance

You can turn on vSphere Fault Tolerance through the vSphere Client.

The option to turn on Fault Tolerance is unavailable (dimmed) if any of these conditions apply:

- The virtual machine resides on a host that does not have a license for the feature.
- The virtual machine resides on a host that is in maintenance mode or standby mode.
- The virtual machine is disconnected or orphaned (its .vmx file cannot be accessed).
- The user does not have permission to turn the feature on.

When Fault Tolerance is turned on, vCenter Server resets the virtual machine's memory limit and sets the memory reservation to the memory size of the virtual machine. While Fault Tolerance remains turned on, you cannot change the memory reservation, size, limit, number of vCPUs, or shares. You also cannot add or remove disks for the VM. When Fault Tolerance is turned off, any parameters that were changed are not reverted to their original values. Connect vSphere Client to vCenter Server using an account with cluster administrator permissions.

1. In the vSphere Client, browse to the virtual machine for which you want to turn on Fault Tolerance.
2. Right-click the virtual machine and select **Fault Tolerance > Turn On Fault Tolerance**.
3. Click **Yes**.
4. Select a datastore on which to place the Secondary VM configuration files. Then click **Next**.
5. Select a host on which to place the Secondary VM. Then click **Next**.
6. Review your selections and then click **Finish**.

The specified virtual machine is designated as a Primary VM, and a Secondary VM is established on another host. The Primary VM is now fault tolerant.

Note: The VM datastores and memory are replicated during the FT Turn On process. This can take several minutes depending on the size of the replicated data. The VM state does not appear as protected until replication is complete.

Turn Off Fault Tolerance

Turning off vSphere Fault Tolerance deletes the secondary virtual machine, its configuration, and all history.

Use the **Turn Off Fault Tolerance** option if you do not plan to reenable the feature. Otherwise, use the **Suspend Fault Tolerance** option.

Note: If the Secondary VM resides on a host that is in maintenance mode, disconnected, or not responding, you cannot use the **Turn Off Fault Tolerance** option. In this case, you should suspend and resume Fault Tolerance instead.

1. In the vSphere Client, browse to the virtual machine for which you want to turn off Fault Tolerance.
2. Right-click the virtual machine and select **Fault Tolerance > Turn Off Fault Tolerance**.
3. Click **Yes**.

Fault Tolerance is turned off for the selected virtual machine. The history and the secondary virtual machine for the selected virtual machine are deleted.

Note: Fault Tolerance cannot be turned off when the secondary VM is in the process of being started. Since this involves syncing up the primary VM's full state to the secondary VM, this process may take longer than expected.

Suspend Fault Tolerance

Suspending vSphere Fault Tolerance for a virtual machine suspends its Fault Tolerance protection, but preserves the Secondary VM, its configuration, and all history. Use this option to resume Fault Tolerance protection in the future.

1. In the vSphere Client, browse to the virtual machine for which you want to suspend Fault Tolerance.
2. Right-click the virtual machine and select **Fault Tolerance > Suspend Fault Tolerance**.
3. Click **Yes**.

Fault Tolerance is suspended for the selected virtual machine. Any history and the Secondary VM for the selected virtual machine are preserved and will be used if the feature is resumed.

After you suspend Fault Tolerance, to resume the feature select **Resume Fault Tolerance**.

Migrate Secondary

After vSphere Fault Tolerance is turned on for a Primary VM, you can migrate its associated Secondary VM.

1. In the vSphere Client, browse to the Primary VM for which you want to migrate its Secondary VM.
2. Right-click the virtual machine and select **Fault Tolerance > Migrate Secondary**.
3. Complete the options in the Migrate dialog box and confirm the changes that you made.
4. Click **Finish** to apply the changes.

The Secondary VM associated with the selected fault tolerant virtual machine is migrated to the specified host.

Test Failover

You can induce a failover situation for a selected Primary VM to test your Fault Tolerance protection.

This option is unavailable (dimmed) if the virtual machine is powered off.

1. In the vSphere Client, browse to the Primary VM for which you want to test failover.
2. Right-click the virtual machine and select **Fault Tolerance > Test Failover**.
3. View details about the failover in the Task Console.

This task induces failure of the Primary VM to ensure that the Secondary VM replaces it. A new Secondary VM is also started placing the Primary VM back in a Protected state.

Test Restart Secondary

You can induce the failure of a Secondary VM to test the Fault Tolerance protection provided for a selected Primary VM.

This option is unavailable (dimmed) if the virtual machine is powered off.

1. In the vSphere Client, browse to the Primary VM for which you want to conduct the test.
2. Right-click the virtual machine and select **Fault Tolerance > Test Restart Secondary**.
3. View details about the test in the Task Console.

This task results in the termination of the Secondary VM that provided Fault Tolerance protection for the selected Primary VM. A new Secondary VM is started, placing the Primary VM back in a Protected state.

Upgrade Hosts Used for Fault Tolerance

Use the following procedure to upgrade hosts used for Fault Tolerance.

Verify that you have cluster administrator privileges.

Verify that you have sets of four or more ESXi hosts that are hosting fault tolerant virtual machines that are powered on. If the virtual machines are powered off, the Primary and Secondary VMs can be relocated to hosts with different builds.

Note: This upgrade procedure is for a minimum four-node cluster. The same instructions can be followed for a smaller cluster, though the unprotected interval will be slightly longer.

1. Using vMotion, migrate the fault tolerant virtual machines off of two hosts.
2. Upgrade the two evacuated hosts to the same ESXi build.
3. Suspend Fault Tolerance on the Primary VM.
4. Using vMotion, move the Primary VM for which Fault Tolerance has been suspended to one of the upgraded hosts.
5. Resume Fault Tolerance on the Primary VM that was moved.
6. Repeat 1 to 5 for as many fault tolerant virtual machine pairs as can be accommodated on the upgraded hosts.
7. Using vMotion, redistribute the fault tolerant virtual machines.

All ESXi hosts in a cluster are upgraded.

Activate Fault Tolerance Encryption

You can encrypt Fault Tolerance log traffic.

FT encryption requires SMP-FT. Encryption on Legacy FT (Record-Replay FT) is not supported.

vSphere Fault Tolerance performs frequent checks between a primary VM and secondary VM so that the secondary VM can quickly resume from the last successful checkpoint. The checkpoint contains the VM state that has been modified since the previous checkpoint. You can encrypt Fault Tolerance log traffic.

When you turn on Fault Tolerance, FT encryption is set to **Opportunistic** by default, which means it activates encryption only if both the primary and secondary host are capable of encryption. Follow this procedure if you need to change the FT encryption mode manually.

Note: Fault Tolerance supports vSphere Virtual Machine Encryption with vSphere 7.0 Update 2 and later. In-guest and array-based encryption do not depend on or interfere with VM encryption. Having multiple encryption layers uses additional compute resources, which might impact virtual machine performance. The impact varies with hardware as well as the amount and type of I/O, but overall performance impact is negligible for most workloads. The effectiveness and compatibility of back-end storage features such as deduplication, compression, and replication might also be affected by VM encryption.

1. Select the VM and choose **Edit Settings**.
2. Under **VM Options** select the **Encrypted FT** drop-down menu.
3. Choose one of the following options:

Option	Description
Disabled	Do not turn on encrypted Fault Tolerance logging.
Opportunistic	Turn on encryption only if both sides are capable. A Fault Tolerance VM is allowed to move to an ESXi host which does not support encrypted Fault Tolerance logging.

Option	Description
Required	Choose hosts for Fault Tolerance primary and secondary that both support encrypted FT logging.

Note: While VM encryption is activated, FT encryption mode is set to **Required** by default and cannot be modified.

When FT encryption mode is set to **Required**:

- When you turn on FT, only FT encryption supported hosts are listed for the placement of FT secondary.
- FT failover can only happen on the FT encryption supported hosts.

4. Click **OK**.

Best Practices for Fault Tolerance

To ensure optimal Fault Tolerance results, you should follow certain best practices.

The following recommendations for host and networking configuration can help improve the stability and performance of your cluster.

Host Configuration

Hosts running the Primary and Secondary VMs should operate at approximately the same processor frequencies, otherwise the Secondary VM might be restarted more frequently. Platform power management features that do not adjust based on workload (for example, power capping and enforced low frequency modes to save power) can cause processor frequencies to vary greatly. If Secondary VMs are being restarted on a regular basis, deactivate all power management modes on the hosts running fault tolerant virtual machines or ensure that all hosts are running in the same power management modes.

Host Networking Configuration

The following guidelines allow you to configure your host's networking to support Fault Tolerance with different combinations of traffic types (for example, NFS) and numbers of physical NICs.

- Distribute each NIC team over two physical switches ensuring L2 domain continuity for each VLAN between the two physical switches.
- Use deterministic teaming policies to ensure particular traffic types have an affinity to a particular NIC (active/standby) or set of NICs (for example, originating virtual port-id).
- Where active/standby policies are used, pair traffic types to minimize impact in a failover situation where both traffic types will share a vmnic.
- Where active/standby policies are used, configure all the active adapters for a particular traffic type (for example, FT Logging) to the same physical switch. This minimizes the number of network hops and lessens the possibility of oversubscribing the switch to switch links.

Note: FT logging traffic between Primary and Secondary VMs is unencrypted and contains guest network and storage I/O data, as well as the memory contents of the guest operating system. This traffic can include sensitive data such as passwords in plaintext. To avoid such data being divulged, ensure that this network is secured, especially to avoid 'man-in-the-middle' attacks. For example, you could use a private network for FT logging traffic.

Homogeneous Clusters

vSphere Fault Tolerance can function in clusters with nonuniform hosts, but it works best in clusters with compatible nodes. When constructing your cluster, all hosts should have the following configuration:

- Common access to datastores used by the virtual machines.
- The same virtual machine network configuration.
- The same BIOS settings (power management and hyperthreading) for all hosts.

Run **Check Compliance** to identify incompatibilities and to correct them.

Performance

To increase the bandwidth available for the logging traffic between Primary and Secondary VMs use a 10Gbit NIC, and activate the use of jumbo frames.

You can select multiple NICs for the FT logging network. By selecting multiple NICs, you can take advantage of the bandwidth from multiple NICs even if all of the NICs are not dedicated to running FT.

Store ISOs on Shared Storage for Continuous Access

Store ISOs that are accessed by virtual machines with Fault Tolerance activated on shared storage that is accessible to both instances of the fault tolerant virtual machine. If you use this configuration, the CD-ROM in the virtual machine continues operating normally, even when a failover occurs.

Avoid Network Partitions

A network partition occurs when a vSphere HA cluster has a management network failure that isolates some of the hosts from vCenter Server and from one another. See [Network Partitions](#) . When a partition occurs, Fault Tolerance protection might be degraded.

In a partitioned vSphere HA cluster using Fault Tolerance, the Primary VM (or its Secondary VM) could end up in a partition managed by a primary host that is not responsible for the virtual machine. When a failover is needed, a Secondary VM is restarted only if the Primary VM was in a partition managed by the primary host responsible for it.

To ensure that your management network is less likely to have a failure that leads to a network partition, follow the recommendations in [Best Practices for Networking](#).

Using vSAN Datastores

vSphere Fault Tolerance can use vSAN datastores, but you must observe the following restrictions:

- A mix of vSAN and other types of datastores is not supported for both Primary VMs and Secondary VMs.

To increase performance and reliability when using FT with vSAN, the following conditions are also recommended.

- vSAN and FT should use separate networks.
- Keep Primary and Secondary VMs in separate vSAN fault domains.

Activate Metro Cluster Fault Tolerance

In vSphere 8.0 U3 you can activate Metro Cluster Fault Tolerance in the Fault Tolerance wizard.

In the fault tolerance wizard, you can select a checkbox labeled **Enable Metro Cluster Fault Tolerance** to activate FT Metro Cluster functionality and a drop-down list for choosing a HostGroup as the preferred location of FT VM. By default, the checkbox is unchecked and the drop-down list is deactivated, indicating that FT Metro Cluster is deactivated for the VM (`ConfigInfo.metroFtEnabled` is 'FALSE').

When the checkbox is checked, the drop-down list will be activated for choosing a HostGroup. The wizard will prevent proceeding to the next step if no HostGroup is selected for the VM. To ensure the legality of the selected HostGroup, the wizard will call the function `queryFaultToleranceCompatibleHosts` and retrieve the result through the returned messages.

Before checking the checkbox **Enable Metro Cluster Fault Tolerance**, the drop-down list of HostGroup is deactivated. A signpost button will be added next to the label `Host Group`, the signpost content is: Based on the selected Host Group, FT can divide the hosts in the cluster into two groups and ensure to place FT primary and FT secondary in different groups.

After checking the checkbox **Enable Metro Cluster Fault Tolerance**, the drop-down list is activated in the wizard for the HostGroup selection. FT handles the validation check for the FT Metro Cluster flag and HostGroup, which is triggered by clicking the **NEXT** button. The wizard sets **FaultToleranceConfigSpec.metroFtEnabled** to **TRUE** and **FaultToleranceConfigSpec.preferredLocation** to the selected HostGroup. Then, the wizard fetches the compatible host list.

After checking the checkbox **Enable Metro Cluster Fault Tolerance**, but not selecting any HostGroup, the wizard can't be moved forward with **NEXT** button with the error message: `Please assign a host group for Fault Tolerance VM before enabling Metro Cluster`. FT also checks the HostGroup for FT Metro Cluster activated VM and the task might fail if the HostGroup is not configured.

You can remove the configured host group while FT VM is running. FT Metro Cluster is deactivated in such case. However, the host group name is still displayed on the FT information card, since FT Metro Cluster can be re-activated after configuring the host group again.

If you delete the HostGroup for a running FT VM, the FT Metro Cluster is deactivated. In case of a failover, `Metro Cluster status Not Applicable Host Group is missing`. is displayed on the FT information card. However, if you add the HostGroup back, the FT Metro Cluster will be activated again.

If you delete the HostGroup for a powered-off FT VM, the Fault Tolerance Secondary VM cannot be powered-on.

Legacy Fault Tolerance

Legacy FT VMs can exist only on ESXi hosts that are running on vSphere versions earlier than 6.5.

ESXi hosts prior to version 6.5 supported vSphere Fault Tolerance based on a different technology. If you are using this form of Fault Tolerance and need to continue doing so, we recommend you reserve a vCenter 6.0 instance to manage the pool of pre-6.5 hosts required to run these VMs. vCenter 6.0 was the last version fully capable of managing legacy FT protected VMs. For more information on Legacy Fault Tolerance, see vSphere Availability 6.0 documentation.

Troubleshooting Fault Tolerant Virtual Machines

To maintain a high level of performance and stability for your fault tolerant virtual machines and also to minimize failover rates, you should be aware of certain troubleshooting issues.

The troubleshooting topics discussed focus on problems that you might encounter when using the vSphere Fault Tolerance feature on your virtual machines. The topics also describe how to resolve problems.

You can also see the VMware knowledge base article at <http://kb.vmware.com/kb/1033634> to help you troubleshoot Fault Tolerance. This article contains a list of error messages that you might encounter when you attempt to use the feature and, where applicable, advice on how to resolve each error.

Hardware Virtualization Not Enabled

You must enable Hardware Virtualization (HV) before you use vSphere Fault Tolerance.

When you attempt to power on a virtual machine with Fault Tolerance enabled, an error message might appear if you did not enable HV.

This error is often the result of HV not being available on the ESXi server on which you are attempting to power on the virtual machine. HV might not be available either because it is not supported by the ESXi server hardware or because HV is not enabled in the BIOS.

If the ESXi server hardware supports HV, but HV is not currently enabled, enable HV in the BIOS on that server. The process for enabling HV varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to enable HV.

If the ESXi server hardware does not support HV, switch to hardware that uses processors that support Fault Tolerance.

Compatible Hosts Not Available for Secondary VM

If you power on a virtual machine with Fault Tolerance enabled and no compatible hosts are available for its Secondary VM, you might receive an error message.

You might encounter the following error message:

```
Secondary VM could not be powered on as there are no compatible hosts that can accommodate it.
```

This can occur for a variety of reasons including that there are no other hosts in the cluster, there are no other hosts with HV enabled, Hardware MMU Virtualization is not supported by host CPUs, data stores are inaccessible, there is no available capacity, or hosts are in maintenance mode.

If there are insufficient hosts, add more hosts to the cluster. If there are hosts in the cluster, ensure they support HV and that HV is enabled. The process for enabling HV varies among BIOSes. See the documentation for your hosts' BIOSes for details on how to enable HV. Check that hosts have sufficient capacity and that they are not in maintenance mode.

Secondary VM on Overcommitted Host Degrades Performance of Primary VM

If a Primary VM appears to be executing slowly, even though its host is lightly loaded and retains idle CPU time, check the host where the Secondary VM is running to see if it is heavily loaded.

When a Secondary VM resides on a host that is heavily loaded, the Secondary VM can affect the performance of the Primary VM.

A Secondary VM running on a host that is overcommitted (for example, with its CPU resources) might not get the same amount of resources as the Primary VM. When this occurs, the Primary VM must slow down to allow the Secondary VM to keep up, effectively reducing its execution speed to the slower speed of the Secondary VM.

If the Secondary VM is on an overcommitted host, you can move the VM to another location without resource contention problems. Or more specifically, do the following:

- For FT networking contention, use vMotion technology to move the Secondary VM to a host with fewer FT VMs contending on the FT network. Verify that the quality of the storage access to the VM is not asymmetric.
- For storage contention problems, turn FT off and on again. When you recreate the Secondary VM, change its datastore to a location with less resource contention and better performance potential.
- To resolve a CPU resources problem, set an explicit CPU reservation for the Primary VM at an MHz value sufficient to run its workload at the desired performance level. This reservation is applied to both the Primary and Secondary VMs, ensuring that both VMs can execute at a specified rate. For guidance in setting this reservation, view the performance graphs of the virtual machine (before Fault Tolerance was enabled) to see how many CPU resources it used under normal conditions.

Increased Network Latency Observed in FT Virtual Machines

If your FT network is not optimally configured, you might experience latency problems with the FT VMs.

FT VMs might see a variable increase in packet latency (on the order of milliseconds). Applications that demand very low network packet latency or jitter (for example, certain real-time applications) might see a degradation in performance.

Some increase in network latency is expected overhead for Fault Tolerance, but certain factors can add to this latency. For example, if the FT network is on a particularly high latency link, this latency is passed on to the applications. Also, if the FT network has insufficient bandwidth (fewer than 10 Gbps), greater latency might occur.

Verify that the FT network has sufficient bandwidth (10 Gbps or more) and uses a low latency link between the Primary VM and Secondary VM. These precautions do not eliminate network latency, but minimize its potential impact.

Some Hosts Are Overloaded with FT Virtual Machines

You might encounter performance problems if your cluster's hosts have an imbalanced distribution of FT VMs.

Some hosts in the cluster might become overloaded with FT VMs, while other hosts might have unused resources.

vSphere DRS does not load balance FT VMs (unless they are using legacy FT). This limitation might result in a cluster where hosts are unevenly distributed with FT VMs.

Manually rebalance the FT VMs across the cluster by using vSphere vMotion. Generally, the fewer FT VMs that are on a host, the better they perform, due to reduced contention for FT network bandwidth and CPU resources.

Losing Access to FT Metadata Datastore

Access to the Fault Tolerance metadata datastore is essential for the proper functioning of an FT VM. Loss of this access can cause a variety of problems.

These problems include the following:

- FT can terminate unexpectedly.
- If both the Primary VM and Secondary VM cannot access the metadata datastore, the VMs might fail unexpectedly. Typically, an unrelated failure that terminates FT must also occur when access to the FT metadata datastore is lost by both VMs. vSphere HA then tries to restart the Primary VM on a host with access to the metadata datastore.
- The VM might stop being recognized as an FT VM by vCenter Server. This failed recognition can allow unsupported operations such as taking snapshots to be performed on the VM and cause problematic behavior.

Lack of access to the Fault Tolerance metadata datastore can lead to the undesirable outcomes in the previous list.

When planning your FT deployment, place the metadata datastore on highly available storage. While FT is running, if you see that the access to the metadata datastore is lost on either the Primary VM or the Secondary VM, promptly address the storage problem before loss of access causes one of the previous problems. If a VM stops being recognized as an FT VM by vCenter Server, do not perform unsupported operations on the VM. Restore access to the metadata datastore. After access is restored for the FT VMs and the refresh period has ended, the VMs are recognizable.

Turning On vSphere FT for Powered-On VM Fails

If you try to turn on vSphere Fault Tolerance for a powered-on VM, this operation can fail.

When you select **Turn On Fault Tolerance** for a powered-on VM, the operation fails and you see an `Unknown error` message.

This operation can fail if the host that the VM is running on has insufficient memory resources to provide fault tolerant protection. vSphere Fault Tolerance automatically tries to allocate a full memory reservation on the host for the VM. Overhead memory is required for fault tolerant VMs and can sometimes expand to 1 to 2 GB. If the powered-on VM is running on a host that has insufficient memory resources to accommodate the full reservation plus the overhead memory, trying to turn on Fault Tolerance fails. Subsequently, the `Unknown error` message is returned.

Choose from these solutions:

- Free up memory resources on the host to accommodate the VM's memory reservation and the added overhead.
- Move the VM to a host with ample free memory resources and try again.

FT Virtual Machines not Placed or Evacuated by vSphere DRS

FT virtual machines in a cluster that is activated with vSphere DRS do not function correctly if Enhanced vMotion Compatibility (EVC) is currently deactivated.

Because EVC is a prerequisite for using DRS with FT VMs, DRS does not place or evacuate them if EVC has been deactivated (even if it is later reactivated).

When EVC is deactivated on a DRS cluster, a VM override that deactivates DRS on an FT VM might be added. Even if EVC is later reactivated, this override is not canceled.

If DRS does not place or evacuate FT VMs in the cluster, check the VMs for a VM override that is deactivating DRS. If you find one, remove the override that is deactivating DRS.

Note: For more information on how to edit or delete VM overrides, see *vSphere Resource Management*.

Fault Tolerant Virtual Machine Failovers

A Primary or Secondary VM can fail over even though its ESXi host has not crashed. In such cases, virtual machine execution is not interrupted, but redundancy is temporarily lost. To avoid this type of failover, be aware of some of the situations when it can occur and take steps to avoid them.

Partial Hardware Failure Related to Storage

This problem can arise when access to storage is slow or down for one of the hosts. When this occurs there are many storage errors listed in the VMkernel log. To resolve this problem you must address your storage-related problems.

Partial Hardware Failure Related to Network

If the logging NIC is not functioning or connections to other hosts through that NIC are down, this can trigger a fault tolerant virtual machine to be failed over so that redundancy can be reestablished. To avoid this problem, dedicate a separate NIC each for vMotion and FT logging traffic and perform vMotion migrations only when the virtual machines are less active.

Insufficient Bandwidth on the Logging NIC Network

This can happen because of too many fault tolerant virtual machines being on a host. To resolve this problem, more broadly distribute pairs of fault tolerant virtual machines across different hosts.

Use a 10-Gbit logging network for FT and verify that the network is low latency.

vMotion Failures Due to Virtual Machine Activity Level

If the vMotion migration of a fault tolerant virtual machine fails, the virtual machine might need to be failed over. Usually, this occurs when the virtual machine is too active for the migration to be completed with only minimal disruption to the activity. To avoid this problem, perform vMotion migrations only when the virtual machines are less active.

Too Much Activity on VMFS Volume Can Lead to Virtual Machine Failovers

When a number of file system locking operations, virtual machine power ons, power offs, or vMotion migrations occur on a single VMFS volume, this can trigger fault tolerant virtual machines to be failed over. A symptom that this might be occurring is receiving many warnings about SCSI reservations in the VMkernel log. To resolve this problem, reduce the number of file system operations or ensure that the fault tolerant virtual machine is on a VMFS volume that does not have an abundance of other virtual machines that are regularly being powered on, powered off, or migrated using vMotion.

Lack of File System Space Prevents Secondary VM Startup

Check whether or not your `/(root)` or `/vmfs/datastore` file systems have available space. These file systems can become full for many reasons, and a lack of space might prevent you from being able to start a new Secondary VM.

vCenter High Availability

vCenter High Availability (vCenter HA) protects vCenter Server against host and hardware failures. The active-passive architecture of the solution can also help you reduce downtime significantly when you patch vCenter Server.

After some network configuration, you create a three-node cluster that contains Active, Passive, and Witness nodes. Different configuration paths are available. What you select depends on your existing configuration.

Plan the vCenter HA Deployment

Before you can configure vCenter HA, you have to consider several factors. A deployment with components that use different versions of vSphere requires different considerations than a deployment that includes only vSphere 8.0 components. Resource and software requirements and the networking setup must also be considered carefully.


vCenter Architecture Overview

A vCenter HA cluster consists of three vCenter Server instances. The first instance, initially used as the Active node, is cloned twice to a Passive node and to a Witness node. Together, the three nodes provide an active-passive failover solution.

Deploying each of the nodes on a different ESXi instance protects against hardware failure. Adding the three ESXi hosts to a DRS cluster can further protect your environment.

When vCenter HA configuration is complete, only the Active node has an active management interface (public IP). The three nodes communicate over a private network called vCenter HA network that is set up as part of configuration. The Active node is continuously replicating data to the Passive node.

Figure 140: vCenter Three-Node Cluster



All three nodes are necessary for the functioning of this feature. Compare the node responsibilities.

Table 377: vCenter HA Nodes

Node	Description
Active	- Runs the active vCenter Server instance - Uses a public IP address for the management interface - Uses the vCenter HA network for replication of data to the Passive node. - Uses the vCenter HA network to communicate with the Witness node.
Passive	Is initially a clone of the Active node

Node	Description
	- Constantly receives updates from and synchronizes state with the Active node over the vCenter HA network - Automatically takes over the role of the Active node if a failure occurs
Witness	- Is a lightweight clone of the Active node - Provides a quorum to protect against a split-brain situations

vCenter HA Hardware and Software Requirements

Before you set up vCenter HA, ensure that you have sufficient memory, CPU, and datastore resources, and ensure that you are using versions of vCenter Server and ESXi that support vCenter HA.

Your environment must meet the following requirements.

Table 378: vCenter HA Requirements

Component	Requirements
ESXi	- ESXi 6.0 or later is required. - A minimum of three ESXi hosts is strongly recommended. Each vCenter HA node can then run on a different host for better protection.
Management vCenter Server (if used)	Your environment can include a management vCenter Server system, or you can set up your vCenter Server to manage the ESXi host on which it runs (self-managed vCenter Server) vCenter Server 6.0 or later is required.
vCenter Server	- vCenter Server 6.5 or later is required. - Deployment size Small (4 CPU and 16GB RAM) or bigger is required to meet the RTO. Do not use Tiny in production environments. - vCenter HA is supported and tested with VMFS, NFS, and vSAN datastores. - Ensure you have enough disk space to collect and store support bundles for all three nodes on the Active node. See Collecting Support Bundles for a vCenter HA Node.
Network connectivity	- vCenter HA network latency between Active, Passive, and Witness nodes must be less than 10 ms. - The vCenter HA network must be on a different subnet than the management network.
Licensing required for vCenter HA	- vCenter HA requires a single vCenter Server license. - vCenter HA requires a Standard license.

vCenter HA Configuration Workflow

To configure the passive and witness nodes in your vCenter HA environment, you can use the **Set Up vCenter HA** wizard. With the manual option, you must manually clone the active node and create the passive and witness nodes.

Note: For automatic and manual configurations, verify that your environment is satisfying the hardware and software requirements. See [vCenter HA Hardware and Software Requirements](#).

Automatic Configuration

You can use the **Set Up vCenter HA** wizard to automatically create the passive and witness nodes as part of vCenter HA configuration.

1. You deploy the first vCenter Server, which will become the active node.

2. You add a second network (port group) for vCenter HA traffic on each ESXi host.
3. You start the vCenter HA configuration and supplies the IP addresses, the target ESXi host or cluster, and the datastore for each clone.
4. The system clones the active node and creates a passive node with precisely the same settings, including the same host name.
5. The system clones the active node again and creates a more light-weight witness node.
6. The system sets up the vCenter HA network on which the three nodes communicate, for example, by exchanging heartbeats and other information.

Manual Configuration

If you want more control over your deployment, you can perform a manual vCenter HA configuration, where you clone the active node and create the passive and witness nodes of your configuration. If you select this option and remove the vCenter HA configuration later, you are responsible for deleting the nodes that you created.

1. You deploy the first vCenter Server, which becomes the active node.
2. You add a second network (port group) for vCenter HA traffic on each ESXi host.
3. You must add a second network adapter (NIC) to the active node if the credentials of the active management vCenter Server are unknown.
4. You log in to the vCenter Server (active node) with the vSphere Client.
5. You start the vCenter HA configuration, select the checkbox to manually configure and supply IP address and subnet information for the passive and witness nodes. Optionally, you can override the failover management IP addresses.
6. After the vCenter HA deployment task finishes, you log in to the management vCenter Server and create two clones of the vCenter Server (active node).
7. The system sets up the vCenter HA network on which the three nodes exchange heartbeats and replication information.
8. The vCenter Server is protected by vCenter HA.

For more information, see [Configure vCenter HA With the vSphere Client](#).

Configure the Network

Regardless of the deployment option and inventory hierarchy that you select, you have to set up your network before you can start configuration. To set the foundation for the vCenter HA network, you add a port group to each ESXi host.

- The vCenter Server that later becomes the Active node, is deployed.
- You can access and have privileges to modify that vCenter Server and the ESXi host on which it runs.
- During network setup, you need static IP addresses for the management network. The management and cluster network addresses must be IPv4 or IPv6. They cannot be mixed mode IP addresses.

After configuration is complete, the vCenter HA cluster has two networks, the management network on the first virtual NIC and the vCenter HA network on the second virtual NIC.

Management network

The management network serves client requests (public IP). The management network IP addresses must be static.

vCenter HA network

The vCenter HA network connects the Active, Passive, and Witness nodes and replicates the server state. It also monitors heartbeats.

- The vCenter HA network IP addresses for the Active, Passive, and Witness nodes must be static.

- The vCenter HA network must be on a different subnet than the management network. The three nodes can be on the same subnet or on different subnets.
- Network latency between the Active, Passive, and Witness nodes must be less than 10 milliseconds.
- You must not add a default gateway entry for the cluster network.

Note: vCenter HA setup is not supported when using a Wide Area Network (WAN) for the networking configuration. VCHA requires a dedicated, high-fidelity, and low-latency network. For more information, see [KB 80551](#).

1. Log in to the management vCenter Server and find the ESXi host on which the Active node is running.
2. Add a port group to the ESXi host.
This port group can be on an existing virtual switch or, for improved network isolation, you can create a new virtual switch. It must be different from the management network.
3. If your environment includes the recommended three ESXi hosts, add the port group to each of the hosts.

Configure vCenter HA With the vSphere Client

When you use the vSphere Client, the **Set Up vCenter HA** wizard creates and configures a second network adapter on the vCenter Server, clones the Active node, and configures the vCenter HA network.

- Deploy vCenter Server that you want to use as the initial Active node.
 - The vCenter Server must have a static IP address.
 - SSH must be activated on the vCenter Server.
- Verify that your environment meets the following requirements.
 - The vCenter Server that will become the Active node is managing its own ESXi host and its own virtual machine. This configuration is sometimes called a self-managed vCenter Server.
- Set up the infrastructure for the vCenter HA network. See [Configure the Network](#).
- Determine which static IP addresses to use for the two vCenter Server nodes that will become the Passive node and Witness node.

Note: In order to use an NSX-T segment on the Active node, you must create NIC2/eth1 by using **Edit VM Settings** to add the second NIC with the NSX-T segment. You do not need to specify any resources for Passive or Witness nodes, because the Clone must be created using **Clone VM** after adding necessary Guest customization specifications for Passive and Witness containing NIC1/eth0 and NIC2/eth1 with IP addresses. When you configure VCHA IP addresses for eth1 in vCenter Server, the eth1 on the Active Node is automatically filled in.

1. Log in to the Active node with the vSphere Client.
2. Select the vCenter Server object in the inventory and select the **Configure** tab.
3. Select **vCenter HA** under settings.
4. Click on the **Set Up vCenter HA** button to start the setup wizard.
 - If the vCenter server is self-managed, the **Resource settings** page is displayed. Proceed to step 7.
 - If your vCenter server is managed by another vCenter server in the same SSO domain, proceed to step 7.
 - If your vCenter server is managed by another vCenter server in a different SSO domain, input the location and credential details of that management vCenter server.
5. Click **Management vCenter Server credentials**. Specify the Management vCenter server FQDN or IP address, Single Sign-On user name and password and click **Next**.
If you do not have the Single Sign-On administrator credentials, select the second bullet and click **Next**.

6. You may see a **Certificate warning** displayed. Review the SHA1 thumbprint and select **Yes** to continue.
7. In the **Resource settings** section, first select the vCenter HA network for the active node from the drop-down menu.
Note: The network selector is no longer visible once NIC2/eth1 is created.
8. Click on the checkbox if you want to automatically create clones for Passive and Witness nodes.
Note: If you do not select the checkbox, you must manually create clones for Passive and Witness nodes after you click **Finish**.
9. For the Passive node, click **Edit**.
 - a) Specify a unique name and target location.
 - b) Select the destination compute resource for the operation.
 - c) Select the datastore in which to store the configuration and disk files.
 - d) Select virtual machine Management (NIC 0) and vCenter HA (NIC 1) networks.
If there are issues with your selections, errors or compatibility warnings are displayed.
 - e) Review your selections and click **Finish**.
10. For the Witness node, click **Edit**.
 - a) Specify a unique name and target location.
 - b) Select the destination compute resource for the operation.
 - c) Select the datastore in which to store the configuration and disk files.
 - d) Select vCenter HA (NIC 1) network.
If there are issues with your selections, errors or compatibility warnings are displayed.
 - e) Review your selections and click **Finish**.
11. Click **Next**.
12. In the **IP settings** section, select the IP version from the drop-down menu.
13. Enter the IPv4 address (NIC 1) and Subnet mask or prefix length information for the Active, Passive and Witness nodes.
You can Edit management network settings for the Passive Node. Customizing these settings are optional. By default, the management network settings of the Active node are applied.
14. Click **Finish**.

The Passive and Witness nodes are created. When **Set Up vCenter HA** is complete, vCenter Server has high availability protection. After vCenter HA is activated, you can click **Edit** to enter Maintenance Mode, Enable or Disable vCenter HA. There are separate buttons to remove vCenter HA or initiate vCenter HA failover.

See [Manage the vCenter HA Configuration](#) for a list of cluster management tasks.

For a brief overview of enhancements in the vSphere Client when working with vCenter HA, see:

Manage the vCenter HA Configuration

After you configure your vCenter HA cluster, you can perform management tasks. These tasks include certificate replacement, replacement of SSH keys, and SNMP setup. You can also edit the cluster configuration to deactivate or activate vCenter HA, enter maintenance mode, and remove the cluster configuration.

Set Up SNMP Traps

You can set up Simple Network Management Protocol (SNMP) traps to receive SNMP notifications for your vCenter HA cluster.

The traps default to SNMP version 1.

Set up SNMP traps for the Active node and the Passive node. You tell the agent where to send related traps, by adding a target entry to the `snmpd` configuration.

1. Log in to the Active node by using the Virtual Machine Console or SSH.
2. Run the `vicfg-snmp` command, for example:

```
vicfg-snmp -t 10.160.1.1@1166/public
```

In this example, `10.160.1.1` is the client listening address, `1166` is the client listening port, and `public` is the community string.

3. Activate the SNMP agent (`snmpd`) by running the following command.

```
vicfg-snmp -e
```

You also might find these commands useful.

- To view the complete help for the command, run `vicfg-snmp -h`.
- To deactivate the SNMP agent, run `vicfg-snmp -D`.
- To show the SNMP agent's configuration, run `vicfg-snmp -s`.
- To reset the configuration to the default, run `vicfg-snmp -r`.

Set Up Your Environment to Use Custom Certificates

The machine SSL certificate on each node is used for cluster management communication and for encryption of replication traffic. If you want to use custom certificates, you have to remove the vCenter HA configuration, delete the Passive and Witness nodes, provision the Active node with the custom certificate, and reconfigure the cluster.

If possible, replace certificates in the vCenter Server that will become the Active node before you clone the node.

1. Edit the cluster configuration and select **Remove**.
2. Delete the Passive node and the Witness node.
3. On the Active node, which is now a standalone vCenter Server, replace the machine SSL Certificate with a custom certificate.
4. Reconfigure the cluster.

Manage vCenter HA SSH Keys

vCenter HA uses SSH keys for password-less authentication between the Active, Passive, and Witness nodes. The authentication is used for heartbeat exchange and file and data replication. To replace the SSH keys in the nodes of a vCenter HA cluster, you deactivate the cluster, generate new SSH keys on the Active node, transfer the keys to the passive node, and activate the cluster.

1. Edit the cluster and change the mode to **Disabled**.
2. Log in to the Active node by using the Virtual Machine Console or SSH.
3. Activate the bash shell.

```
bash
```

4. Optionally, back up the contents of the `/home/vcha/.ssh` folder to the Active node.
5. To stop VCHA service on the Active node, run the following command.

```
vmon-cli --stop vcha
```

6. To generate new SSH keys on the Active node, run the following command.

```
/usr/sbin/vcha-reset-ssh-keys
```

- Copy the SSH keys to the Passive node and Witness node.

```
scp /home/vcha/.ssh/*
```

- To start the VCHA service on the Active node, run the following command.

```
vmon-cli --start vcha
```

- Edit the cluster configuration and set the vCenter HA cluster to **Enabled**.

The VCHA state changes to **Healthy**.

Initiate a vCenter HA Failover

You can manually initiate a failover and have the Passive node become the Active node.

A vCenter HA cluster supports two types of failover.

Automatic failover

The Passive node attempts to take over the active role in case of an Active node failure.

Manual failover

The user can force a Passive node to take over the active role by using the Initiate Failover action.

Initiate a manual failover for troubleshooting and testing.

- Log in to the Active node vCenter Server with the vSphere Client and click **Configure** for the vCenter Server where you need to initiate failover.
- Under **Settings** select **vCenter HA** and click **Initiate Failover**.
- Click **Yes** to start the failover.
A dialog offers you the option to force a failover without synchronization. In most cases, performing synchronization first is best.
- After the failover, you can verify that the Passive node has the role of the Active node in the vSphere Client.

Edit the vCenter HA Cluster Configuration

When you edit the vCenter HA cluster configuration, you can deactivate or activate the cluster, place the cluster in maintenance mode, or remove the cluster.

Verify that the vCenter HA cluster is deployed and contains the Active, Passive, and Witness nodes.

The operating mode of a vCenter Server controls the failover capabilities and state replication in a vCenter HA cluster.

A vCenter HA cluster can operate in one of the following modes.

Table 379: vCenter HA Cluster Modes of Operation

Mode	Automatic Failover	Manual Failover	Replication	
Enabled	Yes	Yes	Yes	This default mode of operation protects the vCenter Server from hardware and software failures by performing automatic failover.
Maintenance	No	Yes	Yes	Used for some maintenance tasks. For other tasks, you have to deactivate vCenter HA.

Mode	Automatic Failover	Manual Failover	Replication	
Disabled	No	No	No	If the Passive or Witness nodes are lost or recovering from a failure, a vCenter HA configuration can be deactivated. The Active node continues as a standalone vCenter Server.

Note: If the cluster is operating in either Maintenance or Disabled mode, an Active node can continue serving client requests even if the Passive and Witness nodes are lost or unreachable.

1. Log in to the Active node vCenter Server with the vSphere Client and click **Configure**.
2. Under **Settings** select **vCenter HA** and click **Edit**.
3. Select one of the options.

Option	Result
Enable vCenter HA	Activates replication between the Active and Passive nodes. If the cluster is in a healthy state, your Active node is protected by automatic failover from the Passive node.
Maintenance Mode	In maintenance mode, replication still occurs between the Active and Passive nodes. However, automatic failover is deactivated.
Disable vCenter HA	Deactivates replication and failover. Keeps the configuration of the cluster. You can later activate vCenter HA again.
Remove vCenter HA cluster	Removes the cluster. Replication and failover no longer are provided. The Active node continues to operate as a standalone vCenter Server. See Remove a vCenter HA Configuration for details.

4. Click OK.

Perform Backup and Restore Operations

For additional security, you can back up the Active node in the vCenter HA cluster. You can then restore the node in case of catastrophic failure.

Verify the interoperability of vCenter HA and the backup and restore solution. One solution is vCenter Server file-based restore.

Note: Remove the cluster configuration before you restore the Active node. Results are unpredictable if you restore the Active node and the Passive node is still running or other cluster configuration is still in place.

1. Back up the Active node.
Do not back up the Passive node and Witness node.
2. Before you restore the cluster, power off and delete all vCenter HA nodes.
3. Restore the Active node.
The Active node is restored as a standalone vCenter Server.

Remove a vCenter HA Configuration

You can remove a vCenter HA configuration from the vSphere Client.

1. Log in to the Active node vCenter Server and click **Configure**.
2. Under **Settings** select **vCenter HA** and click **Remove VCHA**.
 - The vCenter HA cluster's configuration is removed from the Active, Passive, and Witness nodes.
 - You can choose to delete the Passive and Witness nodes.
 - The Active node continues to run as a standalone vCenter Server.
 - You cannot reuse the Passive and Witness nodes in a new vCenter HA configuration.
 - If you performed a manual configuration, or if the Passive and Witness nodes are not discoverable, you must delete these nodes explicitly.
 - Even if the second virtual NIC was added by the configuration process, the removal process does not remove the virtual NIC.

Reboot All vCenter HA Nodes

If you have to shut down and reboot all nodes in the cluster, you must follow a specific shutdown order to prevent the Passive node from assuming the role of Active node.

1. Shut down the nodes in this order.
 - Passive node
 - Active node
 - Witness node
2. Restart each node.

You can restart nodes in any order.
3. Verify that all nodes join the cluster successfully, and that the previous Active node resumes that role.

Change the Server Environment

When you deploy a vCenter Server, you select an environment. For vCenter HA, Small, Medium, Large, and X-Large are supported for production environments. If you need more space and want to change the environment, you have to delete the Passive node virtual machine before you change the configuration.

1. Log in to the Active node with the vSphere Client, edit the cluster configuration, and select **Disable**.
2. Delete the Passive node virtual machine.
3. Change the vCenter Server configuration for the Active node, for example, from a Small environment to a Medium environment.
4. Reconfigure vCenter HA.

Collecting Support Bundles for a vCenter HA Node

Collecting a support bundle from all the nodes in a vCenter HA cluster helps with troubleshooting.

When you collect a support bundle from the Active node in a vCenter HA cluster, the system proceeds as follows.

- Collects support bundle information from the Active node itself.
- Collects support bundles from Passive and Witness nodes and places them in the `commands` directory on the Active node support bundle.

Note: The collection of support bundles from the Passive and Witness nodes is a best effort and happens if the nodes are reachable.

Troubleshoot Your vCenter HA Environment

In case of problems you can troubleshoot your environment. The task you need to perform depends on the failure symptoms. For additional troubleshooting information, see the VMware Knowledge Base system.

vCenter HA Clone Operation Fails During Deployment

If the vCenter HA configuration process does not create the clones successfully, you have to resolve that cloning error. Clone operation fails.

Note: Cloning a Passive or Witness VM for a VCHA deployment to the same NFS 3.1 datastore as the source Active node VM fails. You must use NFS4 or clone the Passive and Witness VMs to a datastore different from the Active VM.

Look for the clone exception. It might indicate one of the following problems.

- You have a DRS-enabled cluster, but do not have three hosts.
- The host or database connection is lost.
- Not enough disk space.
- Other **Clone Virtual Machine** errors

1. Resolve the error that caused the problem.
2. Remove the cluster and start configuration again.

Redeploy the Passive or Witness node

If the passive or witness node fails and vCenter HA cluster was configured using the automatic cloning method, you can redeploy it in the **vCenter HA Settings** page.

1. Log in to the Active node with the vSphere Client.
2. Select the vCenter Server object in the inventory and select the **Configure** tab.
3. Select **vCenter HA** under **Settings**.
4. Click on the **REDEPLOY** button next to the node to start the Redeploy wizard.
5.
 - If your vCenter server is managed by another vCenter server in the same SSO domain, proceed to step 6.
 - If your vCenter server is managed by another vCenter server in a different SSO domain, input the location and credential details of that management vCenter server. Enter the **Management vCenter Server FQDN or IP address** and **Single Sign-On** credentials.
6. Specify a unique name and target location.
7. Select the destination compute resource for the operation.
8. Select the datastore in which to store the configuration and disk files.
9. Configure the virtual machine networks.
 - If you are redeploying the Passive node, select virtual machine Management (NIC 0) and vCenter HA (NIC 1) networks.
 - If you are redeploying the Witness node, select vCenter HA (NIC 1) network.

If there are issues with your selections, errors or compatibility warnings are displayed.

10. Review your selections and click **Finish** to redeploy the node.

vCenter HA Deployment Fails with an Error

Deployment failures can be caused by configuration issues, especially problems with the networking setup.

You start a vCenter HA cluster configuration, and it fails with an error. The error might show the cause of the problem, for example, you might see an SSH Connection Failed message.

If deployment fails, take steps to resolve the network issues.

1. Verify that the Passive and Witness nodes can be reached from the Active node.
2. Verify that routing between the nodes is set up correctly.
3. Check network latency.

Troubleshooting a Degraded vCenter HA Cluster

For a vCenter HA cluster to be healthy, each of the Active, Passive, and Witness nodes must be fully operational and be reachable over the vCenter HA cluster network. If any of the nodes fails, the cluster is considered to be in a degraded state.

If the cluster is in a degraded state, failover cannot occur. For information about failure scenarios while the cluster is in a degraded state, see [Failover Failures](#).

The cluster can be in a degraded state for a number of reasons.

One of the nodes fails

- If the Active node fails, a failover of the Active node to the Passive node occurs automatically. After the failover, the Passive node becomes the Active node.
At this point, the cluster is in a degraded state because the original Active node is unavailable.
After the failed node is repaired or comes online, it becomes the new Passive node and the cluster returns to a healthy state after the Active and Passive nodes synchronize.
- If the Passive node fails, the Active node continues to function, but no failover is possible and the cluster is in a degraded state.
If the Passive node is repaired or comes online, it automatically rejoins the cluster and the cluster state is healthy after the Active and Passive nodes synchronize.
- If the Witness node fails, the Active node continues to function and replication between Active and Passive node continues, but no failover can occur.
If the Witness node is repaired or comes online, it automatically rejoins the cluster and the cluster state is healthy.

Database replication fails

If replication fails between the Active and Passive nodes, the cluster is considered degraded. The Active node continues to synchronize with the Passive node. If it succeeds, the cluster returns to a healthy state. This state can result from network bandwidth problems or other resource shortages.

Configuration file replication issues

If configuration files are not properly replicated between the Active and Passive nodes, the cluster is in a degraded state. The Active node continues to attempt synchronization with the Passive node. This state can result from network bandwidth problems or other resource shortages.

How you recover depends on the cause of the degraded cluster state. If the cluster is in a degraded state, events, alarms, and SNMP traps show errors.

If one of the nodes is down, check for hardware failure or network isolation. Check whether the failed node is powered on.

In case of replication failures, check if the vCenter HA network has sufficient bandwidth and ensure network latency is 10 ms or less.

Recovering from Isolated vCenter HA Nodes

If all nodes in a vCenter HA cluster cannot communicate with each other, the Active node stops serving client requests.

Node isolation is a network connectivity problem.

1. Attempt to resolve the connectivity problem. If you can restore connectivity, isolated nodes rejoin the cluster automatically and the Active node starts serving client requests.
2. If you cannot resolve the connectivity problem, you have to log in to Active node's console directly.
 - a) Power off and delete the Passive node and the Witness node virtual machines.
 - b) Log in to the Active node by using SSH or through the Virtual Machine Console.
 - c) To enable the Bash shell, enter `shell` at the `appliance$` prompt.
 - d) Run the following command to remove the vCenter HA configuration.


```
vcha-destroy -f
```
 - e) Reboot the Active node.

The Active node is now a standalone vCenter Server.
 - f) Perform vCenter HA cluster configuration again.

Resolving Failover Failures

When a Passive node does not become the Active node during a failover, you can force the Passive node to become the Active node.

The Passive node fails while trying to assume the role of the Active node.

A vCenter HA failover might not succeed for these reasons.

- The Witness node becomes unavailable while the Passive node is trying to assume the role of the Active node.
- An server state synchronization issue between the nodes exists.

You recover from this issue as follows.

1. If the Active node recovers from the failure, it becomes the Active node again.
2. If the Witness node recovers from the failure, follow these steps.
 - a. Log in to the Passive node through the Virtual Machine Console.
 - b. To enable the Bash shell, enter `shell` at the `appliance$` prompt.
 - c. Run the following command.


```
vcha-reset-primary
```
 - d. Reboot the Passive node.

3. If both Active node and Witness node cannot recover, you can force the Passive node to become a standalone vCenter Server.
 - a. Delete the Active node and Witness node virtual machines.
 - b. Log in to the Passive node through the Virtual Machine Console.
 - c. To enable the Bash shell, enter `shell` at the `appliance$` prompt.
 - d. Run the following command.


```
vcha-destroy
```
 - e. Reboot the Passive node.

VMware vCenter® HA Alarms and Events

If a vCenter HA cluster is in a degraded state, alarms and events show errors.

Table 380: The following events will raise VCHA health alarm in vpxd:

Event Name	Event Description	Event Type	Category
vCenter HA cluster state is currently healthy	vCenter HA cluster state is currently healthy	com.vmware.vcha.cluster.state.health	info
vCenter HA cluster state is currently degraded	vCenter HA cluster state is currently degraded	com.vmware.vcha.cluster.state.degraded	degraded
vCenter HA cluster state is currently isolated	vCenter HA cluster state is currently isolated	com.vmware.vcha.cluster.state.isolated	error
vCenter HA cluster is destroyed	vCenter HA cluster is destroyed	com.vmware.vcha.cluster.state.destroyed	destroyed

Table 381: The following events will raise PSC HA health alarm in vpxd:

Event Name	Event Description	Event Type	Category
PSC HA state is currently healthy	PSC HA state is currently healthy	com.vmware.vcha.psc.ha.health.healthy	healthy
PSC HA state is currently degraded	PSC HA state is currently degraded	com.vmware.vcha.psc.ha.health.degraded	degraded
PSC HA is not monitored after vCenter HA cluster is destroyed	PSC HA state is not being monitored	com.vmware.vcha.psc.ha.health.unknown	unknown

Table 382: Cluster Status Related Events

Event Name	Event Description	Event Type	Category
Node {nodeName} joined back to the cluster	One node joined back to the cluster	com.vmware.vcha.node.joined	info
Node {nodeName} left the cluster	One node left the cluster	com.vmware.vcha.node.left	warning
Failover succeeded	Failover succeeded	com.vmware.vcha.failover.succeeded	info
Failover cannot proceed when cluster is in disabled mode	Failover cannot proceed when cluster is in disabled mode	com.vmware.vcha.failover.failed.disabledmode	disabledmode

Event Name	Event Description	Event Type	Category
Failover cannot proceed when cluster does not have all three nodes connected	Failover cannot proceed when cluster does not have all three nodes connected	com.vmware.vcha.failover.failed.nodes	monitoring
Failover cannot proceed when vPostgres on Passive node is not ready to takeover	Failover cannot proceed when Passive node is not ready to takeover	com.vmware.vcha.failover.failed.passive	passive not ready
vCenter HA cluster mode changed to {clusterMode}	vCenter HA cluster mode changed	com.vmware.vcha.cluster.mode.changed	changed

Table 383: Database replication-related events

Event Name	Event Description	Event Type	Category
Database replication mode changed to {newState}	Database replication state changed: sync, async or no replication	com.vmware.vcha.DB.replication.state	state changed

Table 384: File replication-related events

Event Name	Event Description	Event Type	Category
Appliance {fileProviderType} is {state}	Appliance File replication state changed	com.vmware.vcha.file.replication.state	state changed

Patching a vCenter High Availability Environment

You can patch a vCenter Server which is in a vCenter High Availability cluster by using the `software-packages` utility available in the vCenter Server shell.

For more information, see *Patch a vCenter High Availability Environment* in *vSphere Upgrade*.

Reduced Downtime Upgrade for vCenter HA

In vSphere 8.0 U3, Reduced Downtime Upgrade is integrated with automatic vCenter HA deployment.

Reduced Downtime Upgrade (RDU) is migration-based VCSA upgrade with the primary goal of reducing downtime for upgrades. During RDU upgrade VCSA configuration, VCDB database and the network information is copied from the old VCSA to the new VCSA version before shutting down the source VCSA and switching over to the target VCSA. During migration upgrade, in the Staging phase before RDU switchover, VCHA is automatically undeployed. There is no VCHA for the duration of the upgrade. After successful RDU switchover, VCHA is redeployed on the target node.

RDU upgrade is integrated with automatic vCenter HA deployment, including self-managed vCenter and non-self-managed vCenter. You can upgrade vCenter without vCenter credentials if the vCenter is self-managed. You must use service account credentials for the managing vCenter provided by the RDU framework when upgrading a non-self-managed vCenter. You can upgrade vCenter appliance that includes VCHA, without remove or setup vCenter HA before and after the upgrade. If you cancel the upgrade, RDU rollback results in a functional vCenter HA deployment as it existed prior to the upgrade.

vSphere Monitoring and Performance

VMware provides several tools to help you monitor your virtual environment and to locate the source of potential issues and current problems.

Performance Charts

Allow you to see performance data on a variety of system resources including CPU, Memory, Storage, and so on.

Performance Monitoring Command-line Utilities

Allow you to access detailed information on system performance through the command line.

Host Health

Allows you to quickly identify which hosts are healthy and which are experiencing problems.

Events, Alerts, and Alarms

Allow you to configure alerts and alarms and to specify the actions the system should take when they are triggered.

System Log Files

System logs contain additional information about activities in your vSphere environment.

Intended Audience

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

The content in this section is intended for vSphere administrators who perform the following tasks:

- Monitor the health and performance of physical hardware backings for the virtual environment.
- Monitor the health and performance of virtual devices in the virtual environment.
- Troubleshoot problems in the system.
- Configure alarms.
- Configure SNMP messages.
- Use vCenter Events for forensic analysis and auditing of the user actions in the virtual environment.

Virtual machine administrators also might find the section on [Monitoring Guest Operating System Performance](#) helpful.

Monitoring Inventory Objects with Performance Charts

The vSphere statistics subsystem collects data on the resource usage of inventory objects. Data on a wide range of metrics is collected at frequent intervals, processed, and archived in the vCenter Server database. You can access statistical information through command-line monitoring utilities or by viewing performance charts in the vSphere Client.

Counters and Metric Groups

vCenter Server systems and hosts use data counters to query for statistics. A data counter is a unit of information relevant to a given inventory object or device. Each counter collects data for a different statistic in a metric group. For example, the disk metric group includes separate data counters to collect data for disk read rate, disk write rate, and disk usage. Statistics for each counter are rolled up after a specified collection interval. Each data counter consists of several attributes that are used to determine the statistical value collected.

For a complete list and description of performance metrics, see the *vSphere API Reference*.

Note: Counters that are introduced in later versions might not contain data from hosts of earlier versions. For details, see the VMware Knowledge Base.

Collection Levels and Collection Intervals

Collection levels determine the number of counters for which data is gathered during each collection interval. Collection intervals determine the time period during which statistics are aggregated, calculated, rolled up, and archived in the vCenter Server database. Together, the collection interval and collection level determine how much statistical data is collected and stored in your vCenter Server database.

Data Availability

Real-time data appears in the performance charts only for hosts and virtual machines that are powered on. Historical data appears for all supported inventory objects, but might be unavailable during certain circumstances.

Performance Chart Types

Performance metrics are displayed in different types of charts, depending on the metric type and object.

Table 385: Performance Chart Types

Chart Type	Description
Line chart	Displays metrics for a single inventory object. The data for each performance counter is plotted on a separate line in the chart. For example, a network chart for a host can contain two lines: one showing the number of packets received, and one showing the number of packets transmitted.
Bar chart	Displays storage metrics for datastores in a selected data center. Each datastore is represented as a bar in the chart. Each bar displays metrics based on the file type: virtual disks, snapshots, swap files, and other files.
Pie chart	Displays storage metrics for a single object, based on the file types, or virtual machines. For example, a pie chart for a datastore can display the amount of storage space occupied by the virtual machines taking up the largest space.
Stacked chart	Displays metrics for the child objects that have the highest statistical values. All other objects are aggregated, and the sum value is displayed with the term Other . For example, a host's stacked CPU usage chart displays CPU usage metrics for the 10 virtual machines on the host that are consuming the most CPU. The Other amount contains the total CPU usage of the remaining virtual machines. The metrics for the host itself are displayed in separate line charts. Stacked charts are useful in comparing the resource allocation and usage across multiple hosts or virtual machines. By default, the 10 child objects with the highest data counter values are displayed.

Data Counters

Each data counter includes several attributes that are used to determine the statistical value collected. See the *vSphere API Reference* for a complete list and description of supported counters.

Table 386: Data Counter Attributes

Attribute	Description
Unit of Measurement	Standard in which the statistic quantity is measured. - Kilobytes (KB) – 1024 bytes Note: Technically, 1 Kilobytes (KB) = 1000 bytes and 1 Kibibyte (KiB) = 1024 bytes. However, based on the context, kilobyte is also used for 1024 bytes in the computer science literature interchangeably. - Kilobytes per second (KBps) – 1024 bytes per second - Kilobits (kb) – 1000 bits - Kilobits per second (kbps) – 1000 bits per second - Megabytes (MB) - Megabytes per second (MBps)

Attribute	Description
	• Megabits (Mb), megabits per second (Mbps) • Megahertz (MHz) • Microseconds (µs) • Milliseconds (ms) • Number (#) • Percent (%) • Seconds (s) • Watts (watt) • Joules (joule) • Terabytes (teraBytes) • Temperature in celsius (celsius)
Description	Text description of the data counter.
Statistics Type	Measurement used during the statistics interval. Related to the unit of measurement. • Rate – Value over the current statistics interval. For example: CPU Usage: Amount of CPU actively used by the host, resource pool, or virtual machine in the cluster. Counter: usagemhz Stats Type: Rate Unit: MegaHertz (MHz) • Delta – Change from previous statistics interval. For example: CPU System Time for vCPU (%): Amount of time spent on system processes on each virtual CPU in the virtual machine. Note: This is the host view of the CPU usage, not the guest operating system view. Counter: system Stats Type: Delta Unit: Percentage (%) • Absolute – Absolute value (independent of the statistics interval). For example: Memory: Amount of host machine memory used by all powered on virtual machines in the cluster. A cluster's consumed memory consists of virtual machine consumed memory and overhead memory. It does not include host-specific overhead memory, such as memory used by the service console or VMkernel. Counter: consumed Stats Type: Absolute Unit: MegaBytes (MB)
Rollup Type	Calculation method used during the statistics interval to aggregate data. Determines the type of statistical values that are returned for the counter. • Average – Data collected during the interval is aggregated and averaged. For example: Virtual Disk Read Requests: Number of virtual disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk read commands is also displayed in the chart. Counter: numberRead Stats Type: Absolute Unit: Number Rollup Type: Average • Minimum – The minimum value is rolled up.

Attribute	Description
	• Maximum – The maximum value is rolled up. The Minimum and Maximum values are collected and displayed only in statistics level 4. Minimum and maximum rollup types are used to capture peaks in data during the interval. For real-time data, the value is the current minimum or current maximum. For historical data, the value is the maximum or minimum of the aggregated values. For example, the following information for the CPU usage chart shows that the average is collected at statistics level 1. The minimum and maximum values are collected at statistics level 4. – Counter: usage – Unit: Percentage (%) – Rollup Type: Average (Minimum/Maximum) – Collection Level: 1 (4) • Summation – Data collected is summed. The measurement displayed in the chart represents the sum of data collected during the interval. For example: Network Packets Received: Number of network packets received across the top ten physical NIC instances on the host. The chart also displays the aggregated value for all NICs. Counter: packetRx Stats Type: Absolute Unit: Number Rollup Type: Summation • Latest – Data collected during the interval is a set value. The value displayed in the performance charts represents the current value. For example: Space in GB (Allocated): Total amount of logical datastore space provisioned by an administrator for the virtual machine. It is the storage size up to which the virtual machine files on datastores can grow. This includes log files, VMX files, and other miscellaneous files. Allocated space is not always in use. Counter: provisioned Stats Type: Absolute Unit: Gigabytes (GB) Rollup Type: Latest
Collection level	A collection level determines the number of data counters used to collect statistics data. Collection levels are also referred as statistics levels. These collection levels vary from 1 to 4, with Level 4 having the most counters. • Level 1 is the least detailed statistics level and only includes the most critical statistics, such as aggregate CPU, memory, and network usage. • Level 2 introduces a number of additional statistics • Level 3 incorporates per instance statistics, for example CPU usage of a host on a per-CPU basis. • Level 4 is the most detailed and is inclusive of all the other levels For more information on collection levels, see Data Collection Levels . Note: Be careful when you set a higher collection level, as the process requires significant increase of resource usage.

Metric Groups in vSphere

The performance data collection subsystem for vSphere collects performance data on various inventory items and their devices. Data counters define individual performance metrics. Performance metrics are organized into logical groups based on the object or object device. Statistics for one or more metrics can be displayed in a chart.

Table 387: Metric Groups

Metric group	Description
Cluster Services	Performance statistics for clusters configured by using vSphere Distributed Resource Scheduler, vSphere High Availability, or both.
CPU	CPU utilization per host, virtual machine, resource pool, or compute resource.
Datastore	Statistics for datastore utilization. Note: From VC 4.1 onwards, NFS statistics are collected under datastore statistics.
Disk	Disk utilization per host, virtual machine, or datastore. Disk metrics include I/O performance, such as latency and read/write speeds, and utilization metrics for storage as a finite resource.
GPU	GPU utilization per cluster, or host. GPU metrics include GPU utilization in percentage of the total, memory usage, temperature of the GPU devices.
Memory	Memory utilization per host, virtual machine, resource pool, or compute resource. The value obtained is one of the following: - For virtual machines, memory refers to the guest physical memory. Guest physical memory is the amount of physical memory presented as a virtual-hardware component to the virtual machine, at creation time, and made available when the virtual machine is running. - For hosts, memory refers to the machine memory. Machine memory is the RAM that is installed on the hardware that comprises the host.
Network	Network utilization for both physical and virtual network interface controllers (NICs) and other network devices. The virtual switches that support connectivity among all components, such as hosts, virtual machines, VMkernel.
Power	Energy usage statistics per host.
Storage Adapter	Data traffic statistics per host bus adapter (HBA).
Storage Path	Data traffic statistics per path.
System	Overall system availability, such as the system heartbeat and uptime. These counters are available directly from hosts and from vCenter Server.
Virtual Disk	Disk utilization and disk performance metrics for virtual machines.
Virtual Flash	Virtual flash counters.
Virtual Machine Operations	Virtual machine power and provisioning operations in a cluster or data center.
vSphere Replication	Statistics for the virtual machine replication performed by VMware vCenter Site Recovery Manager.

Data Collection Intervals

Collection intervals determine the duration for which statistics are aggregated, calculated, rolled up, and archived. Together, the collection interval and collection level determine how much statistical data is gathered and stored in your vCenter Server database.

Table 388: Collection Intervals

Collection Interval/ Archive Length	Collection Frequency	Default Behavior
1 Day	5 Minutes	Real-time (20s) statistics are rolled up to create one data point every 5 minutes. The result is 12 data points every hour and 288 data points every day. After 30 minutes, the six data points collected are aggregated and rolled up as a data point for the 1-Week time range. You can change the interval duration and archive length of the 1-Day collection interval by configuring the statistics settings.
1 Week	30 Minutes	1-Day statistics are rolled up to create one data point every 30 minutes. The result is 48 data points every day and 336 data points every week. Every 2 hours, the 12 data points collected are aggregated and rolled up as a data point for the 1-Month time range. You cannot change the default settings of the 1-Week collection interval.
1 Month	2 Hours	1-Week statistics are rolled up to create one data point every 2 hours. The result is 12 data points every day and 360 data points every month (assuming a 30-day month). After 24 hours, the 12 data points collected are aggregated and rolled up as a data point for the 1-Year time range. You cannot change the default settings of the 1-Month collection interval.
1 Year	1 Day	1-Month statistics are rolled up to create one data point every day. The result is 365 data points each year. You can change the archive length of the 1-Year collection interval by configuring the statistics settings.

Note: If you change the duration of data collection intervals you might need to allocate more storage resources.

Data Collection Levels

Each collection interval has a default collection level that determines the amount of data gathered and which counters are available for display in the charts. Collection levels are also referred to as statistics levels.

Table 389: Statistics Levels

Level	Metrics	Best Practice
Level 1	- Cluster Services (VMware Distributed Resource Scheduler) – all metrics - CPU – cpuentitlement, totalmhz, usage (average), usagemhz - Disk – capacity, maxTotalLatency, provisioned, unshared, usage (average), used - Memory – consumed, mementitlement, overhead, swapinRate, swapoutRate, swapused, totalmb, usage (average), vmmemctl (balloon), totalbandwidth (DRAM or PMem) - Network – usage (average), IPv6 - System – heartbeat, uptime - Virtual Machine Operations – numChangeDS, numChangeHost, numChangeHostDS	Use for long-term performance monitoring when device statistics are not required. Level 1 is the default Collection Level for all Collection Intervals.
Level 2	- Level 1 metrics - CPU – idle, reservedCapacity - Disk – All metrics, excluding numberRead and numberWrite. - Memory – All metrics, excluding memUsed, maximum and minimum rollup values, read or write latency (DRAM or PMem).	Use for long-term performance monitoring when device statistics are not required but you want to monitor more than the basic statistics.

Level	Metrics	Best Practice
	Virtual Machine Operations – All metrics	
Level 3	- Level 1 and Level 2 metrics - Metrics for all counters, excluding minimum and maximum rollover values. - Device metrics	Use for short-term performance monitoring after encountering problems or when device statistics are required.
Level 4	All metrics supported by the vCenter Server, including minimum and maximum rollover values.	Use for short-term performance monitoring after encountering problems or when device statistics are required.

Note:

When the statistics levels, level 3 or level 4 are used beyond the default value, it may cause one particular process, vpxd, to sustain memory growth, if it cannot save the statistics information to the database as quickly as required. If the usage limit of these statistics levels is not monitored closely, it may cause vpxd to grow out of memory and eventually crash.

So, in case the administrator decides to elevate any of these levels, it is necessary for the administrator to monitor the size of the vpxd process to make sure that is not growing boundlessly after the change.

View Performance Charts

The vCenter Server statistics settings, the type of object selected, and the features that are enabled on the selected object determine the amount of information displayed in charts. Charts are organized into views. You can select a view to see related data together on one screen. You can also specify the time range, or data collection interval. The duration extends from the selected time range to the present time.

Overview charts display multiple data sets in one panel to evaluate different resource statistics, display thumbnail charts for child objects. It also displays charts for a parent and a child object. Advanced charts display more information than overview charts, are configurable, and can be printed or exported. You can export data in the PNG, JPEG, or CSV formats. See [View Advanced Performance Charts](#).

1. Select a valid inventory object in the vSphere Client.

Overview and advanced performance charts are available for datacenter, cluster, host, resource pool, vApp, and virtual machine objects. Overview charts are also available for datastores and datastore clusters. Performance charts are not available for network objects.

2. Click the **Monitor** tab, and click **Performance**.

3. Select a view.

Available views depend on the type of object. For views that might contain many charts in a large environment, the vSphere Client displays the charts distributed on multiple pages. You can use the arrow buttons to navigate between pages.

4. Select a predefined or custom time range.

Performance Charts Options Available Under the View Menu

The performance chart options that you can access under the **View** menu vary depending on the type of inventory object you select.

For example, the **Virtual Machines** view is available when you view host performance charts only if there are virtual machines on the selected host. Likewise, the **Fault Tolerance** view for virtual machine performance charts is available only when that feature is enabled for the selected virtual machine.

Table 390: Performance Chart Views by Inventory Object

Object	View List Items
Data center	• Storage - space utilization charts for datastores in the data center, including space by file type and storage space used by each datastore in the data center. • Clusters - thumbnail CPU and memory charts for each cluster, and stacked charts for total CPU and memory usage in the data center. This view is the default.
Datastore and datastore cluster	• Space - space utilization charts for the datastore: — space utilization by file type — space utilization by virtual machine — space usage • Performance - performance charts for the datastore or datastore cluster and for virtual machine disks on the resource. Note: The Performance view for datastores is only available when all hosts that are connected to the datastores are ESX/ESXi 4.1 or greater. The Performance view for datastore clusters is only available when the Storage DRS is enabled.
Cluster	• Home - CPU and memory charts for the cluster. • Resource Pools & Virtual Machines - thumbnail charts for resource pools and virtual machines, and stacked charts for total CPU and memory usage in the cluster. • Hosts - thumbnail charts for each host in the cluster, and stacked charts for total CPU, memory, disk usage, and network usage.
Host	• Home - CPU, memory, disk, GPU compute, GPU memory and network charts for the host. • Virtual Machines - thumbnail charts for virtual machines, and stacked charts for total CPU usage and total memory usage on the host. • Disks - read rates, write rates, and disk usage for the host. • Network Adapters - network bandwidth on a host. • GPUs - thumbnail charts for GPUs compute usage, memory allocation, and temperature.
Resource Pool and vApps	• Home - CPU and memory charts for the resource pool. • Resource Pools & Virtual Machines - thumbnail charts for resource pools, and virtual machines and stacked charts for CPU and memory usage in the resource pool or vApp.
Virtual Machine	• Storage - space utilization charts for the virtual machine: space by file type, space by datastore, and total gigabytes. • Fault Tolerance - CPU and memory charts that display comparative metrics for the fault-tolerant primary and secondary virtual machines. • Home - CPU, memory, network, host (thumbnail charts), and disk usage charts for the virtual machine.

Overview Performance Charts

The overview performance charts display the metrics for an object in the inventory. Use these charts to monitor and troubleshoot performance problems.

The metrics provided in Overview performance charts are a subset of those collected for hosts and the vCenter Server. For a complete list of all metrics collected by hosts and the vCenter Server, see the *vSphere API Reference*.

Clusters

The cluster charts contain information about CPU, disk, memory, and network usage for clusters. The help topic for each chart contains information about the data counters displayed in that chart. The collection level set for vCenter Server determines the available counters.

CPU (MHz)

The CPU (MHz) chart displays CPU usage for the cluster.

Cluster Counters

This chart is located in the Home view of the Cluster **Performance** tab.

Table 391: Data Counters

Chart Label	Description
Usage	Sum of the average CPU usage values, in Megahertz, of all virtual machines in the cluster. - Counter: usagemhz - Stats Type: Rate - Unit: Megahertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Total	Total amount of CPU resources available in the cluster. The maximum value is equal to the number of cores multiplied by the frequency of the processors. For example, a cluster has two hosts, each of which has four CPUs that are 3 GHz each, and one virtual machine that has two virtual CPUs. VM totalmhz = 2 vCPUs * 3000 MHz = 6000 MHz Host totalmhz = 4 CPUs * 3000 MHz = 12000 MHz Cluster totalmhz = 2 x 4 * 3000 MHz = 24000 MHz - Counter: totalmhz - Stats Type: Rate - Unit: Megahertz (MHz) - Rollup Type: Summation - Collection Level: 1

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

Table 392: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: 1. Select the cluster, and click the Configure tab. 2. Under Services, click vSphere DRS. 3. click Edit. An Edit Cluster Settings dialog box opens. 4. Click Turn ON vSphere DRS, and click OK.
3	If the cluster is a DRS cluster: • Increase the number of hosts, and migrate one or more virtual machines to the new host.

#	Resolution
	Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster.
4	Migrate one or more virtual machines to a new host.
5	Upgrade the physical CPUs or cores on each host in the cluster if necessary.
6	Enable CPU-saving features, such as TCP Segmentation Offload.
7	Replace software I/O with the dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

CPU Usage

The cluster CPU Usage charts monitor the CPU utilization of the hosts, resource pools, and virtual machines in the cluster. This chart displays the 10 child objects in the cluster with the most CPU usage.

This chart is located in the Resource Pools and Virtual Machines view of the Cluster **Performance** tab.

Table 393: Data Counters

Chart Label	Description
<host>, <resource pool>, or <virtual machine>	Amount of CPU actively used by the host, resource pool, or virtual machine in the cluster. - Counter: usagemhz - Stats Type: Rate - Unit: MegaHertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

Table 394: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: - Select the cluster, and click the Configure tab. - Under Services, click vSphere DRS. - click Edit. An Edit Cluster Settings dialog box opens. - Click Turn ON vSphere DRS, and click OK.
3	If the cluster is a DRS cluster: - Increase the number of hosts, and migrate one or more virtual machines to the new host. - Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster.
4	Migrate one or more virtual machines to a new host.

#	Resolution
5	Upgrade the physical CPUs or cores on each host in the cluster if necessary.
6	Enable CPU-saving features, such as TCP Segmentation Offload.
7	Replace software I/O with the dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

Disk (KBps)

The Disk (KBps) chart displays the disk I/O of the 10 hosts in the cluster with the most disk usage.

This chart is located in the Hosts view of the cluster **Performance** tab.

Table 395: Data Counters

Chart Label	Description
<i>host_name</i>	Average data I/O rate across all hosts in the cluster. - Counter: usage - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 396: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access.

#	Resolution
	To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Memory (MB)

The Memory (MB) chart displays consumed memory for the cluster. The chart appears only at collection level 1.

This chart is located in the Home view of the cluster **Performance** tab.

Table 397: Data Counters

Chart Label	Description
Consumed	Amount of host machine memory used by all powered on virtual machines in the cluster. A cluster's consumed memory consists of virtual machine consumed memory and overhead memory. It does not include host-specific overhead memory, such as memory used by the service console or VMkernel. - Counter: consumed - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Total	Total amount of machine memory of all hosts in the cluster that is available for virtual machine memory (physical memory for use by the Guest OS) and virtual machine overhead memory. Memory Total = Aggregate host machine memory - (VMkernel memory + Service Console memory + other service memory)

Chart Label	Description
	Note: The totalmb data counter is the same as the effectivemem data counter, which is supported only for backward compatibility. - Counter: totalmb - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Memory usage is not necessarily an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

Table 398: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.
4	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: 1. Select the cluster, and click the Configure tab. 2. Under Services, click vSphere DRS. 3. click Edit. An Edit Cluster Settings dialog box opens. 4. Click Turn ON vSphere DRS, and click OK.
5	If the cluster is a DRS cluster: • Increase the number of hosts, and migrate one or more virtual machines to the new host. • Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.
6	Add more physical memory to one or more hosts.

Memory (MB)

The Memory (MB) chart displays memory data counters for clusters. The chart appears at all collection levels except level 1.

Description

This chart is located in the **Home** view of the cluster **Performance** tab.

Note: These data counter definitions are for hosts. At the cluster level, the values are collected and totaled. The counter values in the chart represent the aggregate amounts of the host data. The counters that appear in the chart depend on the collection level set for your vCenter Server.

Table 399: Data Counters

Chart Label	Description
Active	Sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMkernel applications. Active memory is estimated by the VMkernel. - Counter: active - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Balloon	Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host. - Counter: vmmemctl - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Consumed	Amount of machine memory used on the host. Consumed memory includes virtual machine memory, service console memory, and VMkernel memory. consumed memory = total host memory - free host memory - Counter: consumed - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Granted	Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory. Granted memory for a host includes the shared memory of each virtual machine on the host. - Counter: granted - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Swap Used	Sum of the memory swapped by all powered on virtual machines on the host. - Counter: swapused - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Total	Aggregate total memory available to the cluster. - Counter: totalmb - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

Table 400: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory Consumed

The Memory Consumed chart displays memory usage for the 10 child objects in the cluster with the most consumed memory.

For resource pools and virtual machines in a cluster, this chart is located in the **Resource Pools & Virtual Machines** view of the cluster **Performance** tab. For hosts in a cluster, this chart is located in the **Hosts** view of the cluster **Performance** tab.

Table 401: Data Counters

Chart Label	Description
<i>resource_pool, virtual_machine, or host</i>	Amount of machine memory used by all resource pools and virtual machines in the cluster or by all hosts in the cluster, depending on the cluster view. Consumed memory includes virtual machine memory, service console memory, and VMkernel memory. consumed memory = total host memory - free host memory - Counter: consumed - Stats Type: Absolute - Unit: MegaBytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Memory usage is not necessarily an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

Table 402: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.
4	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: 1. Select the cluster, and click the Configure tab. 2. Under Services, click vSphere DRS. 3. click Edit. An Edit Cluster Settings dialog box opens. 4. Click Turn ON vSphere DRS, and click OK.
5	If the cluster is a DRS cluster: • Increase the number of hosts, and migrate one or more virtual machines to the new host. • Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.
6	Add more physical memory to one or more hosts.

Network (Mbps)

The Network (Mbps) chart displays network speed for the 10 hosts in the cluster with the most network usage.

This chart is located in the **Hosts** view of the Cluster **Performance** tab.

Table 403: Data Counters

Chart Label	Description
<host>	Average rate at which data is transmitted and received across all NIC instances on the host. - Counter: usage - Stats Type: Rate - Unit: Megabits per second (Mbps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 404: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.

#	Resolution
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Data centers

The data center charts contain information about CPU, disk, memory, and storage usage for data centers. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

CPU (MHz)

The CPU (MHz) chart displays CPU usage for the 10 clusters in the data center with the most CPU usage.

This chart is located in the Clusters view of the Datacenters **Performance** tab.

Table 405: Data Counters

Chart Label	Description
<cluster>	Amount of CPU currently in use by the cluster. The active CPU usage is approximately equal to the ratio of the used CPU cycles to the available CPU cycles. The maximum possible value is the frequency of the processors multiplied by the number of cores. For example, a two-way SMP virtual machine using 4000MHz on a host that has four 2GHz processors is using 50% of the CPU ($4000 \div 4 \times 2000 = 0.5$). - Counter: usagemhz - Stats Type: Rate - Unit: MegaHertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of cluster resources. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the hosts in the cluster.

If performance is impacted, consider taking the following actions.

Table 406: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: 1. Select the cluster, and click the Configure tab. 2. Under Services, click vSphere DRS. 3. click Edit. An Edit Cluster Settings dialog box opens. 4. Click Turn ON vSphere DRS, and click OK.
3	If the cluster is a DRS cluster:

#	Resolution
	• Increase the number of hosts, and migrate one or more virtual machines to the new host. • Check the aggressiveness threshold. If the value is low, increase the threshold. This might help avoid hot spots in the cluster.
4	Migrate one or more virtual machines to a new host.
5	Upgrade the physical CPUs or cores on each host in the cluster if necessary.
6	Enable CPU-saving features, such as TCP Segmentation Offload.
7	Replace software I/O with the dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

Memory (MB)

The Memory (MB) chart displays the average amount of consumed memory for the 10 clusters in the data center with the most consumed memory.

This chart is located in the **Clusters** view of the Datacenters **Performance** tab.

Table 407: Data Counters

Chart Label	Description
<cluster>	Amount of host machine memory used by all powered on virtual machines in the cluster. • Counter: consumed • Stats Type: Absolute • Unit: MegaBytes (MB) • Rollup Type: Average (Minimum/Maximum) • Collection Level: 1 (4)

Chart Analysis

A cluster's consumed memory consists of virtual machine consumed memory and overhead memory. It does not include host-specific overhead memory, such as memory used by the service console or VMkernel.

If you experience problems with cluster memory usage, use the thumbnail cluster charts to examine memory usage for each cluster and increase memory resources if needed.

If the cluster is a DRS cluster, check the aggressiveness threshold. If the value is low, increase the threshold. Increasing the threshold might help avoid hot spots in the cluster.

Space Utilization By File Type

The Space Utilization By File Type chart displays datastore space usage for virtual disks, swap files, snapshot files, and other virtual machine files.

Note: This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Storage** view of the data center **Performance** tab.

Datastore Counters

Table 408: Data Counters

File Type	Description
Virtual Disks	Amount of disk space used by virtual disk files. Virtual disk files store the contents of the virtual machine's hard disk drive. It includes information that you write to a virtual machine's hard disk, such as the operating system, program files, and data files. The files have the extension <code>.vmdk</code> and appear as a physical disk drive to a guest operating system. Note: Delta disks, which also have an extension <code>.vmdk</code>, are not included in this file type. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Swap Files	Amount of disk space used by swap files. Swap files back up the virtual machine's physical memory. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Snapshots	Amount of disk space used by virtual machine snapshot files. Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension <code>.vmsn</code>. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Other VM Files	Amount of disk space used by all other virtual machine files, such as configuration files and log files. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Other	Amount of disk space used by all other non-virtual machine files, such as documentation files and backup files.
Free Space	Amount of disk space not currently in use.
Total Space	Amount of disk space available to the datastore. It defines the datastore capacity. The chart displays the information for datastores but not for data centers. total space = virtual disk space + swap file space + snapshot space + other VM file space + other space + free space

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Datastores and Datastore Clusters

The datastore charts contain information about disk usage for datastores or the datastores that are part of a cluster. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

Space in GB

The Space in GB chart displays space usage data counters for datastores.

This chart is located in the **Space** view of the datastore or datastore cluster **Performance** tab.

Table 409: Data Counters

Chart Label	Description
Allocated	Amount of physical space provisioned by an administrator for the datastore. It is the storage size up to which files on the datastore can grow. Allocated space is not always in use. - Counter: provisioned - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1
Used	Amount of physical datastore space in use. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1
Capacity	Maximum capacity of the datastore. capacity = virtual machine file space + non-virtual machine file space + free space Note: Storage data is collected and updated in the overview charts every 30 minutes. Therefore, if you refresh the datastore, the capacity value might only be updated in the datastore Summary tab, and not in the overview charts. - Counter: capacity - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Space Utilization By File Type

The Space Utilization by File Type chart displays space used by virtual disks, swap files, snapshot files, and other virtual machine files on the datastore or the datastore cluster.

Note: This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Storage** view of the datastore **Performance** tab. The counters can also be displayed for datastore cluster charts.

Datastore Counters

Table 410: Data Counters

File Type	Description
Virtual Disks	Amount of disk space used by virtual disk files. Virtual disk files store the contents of the virtual machine's hard disk drive. It includes information that you write to a virtual machine's hard disk, such as the operating system, program files, and data files. The files have the extension <code>.vmdk</code> and appear as a physical disk drive to a guest operating system. Note: Delta disks, which also have an extension <code>.vmdk</code>, are not included in this file type. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Swap Files	Amount of disk space used by swap files. Swap files back up the virtual machine's physical memory. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Snapshots	Amount of disk space used by virtual machine snapshot files. Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension <code>.vmsn</code>. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest

File Type	Description
	Collection Level: 1 (4)
Other VM Files	Amount of disk space used by all other virtual machine files, such as configuration files and log files. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Other	Amount of disk space used by all other non-virtual machine files, such as documentation files and backup files.
Free Space	Amount of disk space not currently in use.
Total Space	Amount of disk space available to the datastore. It defines the datastore capacity. The chart displays the information for datastores but not for data centers. total space = virtual disk space + swap file space + snapshot space + other VM file space + other space + free space

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Space used by Datastore in GB

The Space used by Datastore in GB chart displays the 10 datastores in the data center with the most used disk space.

This chart is located in the **Storage** view of the Datacenter **Performance** tab.

Table 411: Data Counters

Chart Label	Description
<datastore>	Amount of used storage space on the 10 datastores with the most used space. - Counter: used - Stats Type: Absolute - Unit: GigaBytes (GB) - Rollup Type: Latest - Collection Level: 1

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Space Utilization by Virtual Machine

The Space Utilization by Virtual Machine chart displays the amount of space used by the five virtual machines with the most space used on the datastore or the datastores in the cluster.

Note: This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by Virtual Machine chart is located in the **Space** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

Table 412: Data Counters

Chart Label	Description
<i>virtual_machine</i>	Amount of datastore space used by the five virtual machines with the most used datastore space. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Space Allocated by Datastore in GB

The Space allocated by Datastore in GB displays the top 10 datastores, virtual machines in the datastore-cluster with most provisioned space.

This chart is located in the **Space** view of the Datacenter **Performance** tab.

Table 413: Data Counters

Chart Label	Description
<datastore>	Amount of provisioned storage space on the top 10 datastores with the most provisioned space. - Counter: provisioned - Stats Type: Absolute - Unit: KiloBytes (KB) - Rollup Type: Latest - Collection Level: 1

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Space Capacity by Datastore in GB

The Space capacity by Datastore in GB displays the top 10 configured size of the datastores in the datastore cluster.

This chart is located in the **Space** view of the Datacenter **Performance** tab.

Table 414: Data Counters

Chart Label	Description
<datastore>	Configured size of the datastores in the datastore cluster. - Counter: capacity - Stats Type: Absolute - Unit: KiloBytes (KB) - Rollup Type: Latest - Collection Level: 1

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Storage I/O Control Normalized Latency

This chart displays the normalized latency in microseconds on the datastore. Storage I/O Control monitors latency to detect congestion on the datastore. This metric computes a weighted response time across all hosts and VMs accessing the datastore. I/O count is used as the weight for the response time. It captures the device level latency and does not include any queuing inside the hypervisor storage stack or inside the VM. It is adjusted for the I/O size. High latencies that are the result of large I/Os are discounted so as not to make the datastore seem slower than it really is. Data for all virtual machines is combined. This chart displays zero values when Storage I/O Control is deactivated.

This chart is located in the **Performance** view of the datastore **Performance** tab. The sizeNormalizedDatastoreLatency counter can also be displayed for datastore cluster charts.

Table 415: Data Counters

Chart Label	Description
Storage I/O Control Normalized Latency	Storage I/O Control monitors latency to detect congestion on the datastore. - Counter: sizeNormalizedDatastoreLatency - Stats Type: Absolute - Unit: Microseconds - Rollup Type: Average - Collection Level: 3

Storage I/O Control Aggregate IOPs

This chart displays the number of I/O operations per one second on the datastore, aggregated across all hosts, and virtual machines accessing this datastore. The chart displays zero values when Storage I/O Control is deactivated.

This chart is located in the **Performance** view of the datastore or the datastore cluster **Performance** tab. The counter can be displayed for datastore and datastore cluster charts.

Table 416: Data Counters

Chart Label	Description
Storage I/O Control Aggregate IOPs	Number of I/O operations per second on the datastore, aggregated across all hosts, and virtual machines accessing the datastore. - Counter: datastoreIops - Stats Type: Absolute - Unit: Number - Rollup Type: Average - Collection Level: 3

Storage I/O Control Activity

This chart displays the percentage of time during which Storage I/O Control actively controlled latency of the datastore.

This chart is located in the **Performance** views of the datastore **Performance** tabs. The counter can also be displayed for datastore cluster charts.

Table 417: Data Counters

Chart Label	Description
Storage I/O Control Activity	This is the percentage of time during which the Storage I/O Control actively controlled the I/O latency for the datastore. - Counter: siocActiveTimePercentage - Stats Type: Absolute - Unit: Percent - Rollup Type: Average - Collection Level: 3

Average Device Latency per Host

This chart displays the average amount of latency on a host device. The chart displays the 10 hosts with the highest device latency.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 418: Data Counters

Chart Label	Description
Average Device Latency per Host	Measures the amount of time, in milliseconds, to complete a SCSI command issued from the physical device. - Counter: deviceLatency - Stats Type: Absolute - Unit: Milliseconds (ms) - Rollup Type: Average - Collection Level: 3

Maximum Queue Depth per Host

This chart displays maximum queue depth that hosts are currently maintaining for the datastore. When Storage I/O is enabled, queue depth can change over time when congestion is detected at the array.

This chart is located in the **Performance** view of the datastore **Performance** tab. The chart displays information about the ten hosts with the highest values.

Table 419: Data Counters

Chart Label	Description
Max Queue Depth per Host	Maximum queue depth. Queue depth is the number of commands the SCSI driver queues to the HBA. - Counter: maxQueueDepth - Stats Type: Absolute - Unit: Number - Rollup Type: Average - Collection Level: 3

Read IOPs per Host

This chart displays the per-host disk read rates for a datastore. The chart displays information about the ten hosts with the highest values.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 420: Data Counters

Chart Label	Description
Read IOPs per Host	Number of disk read commands completed on each disk on the host, per second. Read rate = blocks read per second × block size Counter: numberReadAveraged

Chart Label	Description
	- Stats Type: Rate - Unit: Number - Rollup Type: Average - Collection Level: 3

Write IOPs Per Host

This chart displays the per-host disk write rates for a datastore. The chart displays information about the 10 hosts with the highest values.

This chart is located in the **Performance** view of the datastore **Performance** tab.

Table 421: Data Counters

Chart Label	Description
Write IOPs per Host	Number of disk write commands completed on each disk on the host, per second. Write rate = blocks written per second × block size - Counter: numberWriteAveraged - Stats Type: Rate - Unit: Number - Rollup Type: Average - Collection Level: 3

Average Read Latency per Virtual Machine Disk

This chart displays the top ten virtual machine disks with the highest average read latency in milliseconds. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

Table 422: Data Counters

Chart Label	Description
Average Read Latency per Virtual Machine Disk	Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request. Total latency = kernelLatency + deviceLatency. - Counter: totalReadLatency - Stats Type: Absolute - Unit: Milliseconds (ms) - Rollup Type: Average - Collection Level: 3

Average Write Latency per Virtual Machine Disk

This chart displays the top ten virtual machine disks with the highest average write latency in milliseconds. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

Table 423: Data Counters

Chart Label	Description
Average Write Latency per Virtual Machine Disk	Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request. Total latency = kernelLatency + deviceLatency. - Counter: totalWriteLatency - Stats Type: Absolute - Unit: Milliseconds (ms) - Rollup Type: Average - Collection Level: 3

Read IOPs per Virtual Machine Disk

This chart displays the top ten virtual machines with the highest number of read operations. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

Table 424: Data Counters

Chart Label	Description
Read IOPs per Virtual Machine Disk	Number of disk read commands completed on each virtual machine disk, per second. Read rate = blocks read per second × block size - Counter: numberReadAveraged - Stats Type: Rate - Unit: Number - Rollup Type: Average - Collection Level: 3

Write IOPs Per Virtual Machine Disk

This chart displays the 10 virtual machines with the highest number of write operations. Data is not displayed when the virtual machine is powered off.

This chart is located in the **Performance** view of the datastore **Performance** tab. The counter can also be displayed for datastore cluster charts.

Table 425: Data Counters

Chart Label	Description
Write IOPs per Virtual Machine Disk	Number of disk write commands completed on each virtual machine disk on the host. Write rate = blocks read per second × block size Counter: numberWriteAveraged

Chart Label	Description
	- Stats Type: Rate - Unit: Number - Rollup Type: Average - Collection Level: 3

Virtual Machine Observed Latency per Datastore

This chart displays the average datastore latency as observed by the virtual machines.

This chart is located in the **Performance** view of the datastore cluster **Performance** tab.

Table 426: Data Counters

Chart Label	Description
VM observed latency report per Datastore	This is the average datastore latency as observed by the virtual machines in the datastore cluster. - Counter: datastoreVMObservedLatency - Stats Type: Absolute - Unit: Microseconds - Rollup Type: Latest - Collection Level: 3

Hosts

The hosts charts contain information about CPU, disk, memory, network, GPU, and storage usage for hosts. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

CPU (%)

The CPU (%) chart displays CPU usage for the host.

This chart is located in the Home view of the Host **Performance** tab.

Table 427: Data Counters

Chart Label	Description
Usage	Actively used CPU, as a percentage of the total available CPU, for each physical CPU on the host. Active CPU is approximately equal to the ratio of the used CPU to the available CPU. Available CPU = # of physical CPUs × clock rate. 100% represents all CPUs on the host. For example, if a four-CPU host is running a virtual machine with two CPUs, and the usage is 50%, the host is using two CPUs completely. - Counter: usage - Stats Type: Rate - Unit: Percentage (%) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the host resources. However, if the value is constantly high, the host is probably lacking the CPU required to meet the demand. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

If performance is impacted, consider taking the following actions.

Table 428: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
4	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
5	Upgrade the physical CPUs or cores on the host if necessary.
6	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

CPU (MHz)

The CPU (MHz) chart displays CPU usage for the host.

This chart is located in the Home view of the Host **Performance** tab.

Table 429: Data Counters

Chart Label	Description
Usage	The sum, in megahertz, of the actively used CPU of all powered on virtual machines on a host. The maximum possible value is the frequency of the processors multiplied by the number of processors. For example, if you have a host with four 2GHz CPUs running a virtual machine that is using 4000MHz, the host is using two CPUs completely. $4000 \div (4 \times 2000) = 0.50$ - Counter: usagemhz - Stats Type: Rate - Unit: MegaHertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the host resources. However, if the value is constantly high, the host is probably lacking the CPU required to meet the demand. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.

If performance is impacted, consider taking the following actions.

Table 430: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
4	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
5	Upgrade the physical CPUs or cores on the host if necessary.
6	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

CPU (DPU)

The CPU (DPU) performance chart displays CPU usage of DPU.

This chart is available in the **DPU** pane of the **View** dropdown menu of vSphere Client for a host instance at **Monitor > Performance > Overview** tab.

Table 431: Data Counters

Chart Label	Description
Usage (%)	Get the CPU usage of a DPU. - Counter: usage - Stats Type: Rate - Unit: % - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Usage (Mhz)	Get the sum of actively used CPUs of all powered on virtual machines of a DPU. - Counter: usage - Stats Type: Rate - Unit: Mhz - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

CPU Usage

The CPU Usage chart displays CPU usage of the 10 virtual machines on the host with the most CPU usage.

This chart is located in the Virtual Machines view of the host **Performance** tab.

Table 432: Counters

Name	Description
<i>virtual_machine</i>	Amount of CPU actively being used by each virtual machine on the host. 100% represents all CPUs. For example, if a virtual machine has one virtual CPU that is running on a host with four CPUs and the CPU usage is 100%, the virtual machine is using one CPU resource. virtual CPU usage = usagemhz ÷ (number of virtual CPUs × core frequency) Note: The host's view of the CPU usage, not the guest operating system view. - Counter: usage - Stats Type: Rate - Unit: Percentage (%). Precision is to 1/100%. A value between 0 and 100. - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

Table 433: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's Virtual Machine view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

Disk (KBps)

The Disk (KBps) chart displays disk I/O of the host.

This chart is located in the Home view of the host **Performance** tab.

Table 434: Data Counters

Chart Label	Description
Usage	Average data I/O rate across all LUNs on the host. - Counter: usage - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 435: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.

#	Resolution
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Disk Rate (KBps)

The Disk Rate chart displays disk read and write rates for LUNs on a host, including average rates.

This chart is located in the **Home** view of the host **Performance** tab.

Table 436: Data Counters

Chart Label	Description
Read	Number of disk read commands completed on each disk on the host, per second. The aggregate number of all disk read commands is also displayed in the chart. Read rate = blocksRead per second × blockSize - Counter: read - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average - Collection Level: 3
Write	Number of disk write commands completed on each disk on the host, per second. The aggregate number of all disk write commands is also displayed in the chart. Write rate = blocksWritten per second × blockSize - Counter: write - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average - Collection Level: 3

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 437: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.

#	Resolution
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Disk Requests (Number)

The Disk Requests chart displays disk usage for the host.

This chart is located in the **Home** view of the host **Performance** tab.

Table 438: Data Counters

Chart Label	Description
Read Requests	Number of disk read commands completed on each LUN on the host. The aggregate number of all disk read commands is also displayed in the chart. - Counter: numberRead - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3
Write Requests	Number of disk write commands completed on each LUN on the host. The aggregate number of all disk write commands is also displayed in the chart. - Counter: numberWrite - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 439: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Disk (Number)

The Disk (Number) chart displays maximum queue depth for the top ten LUNs on a host.

This chart is located in the **Home** view of the host **Performance** tab.

Table 440: Data Counters

Chart Label	Description
Maximum Queue Depth	Maximum queue depth. Queue depth is the number of commands the SCSI driver queues to the HBA. - Counter: <code>maxQueueDepth</code> - Stats Type: Absolute - Unit: Number

Chart Label	Description
	- Rollup Type: Average - Collection Level: 1

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 441: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.

#	Resolution
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Disk (ms)

The Disk (ms) chart displays the amount of time taken to process commands on a host.

This chart is located in the **Home** view of the host **Performance** tab.

Table 442: Data Counters

Chart Label	Description
Highest Disk Latency	Highest latency value of all disks used by the host. Latency measures the time used to process a SCSI command issued by the guest OS to the virtual machine. The kernel latency is the time VMkernel takes to process an I/O request. The device latency is the time it takes the hardware to handle the request. Total latency = kernelLatency + deviceLatency. - Counter: maxTotalLatency - Stats Type: Absolute - Unit: Milliseconds (ms) - Rollup Type: Latest (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 443: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Disk (KBps)

The Disk (KBps) chart displays disk usage for the 10 virtual machines on the host with the most disk usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

Table 444: Data Counters

Chart Label	Description
<code>virtual_machine</code>	Sum of the data read from the virtual machine. - Counter: usage - Stats Type: Rate - Unit: KiloBytes per second (KBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 445: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.

#	Resolution
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

GPU Compute

The GPU Compute chart displays GPU usage for the host.

This chart is located in the Home view of the host **Performance** tab.

Note: The GPU charts are visible only if the host supports GPUs.

Table 446: Data Counters

Chart Labels	Description
Usage	The utilization of a GPU in percentages. - Counter: Utilization - Stats Type: Absolute - Unit: Percent (%) - Rollup Type: Average - Collection Level: 4

Chart Analysis

GPU Compute Usage (%) is vendor and architecture specific. High utilization can suggest that the GPU is used effectively, or it can indicate potential bottlenecks during peak workloads.

If a virtual machine is experiencing a Quality of Service (QoS) issue, please refer to the GPU vendor QoS documentation. If you use NVIDIA virtual GPU (vGPU) technology, see the NVIDIA [Virtual GPU Software for VMware vSphere](#) documentation for more information.

GPU Memory

The GPU Memory chart displays the amount of GPU memory used for the host in kilobytes and in percentage of the total available.

This chart is located in the Home view of the host **Performance** tab.

Note: The GPUs view is visible only if the host supports GPUs.

Table 447: Data Counters

Chart Labels	Description
Memory Usage in KB for <i>device_id</i>	The amount of GPU memory used in kilobytes. - Counter: Memory used - Stats Type: Absolute - Unit: Kilobytes (KB) - Rollup Type: Average - Collection Level: 4

Chart Labels	Description
Memory Usage for <i>device_id</i>	The amount of GPU memory used in percentage of the total available. - Counter: Memory used - Stats Type: Absolute - Unit: Percent (%) - Rollup Type: Average - Collection Level: 4

Chart Analysis

GPU Memory Usage (in Kilobytes) might not represent the active workload in the virtual machine, but the memory reserved for a virtual GPU.

Monitoring memory usage is essential to prevent out-of-memory errors, which can lead to application crashes or reduced performance.

If a virtual machine is experiencing a Quality of Service (QoS) issue, please refer to the GPU vendor QoS documentation. If you use NVIDIA virtual GPU (vGPU) technology, see the NVIDIA [Virtual GPU Software for VMware vSphere](#) documentation for more information.

GPU Compute Usage

The GPU Compute Usage chart displays GPU compute usage for the top 10 graphics devices on the host with the most GPU compute usage.

This chart is located in the GPUs view of the host **Performance** tab.

Note: The GPUs view is visible only if the host supports GPUs.

Table 448: Data Counters

Chart Labels	Description
Usage for <i>device_id</i>	The utilization of a GPU in percentages. - Counter: Utilization - Stats Type: Absolute - Unit: Percent (%) - Rollup Type: Average - Collection Level: 4

Chart Analysis

GPU Compute Usage (%) is vendor and architecture specific. High utilization may suggest that the GPU is used effectively, or it may indicate potential bottlenecks during peak workloads.

If a virtual machine is experiencing a Quality of Service (QoS) issue, please refer to the GPU vendor QoS documentation. If you use NVIDIA virtual GPU (vGPU) technology, see the NVIDIA [Virtual GPU Software for VMware vSphere](#) documentation for more information.

GPU Memory Allocation

The GPU Memory Allocation chart displays GPU memory allocation for the top 10 graphics devices on the host with the most GPU memory usage.

This chart is located in the GPUs view of the host **Performance** tab.

Note: The GPUs view is visible only if the host supports GPUs.

Table 449: Data Counters

Chart Labels	Description
Usage for <i>device_id</i>	The amount of GPU memory used in percentages of the total available. - Counter: Memory usage - Stats Type: Absolute - Unit: Percent (%) - Rollup Type: Average - Collection Level: 4

Chart Analysis

GPU memory usage tracks how much of the GPU onboard memory (VRAM) is in use.

GPU details measurements can show usage trends related to virtual machine power on failures. Device and driver specific placement constraints of vGPU on physical GPU devices, including fragmentation of Device Memory, exhaustion of Device Memory, or exhaustion of Assignable Processor Engines could contribute to power on failures.

If GPU memory usage is almost at capacity and a virtual machine power on failure occurs with a Virtual GPU that would exceed capacity, increase the number of compatible GPU hosts, and migrate one or more virtual machines to the new host to free up GPU Memory resources for the new VM. For further guidance see Knowledge Base article [KB 92222](#).

If a virtual machine is experiencing a Quality of Service (QoS) issue, please refer to the GPU vendor QoS documentation. If you use NVIDIA virtual GPU (vGPU) technology, see the NVIDIA [Virtual GPU Software for VMware vSphere](#) documentation for more information.

GPU Temperature

The GPU Temperature chart displays GPU temperature in degrees Celsius for the top 10 graphics devices on the host with the highest GPU temperature.

This chart is located in the GPUs view of the host **Performance** tab.

Note: The GPUs view is visible only if the host supports GPUs.

Table 450: Data Counters

Chart Labels	Description
GPU temperature for <i>device_id</i>	The temperature of a GPU in degrees Celsius. - Counter: Temperature - Stats Type: Absolute - Unit: Celsius (°C) - Rollup Type: Average

Chart Labels	Description
	Collection Level: 4

Chart Analysis

Monitoring GPU temperature helps prevent overheating, which can result in throttling and reduced performance. Overheating can also lead to hardware damage.

If a virtual machine is experiencing a Quality of Service (QoS) issue, please refer to the GPU vendor QoS documentation. If you use NVIDIA virtual GPU (vGPU) technology, see the NVIDIA [Virtual GPU Software for VMware vSphere](#) documentation for more information.

Memory (%)

The Memory (%) chart displays host memory usage.

This chart is located in the **Home** view of the host **Performance** tab.

Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

Table 451: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.

#	Resolution
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory (Balloon)

The Memory (Balloon) chart displays balloon memory on a host.

This chart is located in the **Home** view of the host **Performance** tab.

Table 452: Data Counters

Chart Label	Description
Balloon	Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host. - Counter: vmmemctl - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

Table 453: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory (DPU)

The Memory (DPU) performance chart displays memory used on DPU.

This chart is available in the **DPU** pane of the **View** dropdown menu of vSphere Client for a host instance at **Monitor > Performance > Overview** tab.

Table 454: Data Counters

Chart Label	Description
consumed (%)	Amount of memory consumed by a DPU. - Counter: consumed - Stats Type: Rate - Unit: % - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Memory (MBps)

The Memory (MBps) chart displays the swap in and swap out rates for a host.

This chart is located on the **Home** view of the Host **Performance** tab.

Table 455: Data Counters

Chart Label	Description
swpinRate	Average rate at which memory is swapped in from the host swap file. - Counter: swpinRate - Stats Type: Rate - Unit: MegaBytes per second (MBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Label	Description
swapoutRate	Average rate at which memory is swapped out to the host swap file. - Counter: swapoutRate - Stats Type: Rate - Unit: MegaBytes per second (MBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Host memory must be large enough to accommodate virtual machine workload. Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when there is a spike in virtual machine workload.

However, a consistently high memory usage value (94% or greater) indicates the host does not have the memory resources required to meet the demand. If the memory balloon and swap values are not high, performance is probably not affected. If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host requires more memory resources.

If the host is not lacking memory resources, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host is lacking memory resources or you notice a degradation in performance, consider taking the following actions.

Table 456: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory (MB)

The Memory (MB) chart displays memory data counters for hosts.

This chart is located in the **Home** view of the host **Performance** tab.

Note: Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system. Machine memory is the actual physical RAM in the host.

Not all counters are collected at collection level 1.

Table 457: Data Counters

Chart Label	Description
Active	Sum of the active guest physical memory of all powered on virtual machines on the host, plus memory used by basic VMKernel applications. Active memory is estimated by the VMkernel and is based on the current workload of the host. - Counter: active - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Balloon	Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines on the host. - Counter: vmmemctl - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Balloon Target	Sum of the balloon target memory of all powered on virtual machines on the host. If the balloon target value is greater than the balloon value, the VMkernel inflates the balloon, causing more virtual machine memory to be reclaimed. If the balloon target value is less than the balloon value, the VMkernel deflates the balloon, which allows the virtual machine to consume additional memory if needed. Virtual machines initiate memory reallocation. Therefore, it is possible to have a balloon target value of 0 and a balloon value greater than 0. - Counter: vmmemctltarget - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Consumed	Amount of machine memory used on the host. Consumed memory includes virtual machine memory, service console memory, and VMkernel memory. consumed memory = total host memory - free host memory - Counter: consumed - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Granted	Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory. Granted memory for a host includes the shared memory of each virtual machine on the host. - Counter: granted - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Shared Common	Amount of machine memory shared by all powered on virtual machines. Shared common memory consists of the entire pool of memory from which sharing is possible, including the amount of physical RAM required by the guest memory.

Chart Label	Description
	memory shared - memory shared common = amount of memory saved on the host from sharing - Counter: sharedcommon - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Swap Used	Sum of the memory swapped by all powered on virtual machines on the host. - Counter: swapused - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)

Chart Analysis

To ensure best performance, the host memory must be large enough to accommodate the active memory of the virtual machines. The active memory can be smaller than the virtual machine memory size. It allows you to over-provision memory, but still ensures that the virtual machine active memory is smaller than the host memory.

Transient high-usage values usually do not cause performance degradation. For example, memory usage can be high when several virtual machines are started at the same time or when a spike occurs in virtual machine workload. However, a consistently high memory usage value (94% or greater) indicates that the host is probably lacking the memory required to meet the demand. If the active memory size is the same as the granted memory size, the demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the memory usage value is high, and the host has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot handle the demand for memory. It leads to memory reclamation, which might degrade performance.

If the host has enough free memory, check the resource shares, reservation, and limit settings of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machines.

If the host has little free memory available, or if you notice a degradation in performance, consider taking the following actions.

Table 458: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.

#	Resolution
6	Add physical memory to the host.

Memory Usage

The Memory Usage chart displays memory usage for the 10 virtual machines on the host with the most memory usage.

This chart is located in the **Virtual Machines** view of the host **Performance** tab.

Virtual Machine Counters

Note: Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system.

Table 459: Data Counters

Chart Label	Description
Usage	Amount of guest physical memory currently in use on the virtual machine. - Counter: usage - Stats Type: Absolute - Unit: Percentage (%) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 460: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.

#	Resolution
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory Bandwidth (MBps)

The Memory Bandwidth chart displays DRAM and/or PMem Bandwidth graphs of the host.

This chart is available in the **Memory** pane of the **View** dropdown menu of vSphere Client for a host instance at **Performance > Overview** tab. The Memory pane gives information about Memory Bandwidth in addition to the Memory Utilization and Memory Reclamation information. It also gives the Memory Miss Rate, but only in Memory mode.

Note: PMem Bandwidth is only available on supported hosts configured in Memory Mode. If the host is not supported by vMMR, then the **Memory** option is not listed in the **View** dropdown menu.

Table 461: Data Counters

Chart Label	Description
Estimated DRAM Bandwidth	Total read and write bandwidth of the DRAM memory type. - Counter: bandwidth.total - Stats Type: Absolute - Unit: MegaBytes per second (MBps) - Rollup Type: Latest - Collection Level: 1 (4)
Estimated PMem Bandwidth	Total read and write bandwidth of the PMem memory type. - Counter: bandwidth.total - Stats Type: Absolute - Unit: MegaBytes per second (MBps) - Rollup Type: Latest - Collection Level: 1 (4)

Memory Miss Rate (%)

The Memory Miss Rate performance chart displays DRAM Miss Rate graph for the host.

This chart is available in the **Memory** pane of the **View** dropdown menu of vSphere Client for a host instance at **Performance > Overview** tab. The Memory pane gives information about Memory Miss Rate (DRAM) in addition to the memory consumption, memory reclamation, and memory bandwidth information.

Note: Miss rate is only supported on supported hosts in Memory Mode. If the host is not supported by vMMR, then the **Memory** option is not listed in the **View** dropdown menu.

Table 462: Data Counters

Chart Label	Description
Estimated DRAM Miss Rate	Get the current miss rate of DRAM memory type. - Counter: misstrate - Stats Type: Absolute - Unit: % - Rollup Type: Latest - Collection Level: 2 (4)

Network (Mbps)

The Network (Mbps) chart displays network usage for the host.

This chart is located in the **Home** view of the Host **Performance** tab.

Table 463: Host Counters

Chart Label	Description
Usage	Average rate at which data is transmitted and received across all NIC instances connected to the host. - Counter: usage - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 464: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Network Rate (Mbps)

The Network Rate chart displays network bandwidth on a host.

The Network Data Transmitted/Received chart for hosts is located in the **Home** view of the Host **Performance** tab.

Table 465: Data Counters

Chart Label	Description
Data Receive Rate	Rate at which data is received across the top ten physical NIC instances on the host. This represents the bandwidth of the network. The chart also displays the aggregated data receive rate of all physical NICs. - Counter: received - Stats Type: Rate - Unit: Kilobytes per second (Kbps) - Rollup Type: Average - Collection Level: 3 (4)
Data Transmit Rate	Rate at which data is transmitted across the top ten physical NIC instances on the host. This represents the bandwidth of the network. The chart also displays the aggregated data transmit rate of all physical NICs. - Counter: transmitted - Stats Type: Rate - Unit: Kilobytes per second (Kbps) - Rollup Type: Average - Collection Level: 3 (4)

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 466: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Network Packets (Number)

The Network Packets chart displays the network bandwidth on a host.

This chart is located in the **Home** view of the Host **Performance** tab.

Table 467: Data Counters

Chart Label	Description
Packets Received	Number of network packets received across the top ten physical NIC instances on the host. The chart also displays the aggregated value for all NICs. - Counter: packetRx - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3
Packets Transmitted	Number of network packets transmitted across the top ten physical NIC instances on the host. The chart also displays the aggregated value for all NICs. - Counter: packetTx - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 468: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.

#	Resolution
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Network (Mbps)

The Network (Mbps) chart displays network usage for the 10 virtual machines on the host with the most network usage. This chart is located in the **Virtual Machines** view of the host **Performance** tab.

Table 469: Data Counters

Chart Label	Description
<virtual machine>	Sum of the data transmitted and received across all virtual NIC instances connected to the virtual machine. - Counter: usage - Stats Type: Rate - Unit: Megabits per second (Mbps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 470: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Resource Pools

The resource pool charts contain information about CPU and memory usage for resource pools. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

CPU (MHz)

The CPU (MHz) chart displays CPU usage in the resource pool or vApp.

This chart is located in the Home view of the Resource Pool or vApp **Performance** tab.

Counters

Table 471: Data Counters

Chart Label	Description
Usage	CPU usage is the sum of the average CPU usage values of the virtual machines in the resource pool or vApp. CPU usage = number of cores * CPU frequency - Counter: usagemhz - Stats Type: Rate - Unit: Megahertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the resources available. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines in the resource pool. Generally, if the CPU usage value for a virtual machine is above 90% and the CPU ready value for a virtual machine is above 20%, performance is impacted.

If performance is impacted, consider taking the following actions.

Table 472: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Deploy single-threaded applications on uniprocessor virtual machines instead of SMP virtual machines.
3	Migrate one or more virtual machines to a new host.
4	Upgrade the physical CPUs or cores on each host if necessary.
5	Enable CPU-saving features, such as TCP Segmentation Offload.
6	Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

CPU Usage

The CPU Usage chart displays CPU usage of virtual machines in the resource pool or vApp. The chart displays the top 10 virtual machines with the highest CPU usage.

This chart is located in the Resource Pools & Virtual Machines view of the Resource Pool or vApp **Performance** tab.

Table 473: Data Counters

Chart Label	Description
<i>virtual_machine</i>	Amount of CPU actively used by virtual machines. - Counter: usagemhz - Stats Type: Rate - Unit: Megahertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

Table 474: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.

#	Resolution
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's Virtual Machine view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

Memory (MB)

The Memory (MB) chart displays memory usage in the resource pool or vApp.

This chart is located in the **Home** view of the resource pool or vApp **Performance** tab.

Table 475: Data Counters

Chart Label	Description
<i>resource_pool</i> or <i>vApp</i>	Sum of the active memory used by all virtual machines in the resource pool or vApp. Active memory is determined by the VMkernel and includes overhead memory. memory usage = active memory / configured virtual machine memory size - Counter: used - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average - Collection Level: 1

Chart Analysis

Memory usage is not necessarily an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

Table 476: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.
4	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: 1. Select the cluster, and click the Configure tab. 2. Under Services, click vSphere DRS. 3. click Edit. An Edit Cluster Settings dialog box opens. 4. Click Turn ON vSphere DRS, and click OK.
5	If the cluster is a DRS cluster: • Increase the number of hosts, and migrate one or more virtual machines to the new host. • Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.
6	Add more physical memory to one or more hosts.

Memory Consumed

The Memory Consumed chart displays the memory performance of all virtual machines in the resource pool or vApp.

This chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

For resource pools and virtual machines in a resource pool or vApp, this chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

Table 477: Data Counters

Chart Label	Description
<i>virtual_machine</i>	Amount of host memory used by the virtual machine for its guest operating system's physical memory. Memory overhead is not included in consumed memory. consumed memory = memory granted - memory saved from page sharing For example, if a virtual machine has 100 MB of memory that is shared equally with three other virtual machines, its portion of the shared memory is 25 MB (100 MB ÷ 4 VMs). This amount is counted in the memory consumed data counter. • Counter: consumed • Stats Type: Absolute • Unit: Megabytes (MB) • Rollup Type: Average (Minimum/Maximum) • Collection Level: 1 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 478: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory (MB)

The Memory (MB) chart displays memory data counters for resource pools or vApps.

Description

This chart is located in the **Home** view of the resource pool or vApp **Performance** tab.

Note: These data counter definitions are for virtual machines. At the resource pool level, the values are collected and totaled. The counter values in the chart represent the aggregate amounts of the virtual machine data. The counters that appear in the chart depend on the collection level set for your vCenter Server.

Table 479: Data Counters

Chart Label	Description
Active	Sum of the active guest physical memory of all powered on virtual machines in the resource pool. - Counter: active - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Balloon	Sum of the guest physical memory reclaimed by the balloon driver for all powered on virtual machines in the resource pool. - Counter: vmmemctl - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Consumed	Amount of the physical memory consumed by the virtual machine for the guest memory. Consumed memory does not include the overhead memory. It includes the shared memory and memory that might be reserved, but not actually used. consumed memory = memory granted – memory saved due to memory sharing - Counter: consumed - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Granted	Sum of the guest physical memory granted for all powered on virtual machines. Granted memory is mapped to the host's machine memory. - Counter: granted - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Shared	Amount of the guest physical memory shared with other virtual machines in the resource pool.
Swapped	Sum of the memory swapped by all powered on virtual machines in the resource pool. - Counter: swapused - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 480: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

vApps

The vApp charts contain information about CPU and memory usage for vApps. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

CPU (MHz)

The CPU (MHz) chart displays CPU usage in the vApp or resource pool.

This chart is located in the **Home** view of the vApp or resource pool **Performance** tab.

Counters

Table 481: Data Counters

Chart Label	Description
Usage	CPU usage is the sum of the average CPU usage values of the virtual machines in the resource pool or vApp. CPU usage = number of cores * CPU frequency - Counter: usagemhz - Stats Type: Rate - Unit: Megahertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage indicates that you are making the best use of the resources available. However, if the value is constantly high, the CPU demanded is likely greater than the CPU capacity available. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines in the resource pool. Generally, if the CPU usage value for a virtual machine is above 90% and the CPU ready value for a virtual machine is above 20%, performance is impacted.

If performance is impacted, consider taking the following actions.

Table 482: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Deploy single-threaded applications on uniprocessor virtual machines instead of SMP virtual machines.
3	Migrate one or more virtual machines to a new host.
4	Upgrade the physical CPUs or cores on each host if necessary.
5	Enable CPU-saving features, such as TCP Segmentation Offload.
6	Replace software I/O with dedicated hardware, such as iSCSI HBAs or TCP Segmentation Offload NICs.

CPU Usage

The CPU Usage chart displays CPU usage of each virtual machine in the vApp or resource pool.

This chart is located in the **Virtual Machines** view of the vApp or resource pool **Performance** tab.

Table 483: Data Counters

Chart Label	Description
<i>virtual_machine</i>	Amount of CPU actively used by virtual machines. - Counter: usagemhz - Stats Type: Rate - Unit: Megahertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

Table 484: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.

#	Resolution
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's Virtual Machine view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

Memory (MB)

The Memory (MB) chart displays memory usage in the vApp or resource pool.

This chart is located in the **Home** view of the vApp or resource pool **Performance** tab.

Table 485: Data Counters

Chart Label	Description
<i>resource_pool</i> or <i>vApp</i>	Sum of the active memory used by all virtual machines in the resource pool or vApp. Active memory is determined by the VMkernel and includes overhead memory. memory usage = active memory / configured virtual machine memory size - Counter: used - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average - Collection Level: 1

Chart Analysis

Memory usage is not necessarily an indicator of performance problems. Memory can be high if a host is swapping or ballooning, which can result in virtual machine guest swapping. In such cases, check for other problems, such as CPU over-commitment or storage latencies.

If you have constantly high memory usage in a cluster, resource pool, or vApp, consider taking the following actions.

Table 486: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	Verify that the balloon driver is enabled. The balloon driver is installed with VMware Tools and is critical to performance. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, it does not impact virtual machine performance.
3	If the balloon value is high, check the resource shares, reservations, and limits for the virtual machines and resource pools on the hosts. Verify that the host's settings are adequate and not lower than those set for the virtual machine. If free memory is available on the hosts and the virtual machines are experiencing high swap or balloon memory, the virtual machine (or resource pool, if it belongs to one) has reached its resource limit. Check the maximum resource limit set on that host.
4	If the cluster is not a DRS cluster, enable DRS. To enable DRS, perform the following tasks: 1. Select the cluster, and click the Configure tab. 2. Under Services, click vSphere DRS. 3. click Edit. An Edit Cluster Settings dialog box opens. 4. Click Turn ON vSphere DRS, and click OK.
5	If the cluster is a DRS cluster: • Increase the number of hosts, and migrate one or more virtual machines to the new host. • Check the aggressiveness threshold. If the value is low, increase the threshold. It might help avoid hot spots in the cluster.
6	Add more physical memory to one or more hosts.

Memory Consumed

The Memory Consumed chart displays the memory performance of the top ten virtual machines in the vApp or resource pool.

This chart is located in the **Virtual Machines** view of the vApp or resource pool **Performance** tab.

For resource pools and virtual machines in a resource pool or vApp, this chart is located in the **Resource Pools & Virtual Machines** view of the resource pool or vApp **Performance** tab.

Table 487: Data Counters

Chart Label	Description
<i>virtual_machine</i>	Amount of host memory used by the virtual machine for its guest operating system's physical memory. Memory overhead is not included in consumed memory. consumed memory = memory granted - memory saved from page sharing For example, if a virtual machine has 100 MB of memory that is shared equally with three other virtual machines, its portion of the shared memory is 25 MB (100 MB ÷ 4 VMs). This amount is counted in the memory consumed data counter. • Counter: consumed • Stats Type: Absolute • Unit: Megabytes (MB) • Rollup Type: Average (Minimum/Maximum) • Collection Level: 1 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 488: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Virtual Machines

The virtual machine charts contain information about CPU, disk, memory, network, storage, and fault tolerance for virtual machines. The help topic for each chart contains information about the data counters displayed in that chart. The counters available are determined by the collection level set for vCenter Server.

CPU (%)

The CPU (%) chart displays virtual machine CPU usage and ready values.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

Table 489: Data Counters

Chart Label	Description
Usage	Amount of actively used virtual CPU as a percentage of total available CPU. CPU usage is the average CPU utilization over all available virtual CPUs in the virtual machine.

Chart Label	Description
	For example, if a virtual machine with one virtual CPU is running on a host that has four physical CPUs and the CPU usage is 100%, the virtual machine is using one physical CPU completely. virtual CPU usage = usagemhz ÷ (number of virtual CPUs × core frequency) Note: This is the host's view of the CPU usage, not the guest operating system view. - Counter: usage - Stats Type: Rate - Unit: Percentage (%). Precision is to 1/100%. A value between 0 and 100. - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Ready	Percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU. CPU ready time is dependent on the number of virtual machines on the host and their CPU loads. At collection level 1, the average CPU ready time of all virtual CPUs on the virtual machine is displayed. At collection level 3, the average CPU ready time of each virtual CPU is also displayed. - Counter: ready - Stats Type: Rate - Unit: Percentage (%) - Rollup Type: Summation - Collection Level: 1

Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

Table 490: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's Virtual Machine view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.

#	Resolution
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

CPU Usage (MHz)

The CPU Usage (MHz) chart displays virtual machine CPU usage.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

Table 491: Data Counters

Chart Label	Description
Usage	Amount of actively used virtual CPU. Note: The host's view of the CPU usage, not the guest operating system view. - Counter: usagemhz - Stats Type: rate - Unit: MegaHertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A short spike in CPU usage or CPU ready indicates that you are making the best use of the virtual machine resources. However, if the CPU usage value for a virtual machine is above 90% and the CPU ready value is above 20%, performance is being impacted.

If performance is impacted, consider taking the following actions.

Table 492: CPU Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on every virtual machine on the host.
2	Set the CPU reservations for all high-priority virtual machines to guarantee that they receive the CPU cycles required.
3	Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked line chart on the host's Virtual Machine view shows the CPU usage for virtual machines on the host.
4	Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
5	Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time doesn't decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
6	Increase the amount of memory allocated to the virtual machine. This decreases disk and or network activity for applications that cache. This might lower disk I/O and reduce the need for the host to virtualize the hardware. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.

#	Resolution
7	Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the hypervisor's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
8	If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
9	Upgrade the physical CPUs or cores on the host if necessary.
10	Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

Disk (KBps)

The Disk (KBps) chart displays disk usage for the virtual machine.

It is located in the **Home** view of the virtual machine **Performance** tab.

Table 493: Data Counters

Chart Label	Description
Usage	Average data I/O rate across all virtual disks on the virtual machine. - Counter: usage - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 494: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Disk Rate (KBps)

The Disk Rate chart displays disk usage for the virtual machine.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

Table 495: Data Counters

Chart Label	Description
Read	Number of disk read commands completed on each virtual disk on the virtual machine, per second. The aggregate number of all disk read commands per second is also displayed in the chart. Read rate = blocksRead per second × blockSize - Counter: read - Stats Type: Rate - Unit: Kilobytes per second (KBps)

Chart Label	Description
	- Rollup Type: Average - Collection Level: 2
Write	Number of disk write commands completed on each virtual disk on the virtual machine, per second. The aggregate number of all disk write commands per second is also displayed in the chart. Write rate = blocksWritten per second × blockSize - Counter: write - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average - Collection Level: 2

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The kernelLatency data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The deviceLatency data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The queueLatency data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 496: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.

#	Resolution
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Disk Requests (Number)

The Disk Requests chart displays disk usage for the virtual machine.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

Table 497: Data Counters

Chart Label	Description
Read Requests	Number of disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all disk read commands is also displayed in the chart. - Counter: numberRead - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3
Write Requests	Number of disk write commands completed on each virtual disk on the virtual machine. The aggregate number of all disk write commands is also displayed in the chart. - Counter: numberWrite - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3

Chart Analysis

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read or write requests, check whether any such applications were running then.

The best ways to determine if your vSphere environment is experiencing disk problems is to monitor the disk latency data counters. You can use the advanced performance charts to view these statistics.

- The `kernelLatency` data counter measures the average amount of time, in milliseconds, that the VMkernel spends processing each SCSI command. For best performance, the value must be 0 -1 milliseconds. If the value is greater than 4 ms, the virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- The `deviceLatency` data counter measures the average amount of time, in milliseconds, to complete a SCSI command from the physical device. Depending on your hardware, a number greater than 15 ms indicates probable problems with the storage array. Move the active VMDK to a volume with more spindles or add disks to the LUN.
- The `queueLatency` data counter measures the average amount of time taken per SCSI command in the VMkernel queue. This value must always be zero. If not, the workload is too high and the array cannot process the data fast enough.

If the disk latency values are high, or if you notice other problems with disk I/O performance, consider taking the following actions.

Table 498: Disk I/O Performance Enhancement Advice

#	Resolution
1	Increase the virtual machine memory. It allows more operating system caching, which reduces I/O activity. Note: It might require you to increase the host memory. Increasing memory might reduce the need to store data because databases can utilize the system memory to cache data and avoid disk access. To verify that virtual machines have adequate memory, check swap statistics in the guest operating system. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
2	Defragment the file systems on all guests.
3	Deactivate antivirus on-demand scans on the VMDK and VMEM files.
4	Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. To increase throughput, consider array-side improvements.
5	Use Storage vMotion to migrate I/O-intensive virtual machines across multiple hosts.
6	Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
7	Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the <code>Disk.SchedNumReqOutstanding</code> parameter. For more information, see <i>vSphere Storage</i> .
8	For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. It alleviates disk spindle contention during periods of high use.
9	On systems with sizable RAM, deactivate memory trimming by adding the line <code>MemTrimRate=0</code> to the virtual machine's VMX file.
10	If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
11	For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select Allocate all disk space now . The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
12	Use the most current host hardware.

Virtual Disk Requests (Number)

The Virtual Disk Requests chart displays virtual disk usage for the virtual machine.

After you click **Overview** on the **Performance** tab of the virtual machine, you can view this chart by selecting **Home** from the **View** drop-down menu. It is available at collection (display) levels 3 and 4.

Table 499: Data Counters

Chart Label	Description
Read Requests	Number of virtual disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk read commands is also displayed in the chart. - Counter: numberRead - Stats Type: Absolute - Unit: Number - Rollup Type: Average - Collection Level: 2
Write Requests	Number of virtual disk write commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk write commands is also displayed in the chart. - Counter: numberWrite - Stats Type: Absolute - Unit: Number - Rollup Type: Average - Collection Level: 2

Virtual Disk Rate (KBps)

The Virtual Disk Rate chart displays virtual disk usage rate for the virtual machine.

After you click **Overview** on the **Performance** tab of the virtual machine, you can view this chart by selecting **Home** from the **View** drop-down menu. It is available only at collection levels 3 and 4.

Table 500: Data Counters

Chart Label	Description
Read Requests	Number of virtual disk read commands completed on each virtual disk on the virtual machine. The aggregate number of all virtual disk read commands per second is also displayed in the chart. Read rate = blocksRead per second × blockSize - Counter: read - Stats Type: Rate - Unit: KiloBytes per second (KBps) - Rollup Type: Average - Collection Level: 3
Write Requests	Number of virtual disk write commands completed on each virtual disk on the virtual machine per second. The aggregate number of all virtual disk write commands per second is also displayed in the chart. Write rate = blocksWritten per second × blockSize Counter: write

Chart Label	Description
	- Stats Type: Rate - Unit: KiloBytes per second (KBps) - Rollup Type: Average - Collection Level: 3

Memory (%)

The Memory (%) chart monitors virtual machine memory usage.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

Virtual Machine Counters

Note: Guest physical memory refers to the virtual hardware memory presented to a virtual machine for its guest operating system.

Table 501: Data Counters

Chart Label	Description
Usage	Amount of guest physical memory currently in use on the virtual machine. - Counter: usage - Stats Type: Absolute - Unit: Percentage (%) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 502: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory (MB)

The Memory (MB) chart displays virtual machine balloon memory.

This chart is located in the **Home** view of the virtual machine **Performance** tab.

Table 503: Data Counters

Chart Label	Description
Balloon	Amount of guest physical memory reclaimed from the virtual machine by the balloon driver. - Counter: vmmemctl - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 504: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory (MBps)

The Memory (MBps) chart displays virtual machine memory swap rates.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

Table 505: Data Counters

Chart Label	Description
swpinRate	Average rate at which memory is swapped into the virtual machine. - Counter: swpinRate - Stats Type: Rate - Unit: MegaBytes per second (MBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
swapoutRate	Average rate at which memory is swapped out of the virtual machine. - Counter: swapoutRate - Stats Type: Rate - Unit: MegaBytes per second (MBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If there is sufficient swap space, a high balloon value is not a performance issue. However, if the swpin and swapout values for the host are large, the host is probably lacking the memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. The host might require more memory resources. If it does not, check the resource shares, reservation, and limit of the virtual

machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If memory usage is high or you notice degradation in performance, consider taking the following actions.

Table 506: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory (MB)

The Memory (MB) chart displays memory data counters for virtual machines.

This chart is located in the **Home** view of the virtual machine **Performance** tab. It appears only at collection levels 2, 3, and 4.

In the following descriptions, the guest physical memory refers to the virtual-hardware memory presented to a virtual machine for its guest operating system. Machine memory is actual physical RAM in the host. Note that not all counters are collected at collection level 1.

Table 507: Data Counters

Chart Label	Description
Active	Amount of guest physical memory in use by the virtual machine. Active memory is estimated by VMkernel statistical sampling and represents the actual amount of memory the virtual machine needs. The value is based on the current workload of the virtual machine. - Counter: active - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Balloon	Amount of guest physical memory reclaimed from the virtual machine by the balloon driver. - Counter: vmmemctl - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Balloon Target	Desired amount of virtual machine balloon memory. Balloon target memory is estimated by the VMkernel.

Chart Label	Description
	If the balloon target amount is greater than the balloon amount, the VMkernel inflates the balloon amount, which reclaims more virtual machine memory. If the balloon target amount is less than the balloon amount, the VMkernel deflates the balloon, which allows the virtual machine to reallocate memory when needed. - Counter: vmmemctltarget - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Consumed	Amount of guest physical memory consumed by the virtual machine for guest memory. Consumed memory does not include overhead memory. It includes shared memory and memory that might be reserved, but not actually used. consumed memory = memory granted - memory saved due to memory sharing - Counter: consumed - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)
Shared	Amount of guest physical memory available for sharing. Memory sharing occurs through transparent page sharing. - Counter: shared - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)
Swapped	The amount of guest physical memory swapped out to the disk by the VMkernel. This data counter measures VMkernel swapping and not to guest OS swapping. swapped = swapout – swapin Note: In some cases, vMotion can skew these values and cause a virtual machine to arrive on a host with some memory already swapped out. As a result, the swapped value can be greater than the swapout – swapin value. - Counter: swapped - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4)

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size,

demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 508: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory Bandwidth (MBps)

The Memory Bandwidth chart for a displays the estimated DRAM Read bandwidth graph and estimated PMem Read Bandwidth graphs for a virtual machine.

This chart is available in the **Memory** pane of the **View** dropdown menu of vSphere Client for a virtual machine at **Performance > Overview** tab. The **Memory** pane gives information about Memory Bandwidth in addition to the Memory Usage and Memory Reclamation information.

Note: If the host supports vMMR and is in memory mode, **Memory** option is listed in the **View** dropdown menu. In other cases, Memory option is not listed.

Table 509: Data Counters

Chart Label	Description
Estimated DRAM Read Bandwidth	Current read bandwidth of DRAM memory type. - Counter: bandwidth.read - Stats Type: Absolute - Unit: MegaBytes per second (MBps) - Rollup Type: Latest - Collection Level: 2 (4)
Estimated PMem Read Bandwidth	Current read bandwidth of PMem memory type. - Counter: bandwidth.read - Stats Type: Absolute - Unit: MegaBytes per second (MBps) - Rollup Type: Latest

Chart Label	Description
	Collection Level: 2 (4)

Network (Mbps)

The Network (Mbps) chart displays network bandwidth for the virtual machine.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab.

Table 510: Virtual Machine Counters

Chart Label	Description
Usage	Average rate at which data is transmitted and received across all virtual NIC instances connected to the virtual machine. - Counter: usage - Stats Type: Rate - Unit: Kilobytes per second (KBps) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 1 (4)

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 511: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.

#	Resolution
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Network Rate (Mbps)

The Network Rate chart displays network usage for virtual machines.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab. It appears only at collection levels 3 and 4.

Table 512: Data Counters

Chart Label	Description
Data Receive Rate	Rate at which data is received across each virtual NIC instance on the virtual machine. - Counter: received - Stats Type: Rate - Unit: Kilobytes per second (Kbps) - Rollup Type: Average - Collection Level: 2 (4)
Data Transmit Rate	Rate at which data is transmitted across each virtual NIC instance on the virtual machine. - Counter: transmitted - Stats Type: Rate - Unit: Kilobytes per second (Kbps) - Rollup Type: Average - Collection Level: 2 (4)

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 513: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Network Packets (Number)

The Network Packets monitors network bandwidth for virtual machines.

This chart is located in the **Home** view of the Virtual Machine **Performance** tab. It appears only at collection levels 3 and 4.

Table 514: Data Counters

Chart Label	Description
Packets Transmitted	Number of network packets transmitted across the top ten virtual NIC instances on the virtual machine. The chart also displays the aggregated value for each NIC. - Counter: packetTx - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3

Chart Label	Description
Packets Received	Number of network packets received across the top ten virtual NIC instances on the virtual machine. The chart also displays the aggregated value for each NIC. - Counter: packetRx - Stats Type: Absolute - Unit: Number - Rollup Type: Summation - Collection Level: 3

Chart Analysis

Network performance depends on the application workload and network configuration. Dropped network packets indicate a bottleneck in the network. To determine whether packets are being dropped, use `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values.

If packets are being dropped, adjust the virtual machine shares. If packets are not being dropped, check the size of the network packets and the data receive and transfer rates. In general, the larger the network packets, the faster the network speed. When the packet size is large, fewer packets are transferred, which reduces the amount of CPU required to process the data. When network packets are small, more packets are transferred but the network speed is slower because more CPU is required to process the data.

Note: In some instances, large packets might result in a high network latency. To check the network latency, use the VMware AppSpeed performance monitoring application or a third-party application.

If packets are not being dropped and the data receive rate is slow, the host is probably lacking the CPU resources required to handle the load. Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different vSwitches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.

If you experience network-related performance problems, also consider taking the following actions.

Table 515: Networking Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine.
2	If possible, use vmxnet3 NIC drivers, which are available with VMware Tools. They are optimized for high performance.
3	If virtual machines running on the same host communicate with each other, connect them to the same vSwitch to avoid transferring packets over the physical network.
4	Assign each physical NIC to a port group and a vSwitch.
5	Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
6	Ensure that the physical NIC capacity is large enough to handle the network traffic on that vSwitch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10 Gbps). Alternatively, consider moving some virtual machines to a vSwitch with a lighter load or to a new vSwitch.
7	If packets are being dropped at the vSwitch port, increase the virtual network driver ring buffers where applicable.
8	Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that NICs with 1 Gbps are not reset to 100 Mbps because they are connected to an older switch.
9	Verify that all NICs are running in full duplex mode. Hardware connectivity problems might result in a NIC resetting itself to a lower speed or half duplex mode.

#	Resolution
10	Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TSO-Jumbo Frames are enabled where possible.

Space in GB

The Space in GB chart displays space utilization data counters for virtual machines.

This chart is located in the **Storage** view of the virtual machine **Performance** tab.

Table 516: Data Counters

Chart Label	Description
Allocated	Total amount of logical datastore space provisioned by an administrator for the virtual machine. It is the storage size up to which the virtual machine files on datastores can grow. This includes log files, VMX files, and other miscellaneous files. Allocated space is not always in use. - Counter: provisioned - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1
Used	Amount of physical datastore space in use by the virtual machine files. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1
Not Shared	Amount of datastore space that belongs only to this virtual machine and is not shared with other virtual machines. Only unshared space is guaranteed to be reclaimed for the virtual machine if, for example, it is moved to a different datastore and then back again. The value is an aggregate of all unshared space for the virtual machine, across all datastores. - Counter: unshared - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Space Utilization by Datastores

The Space Utilization by Datastores chart displays the amount of space used by a virtual machine on different datastores in the data center.

Note: This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by Datastores chart is located in the **Storage** view of the virtual machine **Performance** tab.

Table 517: Data Counters

Chart Label	Description
<i>datastore_name</i>	Amount of disk space in the datastore currently in use by the virtual machine. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Space Utilization By File Type

The Space Utilization by File Type chart displays the datastore usage by virtual machine files.

Note: This chart does not show historical statistics. It only shows the most recently available data, which may be up to 30 minutes late, depending on when the last statistics rollup occurred. In addition, statistics are not collected across all datastores at one time. They are collected asynchronously.

The Space Utilization by File Type chart is located in the **Storage** view of the virtual machine **Performance** tab.

Datastore counters

Table 518: Data Counters

File Type	Description
Virtual Disks	Amount of disk space used by virtual disk files. Virtual disk files store the contents of the virtual machine's hard disk drive, including information that you write to a virtual machine's hard disk - the operating system, program files, and data files. The files have the extension <code>.vmdk</code> and appear as a physical disk drive to a guest operating system. Note: Delta disks, which also have an extension <code>.vmdk</code>, are not included in this file type. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)

File Type	Description
Swap Files	Amount of disk space used by swap files. Swap files back up the virtual machine's physical memory. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Snapshots	Amount of disk space used by virtual machine snapshot files. Snapshot files store information about virtual machine snapshots. They include snapshot state files and delta disk files. A snapshot state file stores the running state of the virtual machine at the time of the snapshot. It has the extension <code>.vmsn</code>. A delta disk file stores the updates made by the virtual machine to the virtual disks after a snapshot is taken. - Counter: used - Stats Type: Absolute - Unit: GigaBytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Other VM Files	Amount of disk space used by all other virtual machine files, such as configuration files and log files. - Counter: used - Stats Type: Absolute - Unit: Gigabytes (GB) - Rollup Type: Latest - Collection Level: 1 (4)
Total Space	Amount of disk space used by the virtual machine. total space = virtual disk space + swap file space + snapshot space + other VM file space

Chart Analysis

The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks. If possible, you can provision more space to the datastore, or you can add disks to the datastore or use shared datastores.

If snapshot files are consuming high datastore space, consider consolidating them to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface. For information about consolidating the data center, see the vSphere documentation.

Fault Tolerance Performance Counters

The fault tolerance charts contain information about CPU and memory for fault tolerant virtual machines.

Note: The performance charts and help topics for fault tolerance are available only when you have vSphere Fault Tolerance enabled. If you select a link for a secondary virtual machine in the thumbnail section of the Resource Pools and Virtual Machines view of the cluster Performance tab, the navigation in the inventory updates to the primary virtual machine. This occurs because secondary machines are not displayed in the inventory.

CPU (MHz)

The CPU (MHz) chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the virtual machine **Performance** tab. It is available only at collection levels 3 and 4.

Table 519: Data Counters

Name	Description
Usage	The average amount of virtual CPU, per CPU instance, in use on the primary and secondary fault tolerant virtual machines. - Counter: usagemhz - Stats Type: Rate - Unit: Megahertz (MHz) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 3 (4)

Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

Table 520: CPU Performance Enhancement Advice

#	Resolution
1	Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level.
2	Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations.
3	Verify that the network and datastore connections for both virtual machines are similar.
4	Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion.
5	Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion.

CPU System Time for vCPU (%)

The CPU System Time chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. It is available only at collection levels 3 and 4.

Table 521: Data Counters

Chart Label	Description
System	Amount of time spent on system processes on each virtual CPU in the virtual machine. Note: This is the host view of the CPU usage, not the guest operating system view. - Counter: system - Stats Type: Delta

Chart Label	Description
	- Unit: Percentage (%) - Rollup Type: Summation - Collection Level: 3

Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

Table 522: CPU Performance Enhancement Advice

#	Resolution
1	Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level.
2	Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations.
3	Verify that the network and datastore connections for both virtual machines are similar.
4	Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion.
5	Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion.

CPU Used Time for vCPU (%)

The CPU Used Time chart displays virtual CPU usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. The chart is available only for collection levels 3 and 4.

Table 523: Data Counters

Chart Label	Description
used	Amount of used virtual CPU as a percentage of total available CPU on the primary and secondary virtual machines. A high value indicates excessive use of CPU resources. - Counter: used - Stats Type: Delta - Unit: Percentage (%) - Rollup Type: Summation - Collection Level: 1

Chart Analysis

A large discrepancy in CPU usage between the primary and secondary virtual machines might indicate performance problems. The CPU ready, system, and wait times of each virtual machine should be synchronized. A large discrepancy in these values might indicate performance problems. Consider taking the following actions.

Table 524: CPU Performance Enhancement Advice

#	Resolution
1	Verify that the primary and secondary hosts are in the same CPU model family and have similar CPU configurations. For best results, use CPUs with the same stepping level.
2	Verify that the CPU resource reservations set for both virtual machines are consistent within the cluster. VMware HA plans for a worst-case scenario by considering all powered on virtual machines in a cluster and finding the maximum memory and CPU reservations.
3	Verify that the network and datastore connections for both virtual machines are similar.
4	Turn off power management (also known as power-capping) in the BIOS. If power management is enabled, the secondary host might enter lower performance, power-saving modes. Such modes can leave the secondary virtual machine with insufficient CPU resources, potentially making it impossible for the secondary to complete all tasks completed on a primary in a timely fashion.
5	Turn off hyperthreading in the BIOS. If hyperthreading is enabled and the secondary virtual machine is sharing a CPU with another demanding virtual machine, the secondary virtual machine might run too slowly to complete all tasks completed on the primary in a timely fashion.

Memory Active (MB)

The Memory Active chart displays active memory usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. It is not available at collection level 1.

Table 525: Data Counters

Chart Label	Description
Active	Amount of guest physical memory in use by the fault tolerant virtual machine. Active memory is estimated by VMkernel statistical sampling and represents the actual amount of memory the virtual machine needs. Additional, unused memory may be swapped out or ballooned with no performance impact. - Counter: active - Stats Type: Absolute - Unit: Megabytes (MB) - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4) Make sure that the primary and secondary virtual machines have enough memory. If the secondary system is not provisioned well, it might slow down performance of the primary virtual machine or fail.

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If sufficient swap space is available, a high balloon value does not cause performance problems. However, if the swapin and swapout values for the host are large, the host is probably lacking the amount of memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. A free memory value of 6% or less indicates that the host cannot meet the memory requirements. This leads to memory reclamation, which might degrade performance. If the active memory size is the same as the granted memory size, demand for memory is greater than the memory resources available. If the active memory is consistently low, the memory size might be too large.

If the host has enough free memory, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If little free memory is available, or if you notice degradation in performance, consider taking the following actions.

Table 526: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Memory Swapout (MB)

The Memory Swapout chart displays the swapout memory usage for fault tolerant virtual machines.

This chart is located in the **Fault Tolerance** view of the Virtual Machine **Performance** tab. It is not available at collection level 1.

Table 527: Data Counters

Chart Label	Description
Swapout	Amount of machine memory written to the VMkernel swap file. - Counter: swapout - Stats Type: Absolute - Unit: MegaBytes - Rollup Type: Average (Minimum/Maximum) - Collection Level: 2 (4) Make sure that the primary and secondary virtual machines have enough memory and that the swapout value is not high. If the secondary system is not provisioned well, it might slow down performance of the primary virtual machine or fail.

Chart Analysis

A virtual machine's memory size must be slightly larger than the average guest memory usage. This enables the host to accommodate workload spikes without swapping memory among guests. Increasing the virtual machine memory size results in more overhead memory usage.

If there is sufficient swap space, a high balloon value is not a performance issue. However, if the swapin and swapout values for the host are large, the host is probably lacking the memory required to meet the demand.

If a virtual machine has high ballooning or swapping, check the amount of free physical memory on the host. The host might require more memory resources. If it does not, check the resource shares, reservation, and limit of the virtual machines and resource pools on the host. Verify that the host settings are adequate and not lower than those set for the virtual machine.

If memory usage is high or you notice degradation in performance, consider taking the following actions.

Table 528: Memory Performance Enhancement Advice

#	Resolution
1	Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
2	Verify that the balloon driver is enabled. The VMkernel regularly reclaims unused virtual machine memory by ballooning and swapping. Generally, this does not impact virtual machine performance.
3	Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
4	If the memory reservation of a virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
5	Migrate one or more virtual machines to a host in a DRS cluster.
6	Add physical memory to the host.

Working with Advanced and Custom Charts

Use advanced charts, or create your own custom charts, to see more performance data. Advanced charts can be useful when you are aware of a problem but need more statistical data to pinpoint the source of the trouble.

Advanced charts include the following features:

- More information. Hover over a data point in a chart and details about that specific data point are displayed.
- Customizable charts. Change chart settings. To create your own charts, save custom settings.
- Export to spreadsheet.
- Save to image file or spreadsheet.

View Advanced Performance Charts in the vSphere Client

Advanced charts support data counters that are not supported in other performance charts.

For example, you can use the Power view to display the `vm.power` counter and monitor virtual machine power information for hosts.

1. Navigate to an inventory object in the vSphere Client.
2. Click the **Monitor** tab, and click **Performance**.
3. Click **Advanced**.
4. Optional: To view a different chart, select an option from the **View** list.
The amount of historical data displayed in a chart depends on the collection interval and statistics level set for vCenter Server.
- 5.

Optional: To view the selected performance chart in a pop-up window, click the **Chart Context Menu** icon () and click **Popup Chart**. The chart opens in a new browser window enabling you to navigate within the vSphere Client while keeping the performance chart open in a separate window. You can learn more about this feature in this video.

Change Advanced Chart Settings

You can customize a performance chart by specifying the objects to monitor, the counters to include, the time range, and chart type. You can customize preconfigured chart views and create chart views.

1. Select an inventory object in the vSphere Client.
2. Click the **Monitor** tab, and click **Performance**.
3. Click **Advanced**.
4. Click **Chart Options**.
5. In Chart Metrics, select a metric group for the chart.
6. Select a timespan range for the metric group from the **Timespan** menu.

Time range options are not active unless you select **Custom interval** in the **Timespan** menu.

If you choose **Custom interval**, do one of the following.

- Select **Last** and set the number of hours, days, weeks, or months for the amount of time to monitor the object.
- Select **From** and select the beginning and end dates.

You can also customize the time range options by customizing the statistics collection interval setting.

7. In Counters, select the data counters to display in the chart.
8. In Target Objects, select the inventory objects to display in the chart.
9. Select the chart type.

When selecting the stacked graph option, consider the following.

- You can select only one item from the list of measurements.
- Per-virtual-machine stacked graphs are available only for hosts.
- To display information about the counter's function and whether the selected metric can be stacked for per-virtual-machine graphs, click a counter description name.

10. Click **OK**.

Create a Custom Advanced Chart

You can create your own charts by saving customized chart settings. New charts are added to the **View** menu and will appear there only when charts for the selected object are being displayed.

1. Select an inventory object in the vSphere Client.
2. Click the **Monitor** tab.
3. Under **Performance**, click **Advanced**, and click **Chart Options**.
4. Customize chart settings by selecting the chart metrics, counters, objects, timespan and chart type that you want to save as a view.
5. Click **Save Options As....**
6. Enter a name for your custom chart options.
7. Click **OK**.

The chart settings are saved and an entry for your chart is added to the **View** menu on the **Advanced Performance** page.

Related Links

[Change Advanced Chart Settings on page 3407](#)

You can customize a performance chart by specifying the objects to monitor, the counters to include, the time range, and chart type. You can customize preconfigured chart views and create chart views.

Delete a Custom Advanced Chart View

You can delete custom chart views from the vSphere Client.

1. Select an inventory object in the vSphere Client.
2. Click the **Monitor** tab, and click **Performance**.
3. Click **Advanced**.
4. Click **Chart Options**.
5. Select a chart and click **Delete Options**.
6. Click **OK** to confirm deletion.

The chart is deleted, and it is removed from the **View** menu.

Save Chart Data to a File

You can save data from the advanced performance charts to a file in various graphic formats or in comma-separated values (CSV) format.

1. In the vSphere Client, select an inventory object.
2. Click the **Monitor** tab, and click **Performance**.
3. Click **Advanced**.
4.  Click the **Chart Context Menu** icon ().
5. Choose a file type to export to and click the respective option.

The file is saved to your local system in the format you specified.

Troubleshoot and Enhance Performance

This section presents tips for identifying and solving performance problems.

The suggestions in this section are not meant to be a comprehensive guide to diagnosing and troubleshooting problems in the virtual environment. It is meant to provide information about some common problems that can be solved without contacting VMware Technical Support.

Solutions for Consistently High CPU Usage

Temporary spikes in CPU usage are not necessarily a concern, but consistently high CPU usage might indicate a problem. You can use the CPU performance charts to monitor CPU usage for hosts, clusters, resource pools, virtual machines, and vApps.

- Host CPU usage constantly is high. A high CPU usage value can lead to increased ready time and processor queuing of the virtual machines on the host.
- Virtual machine CPU usage is above 90% and the CPU ready value is above 20%. Application performance is impacted.
- The host probably is lacking the CPU resources required to meet the demand.
- There might be too many virtual CPUs relative to the number of physical processor cores.
- There might be an IO storage or networking operation that places the CPU in a wait state.
- The Guest OS generates too much load for the CPU.

- Verify that VMware Tools is installed on every virtual machine on the host.
- Compare the CPU usage value of a virtual machine with the CPU usage of other virtual machines on the host or in the resource pool. The stacked bar chart on the host's **Virtual Machine** view shows the CPU usage for all virtual machines on the host.
- Determine whether the high ready time for the virtual machine resulted from its CPU usage time reaching the CPU limit setting. If so, increase the CPU limit on the virtual machine.
- Increase the CPU shares to give the virtual machine more opportunities to run. The total ready time on the host might remain at the same level if the host system is constrained by CPU. If the host ready time does not decrease, set the CPU reservations for high-priority virtual machines to guarantee that they receive the required CPU cycles.
- Increase the amount of memory allocated to the virtual machine. This can potentially decrease disk and/or network activity for applications that cache. This might lower disk I/O and/or network traffic, which could in turn reduce CPU utilization. Virtual machines with smaller resource allocations generally accumulate more CPU ready time.
- Reduce the number of virtual CPUs on a virtual machine to only the number required to execute the workload. For example, a single-threaded application on a four-way virtual machine only benefits from a single vCPU. But the ESXi's maintenance of the three idle vCPUs takes CPU cycles that could be used for other work.
- If the host is not already in a DRS cluster, add it to one. If the host is in a DRS cluster, increase the number of hosts and migrate one or more virtual machines onto the new host.
- Upgrade the physical CPUs on the host if necessary.
- Use the newest version of hypervisor software, and enable CPU-saving features such as TCP Segmentation Offload, large memory pages, and jumbo frames.

Solutions for Memory Performance Problems

Host machine memory is the hardware backing for guest virtual memory and guest physical memory. For best performance, the host machine memory should be at least slightly larger than the combined active memory of the virtual machines on the host. A virtual machine's memory size should be slightly larger than the average guest memory usage. Increasing the virtual machine memory size results in more overhead memory usage.

- Memory usage is constantly high (94% or greater) or constantly low (24% or less).
- Free memory consistently is 6% or less and swapping frequently occurs.
- The host probably is lacking the memory resources required to meet the combined active memory size of all the running VMs.
- Host machine memory resources are not enough to meet the demand, which leads to memory reclamation (like swapping) and degraded performance.
- Verify that VMware Tools is installed on each virtual machine. The balloon driver is installed with VMware Tools and is critical to performance.
- Verify that the balloon driver is enabled. The VMkernel regularly attempts to reclaim the unused virtual machine memory by ballooning and, if necessary, swapping. Generally, this does not impact virtual machine performance.
- Reduce the memory space on the virtual machine, and correct the cache size if it is too large. This frees up memory for other virtual machines.
- If the memory reservation of the virtual machine is set to a value much higher than its active memory, decrease the reservation setting so that the VMkernel can reclaim the idle memory for other virtual machines on the host.
- Migrate one or more virtual machines to a host in a DRS cluster.
- Add physical memory to the host.

Solutions for Storage Performance Problems

Datastores represent storage locations for virtual machine files. A storage location can be a VMFS volume, a directory on Network Attached Storage, or a local file system path. Datastores are platform-independent and host-independent.

- Snapshot files are consuming a lot of datastore space.

- The datastore is at full capacity when the used space is equal to the capacity. Allocated space can be larger than datastore capacity, for example, when you have snapshots and thin-provisioned disks.
- Consider consolidating snapshots to the virtual disk when they are no longer needed. Consolidating the snapshots deletes the redo log files and removes the snapshots from the vSphere Client user interface.
- You can provision more space to the datastore if possible, or you can add disks to the datastore or use shared datastores.

Solutions for Disk Performance Problems

Use the disk charts to monitor average disk loads and to determine trends in disk usage. For example, you might notice a performance degradation with applications that frequently read from and write to the hard disk. If you see a spike in the number of disk read/write requests, check if any such applications were running then.

- The value for the kernelLatency data counter is greater than 4 ms.
- The value for the deviceLatency data counter is greater than 15 ms indicates that there are probably problems with the storage array.
- The queueLatency data counter measures above zero.
- Spikes in latency.
- Unusual increases in read/write requests.
- The virtual machines on the host are trying to send more throughput to the storage system than the configuration supports.
- The storage array probably is experiencing internal problems.
- The workload is too high and the array cannot process the data fast enough.

- The virtual machines on the host are trying to send more throughput to the storage system than the configuration supports. Check the CPU usage, and increase the queue depth.
- Move the active VMDK to a volume with more spindles or add disks to the LUN.
- Increase the virtual machine memory. Doing so should allow for more operating system caching, which can reduce I/O activity. Note: It may require you to increase the host memory. Increasing memory might reduce the need to store data because some workloads can utilize system memory to cache data and avoid disk access.
- Check swap statistics in the guest operating system to verify that virtual machines have adequate memory. Increase the guest memory, but not to an extent that leads to excessive host memory swapping. Install VMware Tools so that memory ballooning can occur.
- Defragment the file systems on all guests.
- If appropriate for your environment, deactivate antivirus on-demand scans on the VMDK and VMEM files.
- Use the vendor's array tools to determine the array performance statistics. When too many servers simultaneously access common elements on an array, the disks might have trouble keeping up. Consider array-side improvements to increase throughput.
- Use Storage vMotion to distribute I/O-intensive virtual machines across multiple hosts.
- Balance the disk load across all physical resources available. Spread heavily used storage across LUNs that are accessed by different adapters. Use separate queues for each adapter to improve disk efficiency.
- Configure the HBAs and RAID controllers for optimal use. Verify that the queue depths and cache settings on the RAID controllers are adequate. If not, increase the number of outstanding disk requests for the virtual machine by adjusting the `Disk.SchedNumReqOutstanding` parameter.
- For resource-intensive virtual machines, separate the virtual machine's physical disk drive from the drive with the system page file. This alleviates disk spindle contention during periods of high use.
- On systems with sizable RAM, deactivate memory trimming by adding the line `MemTrimRate=0` to the virtual machine's VMX file.
- If the combined disk I/O is higher than a single HBA capacity, use multipathing or multiple links.
- For ESXi hosts, create virtual disks as preallocated. When you create a virtual disk for a guest operating system, select **Allocate all disk space now**. The performance degradation associated with reassigning additional disk space does not occur, and the disk is less likely to become fragmented.
- Use the most current hypervisor software.

Solutions for Poor Network Performance

Network performance is dependent on application workload and network configuration. Dropped network packets indicate a bottleneck in the network. Slow network performance can be a sign of load-balancing problems.

Network problems can manifest in many ways:

- Packets are being dropped.
- Network latency is high.
- Data transfer rate is low.

Network problems can have several causes:

- Virtual machine network resource shares are too few.
- Network packet size is too large, which results in high network latency. Use the VMware AppSpeed performance monitoring application or a third-party application to check network latency.
- Network packet size is too small, which increases the demand for the CPU resources needed for processing each packet. Host CPU, or possibly virtual machine CPU, resources are not enough to handle the load.

- Determine whether packets are being dropped by using `esxtop` or the advanced performance charts to examine the `droppedTx` and `droppedRx` network counter values. Verify that VMware Tools is installed on each virtual machine.
- Check the number of virtual machines assigned to each physical NIC. If necessary, perform load balancing by moving virtual machines to different virtual switches or by adding more NICs to the host. You can also move virtual machines to another host or increase the host CPU or virtual machine CPU.
- If possible, use `vmxnet3` NIC drivers, which are available with VMware Tools. They are optimized for high performance.
- If virtual machines running on the same host communicate with each other, connect them to the same virtual switch to avoid the cost of transferring packets over the physical network.
- Assign each physical NIC to a port group and a virtual switch.
- Use separate physical NICs to handle the different traffic streams, such as network packets generated by virtual machines, iSCSI protocols, vMotion tasks.
- Ensure that the physical NIC capacity is large enough to handle the network traffic on that virtual switch. If the capacity is not enough, consider using a high-bandwidth physical NIC (10Gbps) or moving some virtual machines to a virtual switch with a lighter load or to a new virtual switch.
- If packets are being dropped at the virtual switch port, increase the virtual network driver ring buffers where applicable.
- Verify that the reported speed and duplex settings for the physical NIC match the hardware expectations and that the hardware is configured to run at its maximum capability. For example, verify that 1Gbps NICs are not reset to 100Mbps because they are connected to an older switch.
- Verify that all NICs are running in full duplex mode. Hardware connectivity issues might result in a NIC resetting itself to a lower speed or half duplex mode.
- Use vNICs that are TCP Segmentation Offload (TSO)-capable, and verify that TCP Segmentation Offload-Jumbo Frames are enabled where possible.

Empty Performance Charts

No graphics or data are displayed in performance charts.

When data is missing for a performance chart, the chart appears empty and you see the message `No data available`.

The causes for missing data in performance charts that are described herein are based on the assumption that the default roll-up configuration for the vCenter Server system has not changed. The causes include but are not limited to the following scenarios:

- Metrics introduced in ESXi 5.0 are not available for hosts running earlier versions.
- Data is deleted when you remove or add objects to vCenter Server.
- Performance charts data for inventory objects that were moved to a new site by VMware vCenter Site Recovery Manager is deleted from the old site and not copied to the new site.
- Performance charts data is deleted when you use VMware vMotion across vCenter Server instances.
- Real-time statistics are not available for disconnected hosts or powered off virtual machines.
- Real-time statistics are collected on hosts and aggregated every 5 minutes. After six data points are collected for approximately 30 minutes, they are rolled up to the vCenter Server database to create the 1-Day statistic. 1-Day statistics might not be available for 30 minutes after the current time, depending on when the sample period began.
- The 1-Day statistics are rolled up to create one data point every 30 minutes. If a delay occurs in the roll-up operation, the 1-Week statistics might not be available for 1 hour after the current time. It takes 30 minutes for the 1-Week collection interval, plus 30 minutes for the 1-Day collection interval.
- The 1-Week statistics are rolled up to create one data point every two hours. If a delay occurs in the roll-up operations, the 1-Month statistics might not be available for 3 hours. It takes 2 hours for the 1-Month collection interval, plus 1 hour for the 1-Week collection interval.
- The 1-Month statistics are rolled up to create one data point every day. If a delay occurs in the roll-up operations, the statistics might not be available for 1-day and 3 hours. It takes one day for the past year collection interval, plus 3 hours for the past month collection interval. During this time, the charts are empty.

- There is no solution available.

Solutions for Performance Issues With Memory Mode

Intel provides an Intel Optane Persistent Memory Mode (PMem), in which the hardware hides the DRAM as cache and exposes PMem as the system memory. Although PMem is cheaper than DRAM, it has higher access latency which could lead to performance degradation issues.

Performance degradation issues while using PMem in Memory Mode:

- If the active memory is higher than some percentage of the available DRAM memory, the VM performance could degrade as memory accesses may have to go to PMem.
- Any two random VMs can have higher level of page collision due to hardware implementation resulting in degraded VM performance even if the available DRAM memory is fully utilized.

vSphere performs real-time monitoring using vSphere Memory Monitoring and Remediation (vMMR). vMMR collects both Host level and VM level memory statistics like DRAM/PMem bandwidth, latency, miss rate which provides additional insights. This is helpful in analyzing if the host is experiencing issues due to running in Memory Mode and if there is a need to re-distribute workload. If the analysis indicates that some workloads are having performance degradation due to running on systems configured in Memory mode, the VMs can be migrated from the current host to other hosts to balance the load.

- Two preconfigured default alarms have been added based on the newly collected statistics. One at the host level (Host Memory Mode High Active DRAM Usage) and another at the VM level (Virtual Machine High PMem Bandwidth Usage). If the alarm condition is met, an event will be published to trigger the corresponding alarm. If alarm is triggered, then it is indicative that there may be some issue with Memory Mode on this system. You can further analyze, if it is a real problem by using Performance charts.
- You can also create custom alarms based on new performance metrics at cluster/host or VM level. For example, you can create an alarm if the PMem bandwidth is observed to be higher than some value. vMMR alarms only work on systems with Memory Mode. For more information on how to create a custom alarm, see [Create or Edit Alarms](#) section.
- If the host is experiencing a performance issue, then it can be narrowed down to CPU, Memory, Disk, or Network issue by looking at the existing performance charts.

In vSphere Client, a new **Memory** pane is added under Performance tab for both Host and VM. The host level performance chart displays read/write bandwidth, missrate for the different memory types (DRAM, PMem). The VM Level performance chart displays DRAM and PMem read bandwidth of the VM. These performance charts help customers to analyze the statistics and determine if their application workload is regressed due to Memory Mode. For example, if significant higher PMem bandwidth is observed, it is indicative of issues originating due to Memory Mode and can be investigated further.

- You can also plot custom based performance charts at host and VM level by using **Advanced** option and plotting some of Memory Mode related metrics.
- From the VMs tab of an ESXi host, you can view a list containing performance information about all virtual machines that reside on the host. To display information about the Memory Mode impact on a virtual machine, click the view

columns () icon and select the newly added Active Memory, DRAM Read Bandwidth, and PMem Read Bandwidth metrics. This is helpful in identifying the most impacted VMs.

For more information on vMMR, see *vSphere Memory Monitoring and Remediation* document.

Monitoring Guest Operating System Performance

This section describes how to install and view VMware-specific performance data for virtual machines that run Microsoft Windows operating systems. VMware provides performance counters that enable you to view data on many aspects of guest operating system performance for the Microsoft Windows Perfmon utility.

Some virtualization processes dynamically allocate available resources depending on the status, or utilization rates, of virtual machines in the environment. It can make obtaining accurate information about the resource utilization (CPU utilization, in particular) of individual virtual machines, or applications running within virtual machines, difficult. VMware now provides virtual machine-specific performance counter libraries for the Windows Perfmon utility. It enables application administrators to access accurate virtual machine resource utilization statistics from within the Windows Perfmon utility.

You can take advantage of virtualized CPU performance counters to use performance tuning tools inside the guest operating system. See the *vSphere Virtual Machine Administration* documentation.

Enable Statistics Collection for Guest Operating System Performance Analysis

VMware-specific performance objects are loaded into Microsoft Windows Perfmon and enabled when VMware Tools is installed.

To display a performance chart for any performance object, you must add counters. See [View Performance Statistics for Windows Guest Operating Systems](#)

View Performance Statistics for Windows Guest Operating Systems

You can display VMware specific statistics in the Microsoft Windows Perfmon utility.

Verify that a virtual machine with a Microsoft Windows operating system and VMware Tools is installed.

1. Open a console to the virtual machine and log in.
2. Select **Start > Run**.
3. Enter **Perfmon** and press **Enter**.
4. In the Performance dialog box, click **Add**.
5. In the Add Counters dialog box, select **Use local computer counters**.
6. Select a virtual machine performance object.
Virtual machine performance object names begin with **VM**.
7. Select the counters that you want to display for that object.
8. If the performance object has multiple instances, select the instances you want to display.
9. Click **Add**.
The Performance dialog box displays data for the selected performance object.
10. Click **Close** to close the Add Counter dialog box and return to the Performance dialog box.

Monitoring Host Health Status

You can use the vSphere Client to monitor the state of host hardware components, such as CPU processors, memory, fans, and other components.

The host health monitoring tool allows you to monitor the health of a variety of host hardware components including:

- CPU processors
- Memory
- Fans
- Temperature
- Voltage
- Power
- Network
- Battery

- Storage
- Cable/Interconnect
- Software components
- Watchdog
- PCI devices
- Other

The host health monitoring tool presents data gathered using Systems Management Architecture for Server Hardware (SMASH) profiles. The information displayed depends on the sensors available on your server hardware. SMASH is an industry standard specification providing protocols for managing a variety of systems in the data center. For more information, see <http://www.dmtf.org/standards/smash>.

You can monitor host health status by connecting the vSphere Client to a vCenter Server system. You can also set alarms to trigger when the host health status changes.

Note: The interpretation of hardware monitoring information is specific for each hardware vendor. Your hardware vendor can help you understand the results of the host hardware components monitoring.

Monitor Hardware Health Status in the vSphere Client

You can monitor the health status of host hardware in the vSphere Client

1. Select a host in the vSphere Client .
2. Click the **Monitor** tab, and click **Hardware Health**.
3. Select the type of information to view.

Option	Description
Sensors	Displays all sensors arranged in a tree view. If the status is blank, the health monitoring service cannot determine the status of the component.
Storage Sensors	Displays the storage sensors.
Alerts and Warnings	Displays alerts and warnings.
System Event Log	Displays the system event log.

Monitoring and Diagnostics of vSphere Health

You can monitor vSphere Health and identify issues in vSphere.

Skyline Health

Skyline Health for vSphere enables you to identify and resolve potential issues before they have an impact to your environment.

vSphere telemetry data is collected from data-centers globally. This data is further used to analyze pre-conditions in your vSphere environment related to stability and incorrect configurations. These issues are reported under Skyline Health for vSphere and resolution recommendations are provided. It enables VMware to enhance the issue detection without updating the vSphere installation. You can check the health of vSphere Host and vCenter Server.

Check health of the system

You can use the Skyline Health findings for vSphere to monitor the health of the system. You can run health findings and send the data to VMware for advanced analysis.

- You must participate in the Customer Experience Improvement Program to use online health findings.

- To perform the online health checks, vCenter Server must be able to communicate over the Internet.

Note: If Customer Improvement Experience Program (CEIP) is not activated, the Internet connectivity finding is unavailable.

- Navigate to vCenter Server or select a host in the vSphere Client navigator.
- Click the **Monitor** tab, and click **Skyline Health**.
- Expand **Online Health Connectivity** category and select the type of information to view.

Option	Description
Advisor	The Skyline Advisor, included with your Production or Premier Support contract, enhances your proactive support experience with additional features and functionality, including automatic support log bundle transfer with Log Assist.
Audit CEIP Collected Data	When you choose to participate in CEIP, VMware collects technical information about your use of the products and services in CEIP reports on a regular basis. This information does not personally identify you. The latest collected data can be found on the vCenter Server.
Customer experience improvement program (CEIP)	The CEIP check verifies whether the program is enabled for your vCenter Server. If it is not, click the button next to the health check, navigate to the CEIP page, and enroll in the program. To enable CEIP, click Configure CEIP
Online health connectivity	The Internet connectivity check verifies that vCenter Server can communicate with <i>vmware.com</i> over the HTTPS/443 interface. If communication is successful, this check passes. If communication fails, the check indicates that the Internet connectivity is not available.
Number of Online findings performed successfully	This provide information on number of online health findings performed successfully.

- Expand the following categories to view the related health warnings:
 - Security**
 - Self support Diagnostics**
 - Storage**
 - General**
- Click the **RETEST** button to run the health findings and update the results immediately.
You can click the **Ask VMware** button to open a knowledge base article that describes the health finding and provides information about how to resolve the issue.

VMware Skyline Health Diagnostics Tool

VMware Skyline Health Diagnostics tool is a self-service diagnostics platform. This tool helps to detect issues in the vSphere and vSAN product line and provides recommendations in the form of a Knowledge Base article or remediation procedures to address the identified issues. vSphere administrators can use this tool to troubleshoot issues before contacting VMware Global Support Service.

VMware Skyline Health Diagnostics supports VMware vCenter plugin registration from its user interface. You can register with VMware vCenter Server on or above 7.0. After successful plugin registration, you can trigger the diagnostics on VMware vCenter Server, monitor the task status and view the analysis report from vSphere client interface itself.

For more information, refer to the *Installation, Setup and Operations Guide for VMware Skyline Health Diagnostics for vSphere* guide.

Register VMware Skyline Health Diagnostics plugin from vSphere Client

You can register VMware Skyline Health Diagnostics plugin with VMware vCenter server to perform all the VMware Skyline Health Diagnostics operations from vSphere HTML client. As an operator, you might want to register VMware

Skyline Health Diagnostics plugins on vSphere client as single point for performing Health and Diagnostics and view reports.

1. VMware vCenter Server with 8.0 U1 or greater version should be present.
 2. You should have valid vSphere SSO administrator credentials.
1. Log in to the vSphere Client using the supported browser and SSO credentials.
 2. In the top-menu, click **Menu > Skyline Health Diagnostics**. The landing page of Skyline Health Diagnostics plugin gets displayed.
 3. Click on the go to [Skyline Health Diagnostics](#) link next to Get OVA Files, to open **Download Skyline Health Diagnostics OVA** page.
 4. Click **DOWNLOAD NOW** button to start the OVA download.
 5. Navigate to vSphere Client and click **Install Skyline Health Diagnostics** to open the install OVA wizard.
 6. Use the OVA downloaded from local and proceed for installation. Provide all the required details for deploying the OVA and submit the task.
 7. Post successful deployment of OVA and plugin registration, the Skyline Health Diagnostics Plugin landing page changes to Skyline Health Diagnostics Plugin home page.

Deployment of Skyline Health Diagnostics appliance and registration of VMware Skyline Health Diagnostics plugin with VMware vCenter Server is successful and the client plugin of Skyline Health Diagnostics is visible in vSphere client.

For more information, refer to the *Installation, Setup and Operations Guide for VMware Skyline Health Diagnostics for vSphere* guide.

Use VMware Skyline Health Diagnostics Plugin from vSphere Client

You can use the VMware Skyline Health Diagnostic plugin from vSphere Client.

As a vSphere system administrator, you can use this plugin to initiate diagnostics run, run health checks, view the analysis report and manage the configurations. For detailed steps in using this plug-in from vSphere Client, refer to the *Installation, Setup and Operations Guide for VMware Skyline Health Diagnostics for vSphere* guide.

Monitoring Events, Alarms, and Automated Actions

vSphere includes a user-configurable events and alarms subsystem. This subsystem tracks events happening throughout vSphere and stores the data in log files and the vCenter Server database. This subsystem also enables you to specify the conditions under which alarms are triggered. Alarms can change state from mild warnings to more serious alerts as system conditions change, and can trigger automated alarm actions. This functionality is useful when you want to be informed, or take immediate action, when certain events or conditions occur for a specific inventory object, or group of objects.

Events

Events are records of user actions or system actions that occur on objects in vCenter Server or on a host. Actions that might be recorded as events include, but are not limited to, the following examples:

- A license key expires
- A virtual machine is powered on
- A user logs in to a virtual machine
- A host connection is lost

Event data includes details about the event such as who generated it, when it occurred, and what type of event it is.

The types of events are:

Table 529: Event Types

Event Type	Description
Error	Indicates that a fatal problem has occurred in the system and terminates the process or operation.
Warning	Indicates that there is a potential risk to the system which needs to be fixed. This event does not terminate the process or operation.
Information	Describes that the user or system operation is completed successfully.
Audit	Provides important audit log data which is crucial for the security framework. The audit log data includes information about what is the action, who did it, when it occurred, and the IP address of the user. You can learn more about this in the <i>vSphere Security</i> guide.

Alarms

Alarms are notifications that are activated in response to an event, a set of conditions, or the state of an inventory object. An alarm definition consists of the following elements in the vSphere Client:

- Name and description - Provides an identifying label and description.
- Targets - Defines the type of object that is monitored.
- Alarm Rules - Defines the event, condition, or state that triggers the alarm and defines the notification severity. It also defines operations that occur in response to triggered alarms.
- Last modified - The last modified date and time of the defined alarm.

Alarms have the following severity levels:

- Normal – green
- Warning – yellow
- Alert – red

Alarm definitions are associated with the object selected in the inventory. An alarm monitors the type of inventory objects specified in its definition.

For example, you might want to monitor the CPU usage of all virtual machines in a specific host cluster. You can select the cluster in the inventory, and add a virtual machine alarm to it. When enabled, that alarm monitors all virtual machines running in the cluster and triggers when any one of them meets the criteria defined in the alarm. To monitor a specific virtual machine in the cluster, but not others, select that virtual machine in the inventory and add an alarm to it. To apply the same alarms to a group of objects, place those objects in a folder and define the alarm on the folder.

Note: You can activate, deactivate, and modify alarms only from the object in which the alarm is defined. For example, if you defined an alarm in a cluster to monitor virtual machines, you can only modify that alarm through the cluster. You cannot change the alarm at the individual virtual machine level. However, you can disable or re-enable an alarm on a specific leaf inventory entity. For example, if a virtual machine related alarm is enabled on a vCenter level it can be disabled on a specific virtual machine and thus the alarm status for the specific entity will not be calculated and alarm state will be reset to gray.

Alarm Actions

Alarm actions are operations that occur in response to the trigger. For example, you can have an email notification sent to one or more administrators when an alarm is triggered.

Note: Default alarms are not preconfigured with actions. You must manually set what action occurs when the triggering event, condition, or state occurs.

View Events in vSphere Client

You can view events associated with a single object or view all vSphere events. The events list for a selected inventory object includes events associated with child objects. vSphere keeps information about tasks and events for a specified time period. You can configure this time period. By default, this period is set to 30 days.

1. Select an inventory object in the vSphere Client.
2. Click the **Monitor** tab, and click **Events**.
3. Click the expand row icon in each row to view the details for the corresponding event. You can expand multiple rows simultaneously.
4. Optional: Click the **OPEN IN NEW TAB** button at the top right corner of the datagrid to open the events view in a new tab.

When you open the new tab from **Entity > Monitor > Events** menu, the name of the entity is displayed next to the title in the new tab. Click **REFRESH** at the top right of the datagrid in the new tab to refresh the events in the current page.

5. Click the **FILTER** option to filter the events based on event category, type of user, and a selected time range.
6. Optional: Select a event and click **COPY TO CLIPBOARD** option to copy the event to clipboard.
7. Optional: The `Triggering Event` column in the **Alarms** tab of bottom panel adjacent to `Recent Tasks` tab displays "object-name: alarm-name". You can click on the hyperlink text to navigate to a particular object's events view. All the event details related to the selected alarm gets displayed. You can click on the **BACK TO ALL EVENTS** button to return to the default event-list view.

For the alarms, that do not have any underneath event, a plain text is used.

Export Events in vSphere Client

You can export events using the vSphere client into a `.csv` file.

You can export selected or all events using the export option provided by vSphere Client.

1. Select an inventory object in the vSphere Client.
2. Click the **Monitor** tab, and click **Events**.
3. Click **Export** to export events into a `.csv` file.
 - Select **All events** option from the dropdown menu to export all the events.
 - Select **Only selected** option from the dropdown menu to export only the selected events from the event list.
 - Select **Advanced export** option from the dropdown menu to export events based on customized settings, like event category, user type, time range and so on.

View System Event Log

vSphere records events in the vCenter Server database. System log entries include such information as who generated the event, when the event was created, and the type of event.

- Required privilege: **Global > Diagnostics**
1. Select a host in the vSphere Client navigator.
 2. Click **Monitor** tab, and click **Hardware Health**.
 3. Click **SYSTEM EVENT LOG**.

Export System Event Log Data

You can export all or part of the system event log data stored in the vCenter Server database.

Required Role: **Read-only**

1. Select a host in the vSphere Client.
2. Click the **Monitor** tab, and click **Hardware Health**.
3. Click the **SYSTEM EVENT LOG** option.
4. Click **EXPORT**.

Consolidate Identical Events

The event burst filter monitors the incoming stream of events for identical events over a short time. To optimize the storage size of events, events that occur repeatedly are consolidated into a single event before storing them in the database or the remote syslog server.

Event bursts can be seen in various scenarios which include,

- An existing hardware failure.
- An automated solution which logs-in and logs-out of vCenter Server often.

The events burst filter is enabled by default. It consolidates all types of events except those events on the allow list. An event burst is defined as more than one identical event per second. Two events are identical if,

- The events are of the same type.
- The events are on the same inventory object.
- The events are issued by the same user.

Note: The remaining event-specific data is not used to determine if two events are identical.

An event burst is detected after seeing 30 identical events for less than 30 seconds. These 30 events are stored in the VC database or the remote syslog server. Consolidation of the bursting events starts from the 31st event. The bursting events do not go into the database, but the consolidated event gets stored in the database.

In cases of bursting events, the burst filter affects only the events that go to the database and remote syslog stream. Alarms triggered by bursting events and `EventHistoryCollector` objects are not affected.

- `com.vmware.vc.EventBurstStartedEvent` - the start of an event burst.
- `com.vmware.vc.EventBurstEndedEvent` - the end of an event burst.
- `com.vmware.vc.AllEventBurstsEndedEvent` - the end of all the events bursts.
- `com.vmware.vc.EventBurstCompressedEvent` - the consolidated event after the event burst.

Each of this event contains,

- `eventType` - the event type of the bursting event.
- `objectId` - the entity of the bursting event.
- `userName` - the user name of the bursting event.

Also the compressed event contains,

- `count` - the number of identical events since the start of an event burst. The count starts from the 31st event.

- `burstStartTime` - Time of the 31st event of an event burst.

Note: The timestamps of the burst filter events are not related to the bursting events.

Configure Events Burst Filter

You can configure the basic and advanced settings for the events burst filter in the advanced vCenter Server settings.

Burst filter can be in the following configurations.

- **Enabled:** If `compressToDb` or `compressToSyslog` are enabled, then the burst filter detects bursts, posts events for them, and consolidates events before storing them in the database or sending them to the remote syslog server.
- **Disabled:** The burst filter does not detect bursts, does not post events for them, and does not consolidate events before storing them in the database or sending them to the remote syslog server. If both `compressToDb` and `compressToSyslog` are disabled, then the burst filter detects bursts and posts events for them but does not consolidate events.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. **Advanced Settings**
4. Click **Edit**.
5. Click on the filter text box present in the **Name** column of the table header. Type `vpxd.event`, and press Enter.
6. For the basic setting,
 - a) Enable or disable the `vpxd.event.burstFilter.compressToDb` option.
This option enables you to compress the bursting of events in the database. The default value for the setting is enabled.
 - b) Enable or disable the `vpxd.event.burstFilter.compressToSyslog` option.
This option enables you to compress the bursting of events in the syslog server. The default value for the setting is disabled.
7. The burst settings are not visible with their default values in **Advanced Settings**. However, a key-value can be entered for the configuration to take effect.

Note: Adding or modifying advanced settings might result in an unstable environment. Configuration parameters cannot be removed once they are added.

- a) Configure a key-value to enable the `config.vpxd.event.burstFilter.enabled` option.
This parameter enables burst detection of the VC database and syslog. On disabling burst filter, bursts are not detected and the events are stored in the database or remote syslog server by default without compression. The default value is Enabled.
- b) Configure a key-value to enable the `config.vpxd.event.burstFilter.rateEvents` option.
You can configure the number of events after which a burst is detected. If the configured limit is reached, the event sequence is considered a burst and is compressed while storing it in the VC database. The default value is 30. If the default is set, then the first 30 events are not compressed. They are used to detect a burst. If a burst is detected, then a single consolidated event with their count replaces the subsequent events after the first 30 events.
Note: It does not compress the real-time stream of events like alarms and `EventManager` property collector updates.
- c) Configure a key-value to enable the `config.vpxd.event.burstFilter.rateSeconds` option.
This parameter allows you to configure the number of seconds the data is stored for an event from the last time the event occurred. Higher values load the burst filter cache, and the bursts are detected even if the incoming flow of

events is not even. Lower values exert lesser load on burst filter cache, and bursts may not be detected if there is a temporary halt of incoming events. The default value is 30 seconds.

- d) Configure a key-value to enable the `config.vpxd.event.burstFilter.cacheSize` option.

The number of unique events tracked by the burst filter. When the cache is full, the burst filter stops monitoring the incoming new unique events and passes them through the VC database and syslog. The default value is 128000.

Note: Setting a lower value than default reduces the memory footprint but it also reduces the capacity to detect bursts for a large inventory.

- e) Configure a key-value to enable the `config.vpxd.event.burstFilter.whitelist` option.

You can configure the burst filter to compress all types of events except the events that are on the allow list. This parameter allows you to separate event types. It can avoid in monitoring few event types by the burst filter. The default value is

```
<whiteList>vim.event.UserLoginSessionEvent;vim.event.UserLogoutSessionEvent</whiteList>
```

- f) Configure a key-value to enable the `config.vpxd.event.burstFilter.compressRatio` option.

When the burst filter starts compressing an event, it posts a compressed event when the burst ends or on every X event. If a burst occurs for days, it posts some events instead of muting them all until the burst ends. The default value is 3600. The events are logged every 3600 events.

Note: Restart the `vmware-vpxd` service for all changes to take effect. For more information how to restart services in the vSphere Client, see *vCenter Server and Host Management documentation*.

Streaming Events to a Remote Syslog Server

After you enable remote streaming, vCenter Server starts streaming and only the newly generated events are streamed to the remote syslog server.

All syslog messages begin with a specific prefix. You can distinguish the vCenter Server events from other syslog messages by their `Event` prefix.

The syslog protocol limits the length of syslog messages to 1024 characters. Messages that are longer than 1024 characters split into multiple syslog messages.

In the syslog server, events have the following format:

```
<syslog-prefix> : Event [eventId] [partInfo] [createdTime] [eventType] [severity] [user]
[target] [chainId] [desc]
```

Item	Description
syslog-prefix	Displays the syslog prefix. The <code><syslog-prefix></code> is determined by the remote syslog server configuration.
eventId	Displays the unique ID of the event message. The default value is <code>Event</code> .
partInfo	Displays whether the message is split into parts.
createdTime	Displays the time when the event was generated.
eventType	Displays the event type.
severity	Displays whether the event is a piece information, a warning, or an error.
user	Displays the name of the user who generated the event.
target	Displays the object the event refers to.
chainId	Displays information about the parent or the group ID.
desc	Displays the description of the event.

Split of Long Event Message into Multiple Syslog Messages

Events that are longer than 1024 characters split into multiple syslog messages in the following manner:

```
<syslog-prefix> : Event [eventId][1-X][payload-part-1]<syslog-prefix> : Event [eventId][2-X][payload-part-2]
...
<syslog-prefix> : Event [eventId][X-X][payload-part-X]
```

The *X* stands for the number of the event message parts.

Forward vCenter Server Log Files to Remote Syslog Server

You can forward the vCenter Server log files to a remote syslog server to conduct an analysis of your logs.

Log in to the vCenter Server Management Interface as root.

Note:

ESXi can be configured to send log files to a vCenter Server rather than storing them to a local disk. The recommended maximum numbers of supported hosts to collect logs from is 30. See [KB 2003322](#) for information on how to configure ESXi log forwarding. This feature is intended for smaller environments with stateless ESXi hosts. For all other cases, use a dedicated log server. Using vCenter Server to receive ESXi log files might impact vCenter Server performance.

1. In the vCenter Server Appliance Management Interface, select **Syslog**.
2. In the Forwarding Configuration section, click **Configure** if you have not configured any remote syslog hosts. Click **Edit** if you already have configured hosts.
3. In the Create Forwarding Configuration pane, enter the server address of the destination host. The maximum number of supported destination hosts is three.
4. From the **Protocol** drop-down menu, select the protocol to use.

Menu Item	Description
TLS	Transport Layer Security
TCP	Transmission Control Protocol
REL	Reliable Event Logging Protocol
UDP	User Datagram Protocol

5. In the **Port** text box, enter the port number to use for communication with the destination host.
6. In the Create Forwarding Configuration pane, click **Add** to enter another remote syslog server.
7. Click **Save**.
8. Verify that the remote syslog server is receiving messages.
9. In the Forwarding Configuration section, click **Send Test Message**.
10. Verify on the remote syslog server that the test message was received.
The new configuration settings are shown in the Forwarding Configuration section.
11. Optional: To view the logs that are exported, log in to the vCenter Server Appliance Management Interface as root user and view the list of log files in `/etc/vmware-syslog/` directory.

Configure Streaming of Events to a Remote Syslog Server

You can also configure writing of events to the vCenter Server streaming facility. Streaming events is supported only for the vCenter Server. The streaming of events to a remote syslog server is deactivated by default. You can enable

and configure the streaming of vCenter Server events to a remote syslog server from the vCenter Server Management Interface.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Expand **Settings** option, and select **Advanced Settings**.
4. Click **EDIT SETTINGS**.
5. Click on the filter text box present in the **Name** column of the table header. Type `vpzd.event`, and press Enter.
6. Activate or deactivate the `vpzd.event.syslog.enabled` option.
By default, this option is deactivated.
7. Click **SAVE**.

Retention of Events in the vCenter Server Database

You can configure vCenter Server to retain events in the database for a limited period. Discarding events periodically ensures optimal performance of the database.

In new installations of vCenter Server 6.5, the event clean up option is enabled by default and the default number of days to retain event messages in the database is 30. You can change this value to the number of days that you want to retain the event messages in the database.

If you are upgrading or migrating from vCenter Server 7.0 or earlier, and you had the event cleanup option enabled, your setting to retain events is preserved after the upgrade or migration to vCenter Server 6.5.

After the retention period ends, the events are deleted from the database. However, there might be latency in the deletion of the events that are older than the configured retention setting.

Configure Database Settings

You can configure the maximum number of database connections that can occur simultaneously. To limit the growth of the vCenter Server database and save storage space, you can configure the database to discard information about tasks or events periodically.

Note: Do not use the database retention options, if you want to keep a complete history of tasks and events for your vCenter Server.

1. In the vSphere Client, navigate to the vCenter Server instance.
2. Select the **Configure** tab.
3. Under **Settings**, select **General**.
4. Click **Edit**.
5. In Edit vCenter general settings window, click **Database**.
6. In **Maximum connections** field, enter the required number of connections.

Note: Do not change this value unless one of these issues exists in your system.

- If your vCenter Server system performs many operations frequently and performance is critical, increase the number of connections.
- If the database is shared and connections to the database are costly, decrease the number of connections.

7. Enable **Task cleanup** option for vCenter Server to delete the retained tasks periodically.
8. Optional: In **Task retention (days)** field, enter a value (in days).
Information about tasks that are performed on this vCenter Server system is discarded after the specified number of days.
9. Enable **Event cleanup** option for vCenter Server to clean up the retained events periodically.
10. Optional: In **Event retention (days)** field, enter a value (in days).
Information about events for this vCenter Server system is discarded after the specified number of days.
Note: Monitor vCenter Server database consumption and disk partition in the vCenter Server Management Interface.
Warning: Increasing the events retention to more than 30 days results in significant increase of vCenter database size and can shutdown the vCenter Server. Ensure that you increase the vCenter database accordingly.
11. Click **SAVE**.
12. Restart the vCenter Server to apply changes manually.

View Triggered Alarms

Triggered alarms are visible in several locations throughout the vSphere Client.

1. To view alarms triggered on a selected inventory object, click the **Monitor** tab, click **Issues and Alarms**.
2. Click **Triggered Alarms**.

Live Refresh of Recent Tasks and Alarms

Live refresh for recent tasks and alarms that result from operations that other users perform in your environment is enabled by default. Now, all recent tasks and alarms are visible to all of the users with permissions to view them.

Set an Alarm in the vSphere Client

In the vSphere Client you define alarms in the alarm definition wizard. You access the alarm definition wizard from the **Configure** tab, under **More**.

Create or Edit Alarms

To monitor your environment, you can create alarm definitions in the vSphere Client . You can access alarm definitions in the **Configure** tab.

Required Privilege: **Alarms > Create alarm** or **Alarms > Modify alarm**

Create alarms in the **Configure** tab.

1. Select an inventory object, click the **Configure** tab, and click **More**.
2. Click **Alarm Definitions**.
3. Click **Add** to add an alarm.
4. Select an alarm, click **Enable** to activate an alarm.
5. Select an alarm, click **Disable** to deactivate an alarm.
6. Select an alarm, click **Delete** to delete an alarm.
7. Click **Edit** to edit an alarm.
8. You can also edit an alarm definition by selecting **Monitor > Triggered Alarms**
 - a) Select an alarm listed under **Alarm Name**.
 - b) Click **Edit Alarm Definition** to edit an alarm.

Specify Alarm Name, Description, and Target

Settings of an alarm definition include alarm name, description, and target.

- Required Privilege: **Alarms > Create alarm** or **Alarms > Modify alarm**
- In the Alarms Definitions page, click Add. See [Create or Edit Alarms](#).

1. Type a name and description.
2. Select the type of inventory object that this alarm monitors from the **Target type** drop-down menu
Depending on the type of target that you choose to monitor, the summary that follows the **Target**, change.
3. Click **Next**.

Note: Depending on the type of activity that you choose to monitor, the options on the Alarm Rule page, change.

Set alarm rule.

Specify Alarm Rules

You can select and configure the events, states, or conditions that trigger the alarm from the **Alarm Rule** page in the **New Alarm Definition** wizard.

Required Privilege: **Alarms > Create alarm** or **Alarms > Modify alarm**

An alarm definition must contain at least one trigger before you can save it.

1. Select a trigger from the drop-down menu.
The combined event triggers are displayed. You can set the rule for a single event only. You must create multiple rules for multiple events.
2. Click **Add Argument** to select an argument from the drop-down menu.
It supports **ALL** based expression, option to select **ANY** is not available. You must create a separate alarm definition for each trigger. The **OR** operator is not supported in the vSphere Client. However, you can combine more than one condition trigger with **AND** operator.

3. Select an operator from the drop-down menu.
4. Select an option from the drop-down menu to set the threshold for triggering an alarm
5. Select severity of the alarm from the drop-down menu.
 You can set the condition to either **Show as Warning** or **Show as Critical**, but not for both. You must create a separate alarm definition for warning and critical status. **Keep the target's current state** option can be enabled only for an event based alarm condition. For example, if **Cannot power off the VM** and **Keep the target's current state** is selected for an alarm rule, then it cannot be combined with any other alarm rule with different severity like **Warning** or **Critical**. These alarms will not have reset rules as the alarm rule does not change the state of the target object.
6. Send email notifications
 - a) To send email notifications when alarms are triggered, enable **Send email notifications**.
 - b) In the **Email to**, enter recipient addresses. Use commas to separate multiple addresses.
7. To send traps when alarms are triggered on a vCenter Serverinstance, enable **Send SNMP traps**.
8. Run scripts
 - a) To run scripts when alarms are triggered, enable **Run script**.
 - b) In **Run this script** column, enter the full path to the script.

```
/var/myscripts/myAlarmActionScript
```

You can use environment variables to define more complex scripts and attach them to multiple alarms or inventory objects so that the alarm action occurs when the alarm triggers.

Table 530: Alarm Environment Variables

Variable Name	Variable Description	Supported Alarm Type
VMWARE_ALARM_NAME	Name of the triggered alarm.	Condition, State, Event
VMWARE_ALARM_ID	MOID of the triggered alarm.	Condition, State, Event
VMWARE_ALARM_TARGET_NAME	Name of the entity on which the alarm triggered.	Condition, State, Event
VMWARE_ALARM_TARGET_ID	MOID of the entity on which the alarm triggered.	Condition, State, Event
VMWARE_ALARM_OLDSTATUS	Old status of the alarm.	Condition, State, Event
VMWARE_ALARM_NEWSTATUS	New status of the alarm.	Condition, State, Event
VMWARE_ALARM_TRIGGERINGSUMMARY	Multiline summary of the alarm.	Condition, State, Event
VMWARE_ALARM_DECLARINGSUMMARY	Single-line declaration of the alarm expression.	Condition, State, Event
VMWARE_ALARM_ALARMVALUE	Value that triggered the alarm.	Condition, State
VMWARE_ALARM_EVENTDESCRIPTION	Description text of the alarm status change event.	Condition, State
VMWARE_ALARM_EVENTDESCRIPTION	Description of the event that triggered the alarm.	Event
VMWARE_ALARM_EVENT_USERNAME	User name associated with the event.	Event
VMWARE_ALARM_EVENT_DATACENTER	Name of the datacenter in which the event occurred.	Event
VMWARE_ALARM_EVENT_COMPUTERESOURCE	Name of the cluster or resource pool in which the event occurred.	Event

Variable Name	Variable Description	Supported Alarm Type
VMWARE_ALARM_EVENT_HOST	Name of the host on which the event occurred.	Event
VMWARE_ALARM_EVENT_VM	Name of the virtual machine on which the event occurred.	Event
VMWARE_ALARM_EVENT_NETWORK	Name of the network on which the event occurred.	Event
VMWARE_ALARM_EVENT_DATASTORE	Name of the datastore on which the event occurred.	Event
VMWARE_ALARM_EVENT_DVS	Name of the vNetwork Distributed Switch on which the event occurred.	Event

If your script does not make use of the alarm environment variables, include any necessary parameters in the configuration field. Enclose parameters in curly brackets. For example:

```
/var/myscripts/myAlarmActionScript {alarmName} {targetName}
```

Check that the script can be executed:

```
chmod +x /var/myscripts/myAlarmActionScript
```

Starting with vCenter Server 8.0, the owner of the script must have vpxuser privileges :

```
chown vpxd /var/myscripts/myAlarmActionScript
```

For more information about scripts which require root user privileges, see Knowledge Base article [KB 87918](#).

9. (Optional) Configure alarm transitions and frequency.

10. Select an advanced action from the drop-down menu.

You can define the advanced actions for virtual machine and hosts. These advanced actions are applicable only for virtual machines and hosts. There are different sets of advanced actions based on target types of virtual machines and hosts.

You can add multiple advanced actions for an alarm.

11. (Optional) Configure frequency for advanced actions.

12. Click **Add Another Rule** to add an additional rule for an alarm.

13. Click **Duplicate Rule** to create an identical rule for an alarm.

14. Click **Remove Rule** to remove an existing rule set for an alarm.

Click **Next** to set the Reset Rule.

Specify Alarm Reset Rules

You can select and configure the events, states, or conditions that trigger the alarm from the **Reset Rule** page in the **New Alarm Definition** wizard.

Required Privilege: **Alarms > Create alarm** or **Alarms > Modify alarm**

You can set the Alarm Reset Rule

1. Enable the **Reset the alarm to green** option.

2. Select a trigger from the drop-down menu.

The combined event triggers are displayed. You can set the rule for a single event only. You must create multiple rules for multiple events.

3. Click **Add Argument** to select an argument from the drop-down menu.
It supports **ALL** based expression, option to select **ANY** is not available. You must create a separate alarm definition for each trigger. The **OR** operator is not supported in the vSphere Client. However, you can combine more than one condition trigger with **AND** operator.
4. Select an operator from the drop-down menu.
5. Send email notifications
 - a) To send email notifications when alarms are triggered, enable **Send email notifications**.
 - b) In the **Email to**, enter recipient addresses. Use commas to separate multiple addresses.
6. To send traps when alarms are triggered on a vCenter Server instance, enable **Send SNMP traps**.
7. Run scripts
 - a) To run scripts when alarms are triggered, enable **Run script**.
 - b) In **Run this script** column, type script or command information:

If your script does not make use of the alarm environment variables, include any necessary parameters in the configuration field. Enclose parameters in curly brackets. For example:

```
c:\tools\cmd.exe {alarmName} {targetName}
c:\windows\system32\cmd.exe /c c:\tools\cmd.bat {alarmName} {targetName}
```


The script can run on any platform. You must provide the path to the script and argument keys. For example:

```
/var/myscripts/myAlarmActionScript {alarmName} {targetName}
```
8. (Optional) Configure alarm transitions and frequency.
9. Select an advanced action from the **Add Advanced Actions** drop-down menu.
You can add multiple advanced actions for an alarm reset rule. You can define the advanced actions for virtual machines and hosts. These advanced actions are applicable only for virtual machines and hosts. There are different sets of advanced actions based on target types of virtual machines and hosts.
You can add multiple advanced actions for an alarm.
10. (Optional) Configure frequency for advanced actions.
11. Click **Add Another Rule** to add an additional reset rule for an alarm.
12. Click **Duplicate Rule** to create an identical reset rule for an alarm.
13. Click **Remove Rule** to remove an existing reset rule set for an alarm.

Click **Next** to review the alarm definition.

Review and Enable Alarm

You can review and enable the alarm in the vSphere Client

Required Privilege: **Alarms > Create alarm** or **Alarms > Modify alarm**

After setting the alarm rule, review the alarm before enabling the alarm.

1. Review the **Alarm Name**, **Description**, **Targets**, and **Alarm Rule**.
2. (Optional) Configure alarm transitions and frequency.
3. Select **Enable this alarm** to enable the alarm.

The alarm is enabled.

Disable Alarms on Specific Objects

vSphere includes a feature to disable alarms on specific leaf object in the inventory, even if those are defined on one of its parents. This feature is helpful, if a specific alarm is not applicable for a given object.

Required Privilege: Alarms.Disable or enable alarm on entity

You can access alarms in the **Configure** tab.

Note: UI automatically enables "Disable/Enable" buttons for alarm specific leaf entities.

1. Select an inventory object, click the **Configure** tab, and click **More**.
2. Click **Alarm Definitions**.
3. Select an **Alarm**.
4. Click **Disable** to disable the alarm on that entity.

Acknowledge Triggered Alarms

After you acknowledge an alarm in the vSphere Client, its alarm actions are discontinued. Alarms are not cleared, or reset when acknowledged.

Required privilege: **Alarm > Alarm Acknowledge**

Acknowledging an alarm lets other users know that you are taking ownership of the issue. For example, a host has an alarm set to monitor CPU usage. It sends an email to an administrator when the alarm is triggered. The host CPU usage spikes, triggering the alarm which sends an email to the host's administrator. The administrator acknowledges the triggered alarm to let other administrators know the problem is being addressed, and to prevent the alarm from sending more email messages. The alarm, however, is still visible in the system.

- Right-click the alarm in the Alarms panel and select **Acknowledge**.
- Acknowledge the alarm in the **Monitor** tab.
 - a) Select an inventory object in the object navigator.
 - b) Click the **Monitor** tab.
 - c) Click **Issues and Alarms**, and click **Triggered Alarms**.
 - d) Select an alarm and select **Acknowledge**.

Reset Triggered Event Alarms

An alarm triggered by an event might not reset to a normal state if vCenter Server does not retrieve the event that identifies the normal condition. In such cases, reset the alarm manually in the vSphere Client to return it to a normal state.

Required privilege: **Alarm > Set Alarm Status**

- Right-click an alarm in the Alarms sidebar pane and select **Reset to green**.
- Reset triggered alarms in the **Monitor** tab.
 - a) Select an inventory object.
 - b) Click the **Monitor** tab.
 - c) Click **Issues and Alarms**, and click **Triggered Alarms**.
 - d) Select the alarms you want to reset.

Use Shift+left-click or Ctrl+left-click to select multiple alarms is supported in the vSphere Client.
 - e) Right-click an alarm and select **Reset to Green**.

Preconfigured vSphere Alarms

vCenter Server provides a list of default alarms, which monitor the operations of vSphere inventory objects. You must only set up actions for these alarms.

Some alarms are stateless. vCenter Server does not keep data on stateless alarms, does not compute, or display their status. Stateless alarms cannot be acknowledged or reset. Stateless alarms are indicated by an asterisk next to their name.

Note:

This list of default vSphere Alarms is not exhaustive. There might be more preconfigured vSphere alarms depending on your configuration, for example, the services added to your environment.

Table 531: Default vSphere Alarms

Alarm Name	Description
Host connection and power state	Monitors the power state of the host and whether the host is reachable.
Host CPU usage	Monitors host CPU usage.
Host memory usage	Monitors host memory usage.
Virtual machine CPU usage	Monitors virtual machine CPU usage.
Virtual machine memory usage	Monitors virtual machine memory usage.
Datastore usage on disk	Monitors datastore disk usage. Note: This alarm controls the Status value for datastores in vSphere Client. If you deactivate this alarm, the datastore status is displayed as Unknown.
Virtual machine CPU ready	Monitors virtual machine CPU ready time.
Virtual machine total disk latency	Monitors virtual machine total disk latency.
Virtual machine disk commands canceled	Monitors the number of virtual machine disk commands that are canceled.
Virtual machine disk reset	Monitors the number of virtual machine bus resets.
License inventory monitoring	Monitors the license inventory for compliance.
License user threshold monitoring	Monitors whether a user-defined license threshold is exceeded.
License capacity monitoring	Monitors whether a license capacity is exceeded.
The host license edition is not compatible with the vCenter Server license edition	Monitors the compatibility of the vCenter Server and host license editions.
Host flash capacity exceeds the licensed limit for vSAN	Monitors whether the flash disk capacity on the host exceeds the limit of the vSAN license.
Expired vSAN license	Monitors the expiry of the vSAN license and the end of the evaluation period.
Errors occurred on the disk(s) of a vSAN host	Default alarm that monitors whether the host disks in the vSAN cluster have errors.
Timed out starting Secondary VM *	Monitors whether starting a secondary virtual machine has timed out.
No compatible host for Secondary VM	Monitors the availability of compatible hosts on which a secondary virtual machine can be created and run.
Virtual machine Fault Tolerance state changed	Monitors changes in the Fault Tolerance state of a virtual machine.
Virtual Machine Fault Tolerance vLockStep interval Status Changed	Monitors changes in the Fault Tolerance Secondary vLockStep interval.

Alarm Name	Description
Host processor status	Monitors the host processors.
Host memory status	Monitors host memory usage.
Host hardware fan status	Monitors host fans.
Host hardware voltage	Monitors host hardware voltage.
Host hardware temperature status	Monitors the temperature status of the host hardware.
Host hardware power status	Monitors the host power status.
Host hardware system board status	Monitors the status of host system boards.
Host battery status	Monitors the battery status of hosts.
Status of other host hardware objects	Monitors other host hardware objects.
Host storage status	Monitors host connectivity to storage devices.
Host IPMI System Event Log status	Monitors the capacity of the IPMI system event log.
Host Baseboard Management Controller status	Monitors the status of the Baseboard Management Controller.
Host error *	Monitors host error and warning events.
Virtual machine error *	Monitors virtual machine error and warning events.
Host connection failure *	Monitors host connection failures.
Unmanaged workload detected on SIOC-enabled datastore	Monitors the unmanaged I/O workload on a SIOC-enabled datastore.
Thin-provisioned volume capacity threshold exceeded	Monitors whether the thin provisioning threshold on the storage array exceeds for volumes backing the datastore.
Datastore capability alarm	Monitors the change in the capability status for volumes backing the datastore.
VASA provider disconnected	Monitors the changes in the connection state of VASA providers.
VASA Provider certificate expiration alarm	Monitors whether VASA provider certificates are getting close to their expiry date.
VM storage compliance alarm	Monitors the virtual disk compliance with the object-based storage.
Datastore compliance alarm	Monitors whether the virtual disk on the datastore complies with the object-based storage.
Refreshing CA certificates and CRLs for a VASA provider failed	Monitors whether the refreshing of CA certificates and CRLs for some of the VASA providers has failed.
Insufficient vSphere HA failover resources	Monitors the sufficiency of failover cluster resources required for vSphere High Availability.
vSphere HA failover in progress	Monitors the failover progress of vSphere High Availability.
Cannot find vSphere HA primary agent	Monitors whether vCenter Server is able to connect to a vSphere High Availability primary agent.
vSphere HA host status	Monitors the host health status reported by vSphere High Availability.
vSphere HA virtual machine failover failed	Monitors whether a failover operation that uses vSphere High Availability failed.
vSphere HA virtual machine monitoring action	Monitors whether vSphere High Availability has restarted a virtual machine.

Alarm Name	Description
vSphere HA virtual machine monitoring error	Monitors whether vSphere High Availability failed to reset a virtual machine.
vSphere HA VM Component Protection could not power off a virtual machine	Monitors whether vSphere High Availability VM Component Protection cannot power off a virtual machine with an inaccessible datastore.
License error *	Monitors license errors.
Health status changed *	Monitors changes to service and extension health status.
Virtual machine component protection restart error	Monitors whether the vSphere HA VM Component Protection fails to restart a virtual machine.
Storage DRS recommendation	Monitors Storage DRS recommendations.
Storage DRS is not supported on a host	Monitors and alerts when Storage DRS is not supported on a host.
Datastore cluster is out of space	Monitors whether a datastore cluster runs out of disk space.
Datastore is in multiple datacenters	Monitors whether a datastore in a datastore cluster is visible in more than one data center.
vSphere Distributed Switch VLAN trunked status	Monitors changes in vSphere Distributed Switch VLAN trunked status.
vSphere Distributed Switch MTU matched status	Monitors changes in vSphere Distributed Switch MTU matched status.
vSphere Distributed Switch MTU supported status	Monitors changes in vSphere Distributed Switch MTU supported status.
vSphere Distributed Switch teaming matched status	Monitors changes in vSphere Distributed Switch teaming matched status.
Virtual Machine network adapter reservation status	Monitors changes in the reservation status of a virtual machine network adapter.
Virtual machine Consolidation Needed status	Monitors changes in the virtual machine Consolidation Needed status.
Host virtual flash resource status	Monitors the Flash Read Cache resource status on the host.
Host virtual flash resource usage	Monitors the Flash Read Cache resource usage on the host.
Registration/unregistration of a VASA vendor provider on a vSAN host fails	Default alarm that monitors whether the registration or unregistration of a VASA vendor provider on a vSAN host fails.
Registration/unregistration of third-party IO filter storage providers fails on a host	Default alarm that monitors whether vCenter Server fails to register or unregister third-party IO filter storage providers on a host.
Service Control Agent Health Alarm	Monitors the health status of the VMware Service Control Agent.
Identity Health Alarm	Monitors the health status of the Identity Management Service.
vSphere Client Health Alarm	Monitors the health status of the vSphere Client.
ESX Agent Manager Health Alarm	Monitors the health status of the ESX Agent Manager.
Message Bus Config Health Alarm	Monitors the health status of the Message Bus Configuration Service.
Cis License Health Alarm	Monitors the health status of the License Service.
vCenter Server Health Alarm	Monitors the health status of vCenter Server.

Alarm Name	Description
Database Health Alarm	Monitors the database health status. When database space reaches 80%, vCenter Server displays a warning event. When database space reaches 95%, vCenter Server displays an error event and shuts down. You can clean up the database, or increase database storage capacity, and start vCenter Server. The alarm is triggered only for PostgreSQL and Microsoft SQL Server database health issues, and does not work with Oracle databases.
Data Service Health Alarm	Monitors the health status of the Data Service.
RBD Health Alarm	Monitors the health status of the vSphere Auto Deploy Waiter.
vService Manager Health Alarm	Monitors the health status of the vService Manager.
Performance Charts Service Health Alarm	Monitors the health status of the Performance Charts Service.
Content Library Service Health Alarm	Monitors the health status of the VMware Content Library Service.
VMware vSphere ESXi Dump Collector Health Alarm	Monitors the health status of the VMware vSphere ESXi Dump Collector Service.
VMware vAPI Endpoint Service Health Alarm	Monitors the health status of the VMware vAPI Endpoint Service.
VMware vSphere Profile-Driven Storage Service Health Alarm	Monitors the health status of the VMware vSphere Profile-Driven Storage Service.
VMware vFabric Postgres Service Health Alarm	Monitors the health status of the VMware vFabric Postgres Service.
ESXi Host Certificates Update Failure Status	Monitors whether the update of the ESXi host certificates failed.
ESXi Host Certificate Status	Monitors the certificate status of an ESXi host.
ESXi Host Certificate Verification Failure Status	Monitors whether the verification of an ESXi host certificate failed.
vSphere vCenter Host Certificate Management Mode	Monitors changes in the certificate management mode of vCenter Server.
Root Certificate Status	Monitors whether a root certificate is getting close to its expiration date.
GPU ECC Uncorrected Memory Alarm	Monitors the GPU ECC uncorrected memory status.
GPU ECC Corrected Memory Alarm	Monitors the GPU ECC corrected memory status.
GPU Thermal Condition Alarm	Monitors the GPU Thermal condition status.
Network connectivity lost	Monitors the network connectivity on a virtual switch.
Network uplink redundancy lost	Monitors network uplink redundancy on a virtual switch.
Network uplink redundancy degraded *	Monitors network uplink redundancy degradation on a virtual switch.
VMKernel NIC not configured correctly *	Monitors incorrectly configured VMkernel NICs.
Cannot connect to storage *	Monitors host connectivity to a storage device.
Migration error *	Monitors whether a virtual machine cannot be migrated or relocated, or is orphaned.
Exit standby error	Monitors whether a host cannot exit standby mode.

View Alarm Definitions for Preconfigured vSphere Alarms

You can view the definitions of all preconfigured vSphere alarms and all alarms that you configured for each inventory object.

When a preconfigured alarm is triggered, and you want to see what rules triggered the alarm, you can view the alarm definition.

1. Select an object from the inventory, and click the **Configure** tab.
2. In the left navigation pane, click **Alarm Definitions**.
A list with all configured alarms for the selected object appears.
3. Optional: Use the filters above each column to quickly find the alarm definition.
4. Click the arrow to the left of the alarm name.

You can view all information related to the configuration of the alarm, including the description of the alarm, the target object, and the alarm rules which trigger the alarm.

Monitoring Solutions with the vCenter Solutions Manager

A solution is an extension of the vCenter Server that adds new functions to the vCenter Server instance. You can monitor the detailed information of an inventory and the health status of the solutions.

VMware products that integrate with vCenter Server are also considered solutions. For example, vSphere ESX Agent Manager is a VMware solution that lets you manage host agents that add new capabilities to ESX and ESXi hosts.

You can also install a solution to add functionality from third-party technologies to the standard functions of vCenter Server. Solutions are typically delivered as OVF packages. You can install and deploy solutions from the vSphere Client. You can integrate the solutions into the vCenter Solutions Manager, which provides a list of all installed solutions.

If a virtual machine or vApp runs a solution, a custom icon represents it in the inventory of the vSphere Client. Each solution registers a unique icon to show that the solution manages the virtual machine or vApp. The icons show the power states (powered on, paused, or powered off). Solutions display more than one type of icon if they manage more than one type of virtual machine or vApp.

When you power on or power off a virtual machine or vApp, you are notified that you are operating with an object that is managed by the vCenter Solutions Manager. When you attempt another operation on a virtual machine or a vApp that is managed by a solution, an informational warning message appears.

For more information, see the *Developing and Deploying vSphere Solutions, vServices, and ESX Agents* documentation.

View Solutions

vCenter Solutions Manager helps to deploy, monitor, and interact with solutions that are installed in a vCenter Server instance.

The vCenter Solutions Manager displays information about the solution, such as the solution name, the vendor name, and the version of the product. The vCenter Solutions manager also displays information about the health of a solution.

1. Navigate to the vCenter Solutions Manager.
 - a) In the vSphere Client, select **Menu > Administration**.
 - b) Expand **Solutions** and click **vCenter Server Extensions**.
2. Click a solution from the list.
For instance, vService Manager or vSphere ESX Agent Manager.

3. Navigate through the tabs to view information about the solution.

- **Summary:** You can view details about the solution, such as the product name, a short description, and links to the product and vendor websites. You can also view the solution configuration and the solution UI.
Select the vCenter Server link to view the **Summary** page of the Virtual Machines or vApp.
- **Monitor:** You can view the tasks and events related to the solution.
- **VMs:** You can view a list of all the Virtual Machines and vApps that belong to the solution.

Monitoring the Health of Services and Nodes

You can monitor the health status of services and nodes to determine whether problems exist in your environment.

The vSphere Client provides an overview of all services and nodes across the management stack of the vCenter Server system. A list of default services is available for each vCenter Server instance.

View the Health Status of Nodes

In the vSphere Client, you can view the health status of vCenter Server nodes.

Verify that the user you use to log in to the vCenter Server instance is a member of the SystemConfiguration.Administrators group in the vCenter Single Sign-On domain.

vCenter Server instances and machines that run vCenter Server services are considered nodes. Graphical badges represent the health status of nodes.

1. Use the vSphere Client to log in as `administrator@your_domain_name` to the vCenter Server instance.
The address is of the type `http://appliance-IP-address-or-FQDN/ui`.
2. From the vSphere Client menu, select **Administration**.
3. Select **Deployment > System Configuration**.
4. Select a node to view its health status.

Table 532: Health States

Badge Icon	Description
	Good. The health of the object is normal.
	Warning. The object is experiencing some problems.
	Critical. The object is either not functioning properly or will stop functioning soon.
	Unknown. No data is available for this object.

View the Health Status of Services

You can view the health status of various vCenter Services using the vCenter Server Management Interface

- Verify that the vCenter Server is successfully deployed and running.

- If you are using Internet Explorer, verify that TLS 1.0, TLS 1.1, and TLS 1.2 are enabled in the security settings.

Note: The login session expires if you leave the vCenter Server Management Interface idle for 10 minutes.

1. In a Web browser, go to the vCenter Server Management Interface, <https://appliance-IP-address-or-FQDN:5480>.
2. Log in as root.
The default root password is the password that you set while deploying vCenter Server.
3. In the vCenter Server Management Interface, click **Services**.
4. You can view the health of all the listed services in the **Health** column.

Performance Monitoring Utilities: `resxstop` and `esxstop`

The `resxstop` and `esxstop` command-line utilities provide a detailed look at how ESXi uses resources in real time. You can start either utility in one of three modes: interactive (default), batch, or replay.

The fundamental difference between `resxstop` and `esxstop` is that you can use `resxstop` remotely, whereas you can start `esxstop` only through the ESXi Shell of a local ESXi host.

Using the `esxstop` Utility

You can run the `esxstop` utility using the ESXi Shell to communicate with the management interface of the ESXi host. You must have root user privileges.

`esxstop` Utility

To use the `esxstop` utility, type the following command with desired options:

```
esxstop [-h] [-v] [-b] [-l] [-s] [-a] [-c config file] [-R vm-support_dir_path] [-d delay] [-n iterations] [-export-entity entity-file] [-import-entity entity-file]
```

Command-Line Help Option	Description
<code>-h</code>	Prints this help menu.
<code>-v</code>	Prints version.
<code>-b</code>	Enables batch mode.
<code>-l</code>	Locks the <code>esxstop</code> objects to those available in the first snapshot.
<code>-s</code>	Enables secure mode.
<code>-a</code>	Displays all statistics.
<code>-c</code>	Sets the <code>esxstop</code> configuration file, which by default is <code>.esxstop60.rc</code> .
<code>-R</code>	Enables replay mode.
<code>-d</code>	Sets the delay between updates in seconds.
<code>-n</code>	Runs <code>esxstop</code> for only <code>n</code> iterations. Use <code>-n infinity</code> to run <code>esxstop</code> for infinite period of time.
<code>-u</code>	Suppresses server-wide physical CPU statistics.

For example:

```
[root@localhost:~] esxstop
```

The above command will open the interactive screen of `esxstop`, by default.

```
[root@localhost:~] esxtop -b -a -d 10 -n 3 > /vmfs/volumes/localhost/test.csv
```

The above command runs the batch mode of `esxtop` with all counters, updates delay of 10 secs for 3 iterations. The output is written to the `test.csv` file which can be opened with other tools like Windows Perfmon.

esxtop Configuration

The `esxtop` utility reads its default configuration from `.esxtop50rc` on the ESXi system. This configuration file consists of nine lines.

The first eight lines contain lowercase and uppercase letters to specify which fields appear in which order on the CPU, memory, storage adapter, storage device, virtual machine storage, network, interrupt, and CPU power panels. The letters correspond to the letters in the Fields or Order panels for the respective `esxtop` panel.

The ninth line contains information on the other options. Most important, if you saved a configuration in secure mode, you do not get an insecure `esxtop` without removing the `s` from the seventh line of your `.esxtop50rc` file. A number specifies the delay time between updates. As in interactive mode, typing `c`, `m`, `d`, `u`, `v`, `n`, `I`, or `p` determines the panel with which `esxtop` starts.

Note: Do not edit the `.esxtop50rc` file. Instead, select the fields and the order in a running `esxtop` process, make changes, and save this file using the `w` interactive command.

Related Links

[Interactive Mode Command-Line Options on page 3440](#)

You can use various command-line options with `esxtop` and `resxtop` in interactive mode.

Using the resxtop Utility

`resxtop` is a command-line utility or tool that runs on Linux and provides a detailed look at how ESXi uses resources in real time.

Before you can use any `resxtop` commands, you must download and install `resxtop` to your system.

Note: `resxtop` is supported only on Linux.

After it is set up, start `resxtop` from the command line. For remote connections, you can connect to a host either directly or through vCenter Server.

To launch `resxtop` and connect to a remote server, enter this command

```
resxtop --server <hostname> --username <user>
```

The command-line options listed in the following table are the same as for `esxtop` (except for the `R` option) with additional connection options.

Note: `resxtop` does not use all the options shared by other ESXCLI commands.

Table 533: resxtop Command-Line Options

Option	Description
[server]	Name of the remote host to connect to (required). If connecting directly to the ESXi host, use the name of that host. If your connection to the ESXi host is indirect (that is, through vCenter Server), use the name of the vCenter Server system for this option.
[vihost]	If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the host, this option is not used. Note that the host name needs to be the same as what appears in the vSphere Client.

Option	Description
[portnumber]	Port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed.
[username]	User name to be authenticated when connecting to the remote host. The remote server prompts you for a password.

You can also use `resxtp` on a local ESXi host by omitting the `server` option on the command line. The command defaults to `localhost`.

Installing resxtp on Ubuntu Linux

Installing RESXTOP on Linux systems is a simple process that involves extracting the files from the archive, running and install script, and updating the path.

1. Download `resxtp` from <https://developer.broadcom.com/tools/vmware-vsphere-resxtp/latest>.
2. Unpack the archive using the linux tar command `"tar -xvf resxtp-7.0.0-15992393-lin64.tgz"`.
3. Execute the install script by running the linux command `"./install.sh"`.
4. Update the path by running the linux command `"export LD_LIBRARY_PATH=/usr/lib/vmware/resxtp"`.

Related Links

[Interactive Mode Command-Line Options on page 3440](#)

You can use various command-line options with `esxtp` and `resxtp` in interactive mode.

Using esxtp or resxtp in Interactive Mode

By default, `resxtp` and `esxtp` run in interactive mode. Interactive mode displays statistics in different panels.

A help menu is available for each panel.

Interactive Mode Command-Line Options

You can use various command-line options with `esxtp` and `resxtp` in interactive mode.

Table 534: Interactive Mode Command-Line Options

Option	Description
<code>h</code>	Prints help for <code>resxtp</code> (or <code>esxtp</code>) command-line options.
<code>v</code>	Prints <code>resxtp</code> (or <code>esxtp</code>) version number.
<code>s</code>	Calls <code>resxtp</code> (or <code>esxtp</code>) in secure mode. In secure mode, the <code>-d</code> command, which specifies delay between updates, is disabled.
<code>d</code>	Specifies the delay between updates. The default is five seconds. The minimum is two seconds. Change this with the interactive command <code>s</code> . If you specify a delay of less than two seconds, the delay is set to two seconds.
<code>n</code>	Number of iterations. Updates the display <code>n</code> times and exits. Default value is 10000.
<code>server</code>	The name of the remote server host to connect to (required for <code>resxtp</code> only).
<code>vihost</code>	If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the ESXi host, this option is not used. Note that the host name needs to be the same as what is displayed in the vSphere Client.

Option	Description
<code>portnumber</code>	The port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. (<code>resxstop</code> only)
<code>username</code>	The user name to be authenticated when connecting to the remote host. The remote server prompts you for a password, as well (<code>resxstop</code> only).
<code>a</code>	Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default <code>~/esxstop50rc</code> configuration file or a user-defined configuration file.
<code>cfilename</code>	Load a user-defined configuration file. If the <code>-c</code> option is not used, the default configuration filename is <code>~/esxstop50rc</code> . Create your own configuration file, specifying a different filename, using the <code>w</code> single-key interactive command.

Common Statistics Description

Several statistics appear on the different panels while `resxstop` (or `esxstop`) is running in interactive mode. These statistics are common across all four panels.

The Uptime line, found at the top of each of the four `resxstop` (or `esxstop`) panels, displays the current time, time since last reboot, number of currently running worlds and load averages. A world is an ESXi VMkernel schedulable entity, similar to a process or thread in other operating systems.

Below that the load averages over the past one, five, and fifteen minutes appear. Load averages consider both running and ready-to-run worlds. A load average of 1.00 means that there is full utilization of all physical CPUs. A load average of 2.00 means that the ESXi system might need twice as many physical CPUs as are currently available. Similarly, a load average of 0.50 means that the physical CPUs on the ESXi system are half utilized.

Statistics Columns and Order Pages

You can define the order of fields displayed in interactive mode.

If you press `f`, `F`, `o`, or `O`, the system displays a page that specifies the field order on the top line and short descriptions of the field contents. If the letter in the field string corresponding to a field is uppercase, the field is displayed. An asterisk in front of the field description indicates whether a field is displayed.

The order of the fields corresponds to the order of the letters in the string.

From the Field Select panel, you can:

- Toggle the display of a field by pressing the corresponding letter.
- Move a field to the left by pressing the corresponding uppercase letter.
- Move a field to the right by pressing the corresponding lowercase letter.

Interactive Mode Single-Key Commands

When running in interactive mode, `resxstop` (or `esxstop`) recognizes several single-key commands.

All interactive mode panels recognize the commands listed in the following table. The command to specify the delay between updates is disabled if the `s` option is given on the command line. All sorting interactive commands sort in descending order.

Table 535: Interactive Mode Single-Key Commands

Key	Description
h or ?	Displays a help menu for the current panel, giving a brief summary of commands, and the status of secure mode.
space	Immediately updates the current panel.
^L	Erases and redraws the current panel.
f or F	Displays a panel for adding or removing statistics columns (text boxes) to or from the current panel.
o or O	Displays a panel for changing the order of statistics columns on the current panel.
#	Prompts you for the number of statistics rows to display. Any value greater than 0 overrides automatic determination of the number of rows to show, which is based on window size measurement. If you change this number in one <code>resxtop</code> (or <code>esxtop</code>) panel, the change affects all four panels.
s	Prompts you for the delay between updates, in seconds. Fractional values are recognized down to microseconds. The default value is five seconds. The minimum value is two seconds. This command is not available in secure mode.
W	Write the current setup to an <code>esxtop</code> (or <code>resxtop</code>) configuration file. This is the recommended way to write a configuration file. The default filename is the one specified by <code>-c</code> option, or <code>~/esxtop50rc</code> if the <code>-c</code> option is not used. You can also specify a different filename on the prompt generated by this W command.
q	Quit the interactive mode.
c	Switch to the CPU resource utilization panel.
p	Switch to the CPU Power utilization panel.
m	Switch to the memory resource utilization panel.
d	Switch to the storage (disk) adapter resource utilization panel.
u	Switch to storage (disk) device resource utilization screen.
v	Switch to storage (disk) virtual machine resource utilization screen.
n	Switch to the network resource utilization panel.
i	Switch to the interrupt panel.

CPU Panel

The CPU panel displays server-wide statistics as well as statistics for the individual world, resource pool, and virtual machine CPU utilization.

Resource pools, virtual machines that are running, or other worlds are at times called groups. For worlds belonging to a virtual machine, statistics for the virtual machine that is running are displayed. All other worlds are logically aggregated into the resource pools that contain them.

Table 536: CPU Panel Statistics

Line	Description
PCPU USED(%)	A PCPU refers to a physical hardware execution context. It can be a physical CPU core if the hyperthreading is unavailable or deactivated, or a logical CPU (LCPU or SMT thread) if the hyperthreading is activated. PCPU USED(%) displays the following percentages: - percentage of CPU usage per PCPU - percentage of CPU usage averaged over all PCPUs

Line	Description
	CPU Usage (%USED) is the percentage of PCPU nominal frequency that was used since the last screen update. It equals the total sum of %USED for Worlds that ran on this PCPU. Note: If a PCPU is running at frequency that is higher than its nominal (rated) frequency, then PCPU USED(%) can be greater than 100%. If a PCPU and its partner are busy when hyperthreading is activated, each PCPU account for half of the CPU usage.
PCPU UTIL(%)	A PCPU refers to a physical hardware execution context. It can be a physical CPU core if the hyperthreading is unavailable or deactivated, or a logical CPU (LCPU or SMT thread) if the hyperthreading is activated. PCPU UTIL(%) represents the percentage of real time that the PCPU was not idle (raw PCPU utilization). It displays the percentage CPU utilization per PCPU, and the percentage CPU utilization averaged over all PCPUs. Note: PCPU UTIL(%) might differ from PCPU USED(%) due to power management technologies or hyperthreading.
ID	Resource pool ID or virtual machine ID of the resource pool or virtual machine of the world that is running. Alternatively, the world ID of the world that is running.
GID	Resource pool ID of the resource pool or virtual machine of the world that is running.
NAME	Name of the resource pool or virtual machine of the world that is running, or name of the world that is running.
NWLD	Number of members in the resource pool or virtual machine of the world that is running. If a Group is expanded using the interactive command <code>e</code> , then NWLD for all the resulting worlds is 1.
%STATE TIMES	Set of CPU statistics made up of the following percentages. For a world, the percentages are a percentage of one physical CPU core.
%USED	Percentage of physical CPU core cycles used by the resource pool, virtual machine, or world. %USED might depend on the frequency with which the CPU core is running. When running with lower CPU core frequency, %USED can be smaller than %RUN. On CPUs which support the turbo mode, CPU frequency can also be higher than the nominal (rated) frequency, and %USED can be larger than %RUN. $\%USED = \%RUN + \%SYS - \%OVRP$
%SYS	Percentage of time spent in the ESXi VMkernel on behalf of the resource pool, virtual machine, or world to process interrupts and to perform other system activities. This time is part of the time used to calculate %USED. $\%USED = \%RUN + \%SYS - \%OVRP$
%WAIT	Percentage of time the resource pool, virtual machine, or world spent in the blocked or busy wait state. This percentage includes the percentage of time the resource pool, virtual machine, or world was idle. $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
%VMWAIT	The total percentage of time the Resource Pool/World spent in a blocked state waiting for events.
%IDLE	Percentage of time the resource pool, virtual machine, or world was idle. Subtract this percentage from %WAIT to see the percentage of time the resource pool, virtual machine, or world was waiting for some event. The difference, %WAIT - %IDLE, of the VCPU worlds can be used to estimate guest I/O wait time. To find the VCPU worlds, use the single-key command <code>e</code> to expand a virtual machine and search for the world NAME starting with "vcpu". (The VCPU worlds might wait for other events in addition to I/O events, so this measurement is only an estimate.)
%RDY	Percentage of time the resource pool, virtual machine, or world was ready to run, but was not provided CPU resources on which to execute. $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
%MLMTD (max limited)	Percentage of time the ESXi VMkernel deliberately did not run the resource pool, virtual machine, or world because doing so would violate the resource pool, virtual machine, or world's limit setting. Because the resource pool, virtual machine, or world is ready to run when it is prevented from running in this way, the %MLMTD (max limited) time is included in %RDY time.

Line	Description
%SWPWT	Percentage of time a resource pool or world spends waiting for the ESXi VMkernel to swap memory. The %SWPWT (swap wait) time is included in the %WAIT time.
EVENT COUNTS/s	Set of CPU statistics made up of per second event rates. These statistics are for VMware internal use only.
CPU ALLOC	Set of CPU statistics made up of the following CPU allocation configuration parameters.
AMIN	Resource pool, virtual machine, or world attribute Reservation.
AMAX	Resource pool, virtual machine, or world attribute Limit. A value of -1 means unlimited.
ASHRS	Resource pool, virtual machine, or world attribute Shares.
SUMMARY STATS	Set of CPU statistics made up of the following CPU configuration parameters and statistics. These statistics apply only to worlds and not to virtual machines or resource pools.
AFFINITY BIT MASK	Bit mask showing the current scheduling affinity for the world.
HTSHARING	Current hyperthreading configuration.
CPU	The physical or logical processor on which the world was running when <code>resxstop</code> (or <code>esxstop</code>) obtained this information.
HTQ	Indicates whether the world is quarantined or not. N means no and Y means yes.
TIMER/s	Timer rate for this world.
%OVRP	Percentage of system time spent during scheduling of a resource pool, virtual machine, or world on behalf of a different resource pool, virtual machine, or world while the resource pool, virtual machine, or world was scheduled. This time is not included in %SYS. For example, if virtual machine A is being scheduled and a network packet for virtual machine B is processed by the ESXi VMkernel, the time spent appears as %OVRP for virtual machine A and %SYS for virtual machine B. $\%USED = \%RUN + \%SYS - \%OVRP$
%RUN	Percentage of total time scheduled. This time does not account for hyperthreading and system time. On a hyperthreading activated server, the %RUN can be twice as large as %USED. $\%USED = \%RUN + \%SYS - \%OVRP$ $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
%CSTP	Percentage of time a resource pool spends in a ready, co-deschedule state. Note: You might see this statistic displayed, but it is intended for VMware use only. $100\% = \%RUN + \%RDY + \%CSTP + \%WAIT$
POWER	Current CPU power consumption for a resource pool (in Watts).
%LAT_C	Percentage of time the resource pool or world was ready to run but was not scheduled to run because of CPU resource contention.
%LAT_M	Percentage of time the resource pool or world was ready to run but was not scheduled to run because of memory resource contention.
%DMD	CPU demand in percentage. It represents the average active CPU load in the past minute.
CORE UTIL(%)	Percentage of CPU cycles per core when at least one of the PCPUs in this core is unhalted, and its average over all cores. This statistic only appears when hyperthreading is activated. In batch mode, the corresponding CORE UTIL(%) statistic is displayed for each PCPU. For example, PCPU 0 and PCPU 1 have the same the CORE UTIL(%) number, and that is the number for core 0.

You can change the display using single-key commands.

Table 537: CPU Panel Single-Key Commands

Command	Description
e	Toggles whether CPU statistics are displayed expanded or unexpanded. The expanded display includes CPU resource utilization statistics broken down by individual worlds belonging to a resource pool or virtual machine. All percentages for the individual worlds are percentage of a single physical CPU. Consider these examples: • If the %Used by a resource pool is 30% on a two-way server, the resource pool is utilizing 30 percent of one physical core. • If the %Used by a world belonging to a resource pool is 30 percent on a two-way server, that world is utilizing 30% of one physical core.
U	Sorts resource pools, virtual machines, and worlds by the resource pool's or virtual machine's %Used column. This is the default sort order.
R	Sorts resource pools, virtual machines, and worlds by the resource pool's or virtual machine's %RDY column.
N	Sorts resource pools, virtual machines, and worlds by the GID column.
V	Displays virtual machine instances only.
L	Changes the displayed length of the NAME column.

CPU Power Panel

The CPU Power panel displays CPU Power utilization statistics.

On the CPU Power panel, statistics are arranged per physical CPU. A physical CPU is a physical hardware execution context. It is the physical CPU core when hyper-threading is unavailable or deactivated, or a logical CPU (LCPU or SMT thread) when hyper-threading is activated.

Table 538: CPU Power Panel Statistics

Line	Description
Power Usage	Current total power usage (in Watts).
Power Cap	Total power cap (in Watts).
PSTATE MHZ	Clock frequency per state.
%USED	Percentage of physical CPU nominal frequency used since the last screen update. It is the same as PCPU USED(%) shown in the CPU Screen.
%UTIL	Raw physical CPU utilization is the percentage of time that physical CPU was not idle. It is the same as PCPU UTIL(%) shown in the CPU Screen.
%Cx	Percentage of time the physical CPU spent in C-State 'x'.
%Px	Percentage of time the physical CPU spent in P-State 'x'. On systems with Processor Clocking Control, P-states are not directly visible to ESXi. The <code>esxtop</code> shows the percentage of time spent at full speed under the heading 'P0' and the percentage of time spent at any lower speed under 'P1'.
%Tx	Percentage of time the physical CPU spent in T-State 'x'.
%A/MPERF	aperf and mperf are two hardware registers used to keep track of the actual frequency and nominal frequency of the processor. Displays the real-time aperf to mperf ratio in the last esxtop update period. $\%A/MPERF \times \text{nominal frequency of the processor} = \text{current frequency of the processor}$

Memory Panel

The Memory panel displays server-wide and group memory utilization statistics. As on the CPU panel, groups correspond to resource pools, running virtual machines, or other worlds that are consuming memory.

The first line, found at the top of the Memory panel displays the current time, time since last reboot, number of currently running worlds, and memory overcommitment averages. The memory overcommitment averages over the past one, five, and fifteen minutes appear. Memory overcommitment of 1.00 means a memory overcommitment of 100 percent.

Table 539: Memory Panel Statistics

Field	Description
PMEM (MB)	Displays the machine memory statistics for the server. All numbers are in megabytes. total Total amount of the machine memory in the server. vmk Amount of the machine memory being used by the ESXi VMkernel. other Amount of the machine memory being used by everything other than the ESXi VMkernel. free Amount of the machine memory that is free.
VMKMEM (MB)	Displays the machine memory statistics for the ESXi VMkernel. All numbers are in megabytes. managed Total amount of the machine memory managed by the ESXi VMkernel. min free Minimum amount of the machine memory that the ESXi VMkernel aims to keep free. rsvd Total amount of the machine memory currently reserved by resource pools. ursvd Total amount of the machine memory currently unreserved. state Current machine memory availability state. Possible values are high, soft, hard and low. High means that the machine memory is not under any pressure and low means that it is.
NUMA (MB)	Displays the ESXi NUMA statistics. This line appears only if the ESXi host is running on a NUMA server. All numbers are in megabytes. For each NUMA node in the server, two statistics are displayed: - The total amount of machine memory in the NUMA node that is managed by ESXi. - The amount of machine memory in the node that is currently free (in parentheses). Shared memory for the ESXi host might be larger than the total amount of memory if memory is over-committed.
PSHARE (MB)	Displays the ESXi page-sharing statistics. All numbers are in megabytes. shared Amount of the physical memory that is being shared. common Amount of the machine memory that is common across worlds. saving Amount of the machine memory that is saved because of page sharing.

Field	Description
	shared = common + saving
SWAP (MB)	Displays the ESXi swap usage statistics. All numbers are in megabytes. curr Current swap usage. rcldmtgt Where the ESXi system expects the reclaimed memory to be. Memory can be reclaimed by swapping or compression. r/s Rate at which the memory is swapped in by the ESXi system from disk. w/s Rate at which the memory is swapped to disk by the ESXi system.
ZIP (MB)	Displays the ESXi memory compression statistics. All numbers are in megabytes. zipped Total compressed physical memory. saved Saved memory by compression.
MEMCTL (MB)	Displays the memory balloon statistics. All numbers are in megabytes. curr Total amount of the physical memory reclaimed using the <code>vmmemctl</code> module. target Total amount of physical memory the ESXi host attempts to reclaim using the <code>vmmemctl</code> module. max Maximum amount of the physical memory the ESXi host can reclaim using the <code>vmmemctl</code> module.
AMIN	Memory reservation for this resource pool or virtual machine.
AMAX	Memory limit for this resource pool or virtual machine. A value of -1 means Unlimited.
ASHRS	Memory shares for this resource pool or virtual machine.
NHN	Current home node for the resource pool or virtual machine. This statistic is applicable only on NUMA systems. If the virtual machine has no home node, a dash (-) appears.
NRMEM (MB)	Current amount of remote memory allocated to the virtual machine or resource pool. This statistic is applicable only on NUMA systems.
N% L	Current percentage of memory allocated to the virtual machine or resource pool that is local.
MEMSZ (MB)	Amount of physical memory allocated to a resource pool or virtual machine. The values are the same for the VMM and VMX groups. MEMSZ = GRANT + MCTLSZ + SWCUR + "never touched"
GRANT (MB)	Amount of guest physical memory mapped to a resource pool or virtual machine. The consumed host machine memory is equal to GRANT - SHRDSVD. The values are the same for the VMM and VMX groups.
CNSM	Amount of the memory currently consumed by the virtual machine. The memory currently consumed by the virtual machine is equal to the amount of memory that the VM guest operating system currently uses, excluding the amount of memory saved for sharing if memory sharing is enabled on the VM, excluding the amount of memory saved if some of the VM memory is compressed. For more information on memory sharing and memory compression, see the <i>vSphere Resource Management</i> documentation.
SZTGT (MB)	Amount of machine memory the ESXi VMkernel wants to allocate to a resource pool or virtual machine. The values are the same for the VMM and VMX groups.

Field	Description
TCHD (MB)	Working set estimate for the resource pool or virtual machine. The values are the same for the VMM and VMX groups.
%ACTV	Percentage of guest physical memory that is being referenced by the guest. This is an instantaneous value.
%ACTVS	Percentage of guest physical memory that is being referenced by the guest. This is a slow moving average.
%ACTVF	Percentage of guest physical memory that is being referenced by the guest. This is a fast moving average.
%ACTVN	Percentage of guest physical memory that is being referenced by the guest. This is an estimation. (You might see this statistic displayed, but it is intended for VMware use only.)
MCTL?	Memory balloon driver is installed or not. N means no, Y means yes.
MCTLSZ (MB)	Amount of physical memory reclaimed from the resource pool by way of ballooning.
MCTLTGT (MB)	Amount of physical memory the ESXi system attempts to reclaim from the resource pool or virtual machine by way of ballooning.
MCTLMAX (MB)	Maximum amount of physical memory the ESXi system can reclaim from the resource pool or virtual machine by way of ballooning. This maximum depends on the guest operating system type.
SWCUR (MB)	Current swap usage by this resource pool or virtual machine.
SWTGT (MB)	Target where the ESXi host expects the swap usage by the resource pool or virtual machine to be.
SWR/s (MB)	Rate at which the ESXi host swaps in memory from disk for the resource pool or virtual machine.
SWW/s (MB)	Rate at which the ESXi host swaps resource pool or virtual machine memory to disk.
LLSWR/s (MB)	Rate at which memory is read from the host cache. The reads and writes are attributed to the VMM group only, LLSWAP statistics are not displayed for the VM group.
LLSWW/s (MB)	Rate at which memory is written to the host cache from various sources. The reads and writes are attributed to the VMM group only, LLSWAP statistics are not displayed for the VM group.
CPTRD (MB)	Amount of data read from checkpoint file.
CPTTGT (MB)	Size of checkpoint file.
ZERO (MB)	Resource pool or virtual machine physical pages that are zeroed.
SHRD (MB)	Resource pool or virtual machine physical pages that are shared.
SHRDSVD (MB)	Machine pages that are saved because of resource pool or virtual machine shared pages.
OVHD (MB)	Current space overhead for resource pool.
OVHDMAX (MB)	Maximum space overhead that might be incurred by resource pool or virtual machine.
OVH DUW (MB)	Current space overhead for a user world. (You might see this statistic displayed, but it is intended for VMware use only.)
GST_NDx (MB)	Guest memory allocated for a resource pool on NUMA node x. This statistic is applicable on NUMA systems only.
OVD_NDx (MB)	VMM overhead memory allocated for a resource pool on NUMA node x. This statistic is applicable on NUMA systems only.

Field	Description
TCHD_W (MB)	Write working set estimate for resource pool.
CACHESZ (MB)	Compression memory cache size.
CACHEUSD (MB)	Used compression memory cache.
ZIP/s (MB/s)	Compressed memory per second.
UNZIP/s (MB/s)	Decompressed memory per second.

Table 540: Memory Panel Interactive Commands

Command	Description
M	Sort resource pools or virtual machines by MEMSZ column. This is the default sort order.
B	Sort resource pools or virtual machines by Group Memctl column.
N	Sort resource pools or virtual machines by GID column.
V	Display virtual machine instances only.
L	Changes the displayed length of the NAME column.

Storage Adapter Panel

Statistics in the Storage Adapter panel are aggregated per storage adapter by default. Statistics can also be viewed per storage path.

Table 541: Storage Adapter Panel Statistics

Column	Description
ADAPTR	Name of the storage adapter.
PATH	Storage path name. This name is only visible if the corresponding adapter is expanded. See interactive command e in <i>Storage Adapter Panel Interactive Commands</i> .
NPTH	Number of paths.
AQLEN	Current queue depth of the storage adapter.
CMDS/s	Number of commands issued per second.
READS/s	Number of read commands issued per second.
WRITES/s	Number of write commands issued per second.
MBREAD/s	Megabytes read per second.
MBWRTN/s	Megabytes written per second.
RESV/s	Number of SCSI reservations per second.
CONS/s	Number of SCSI reservation conflicts per second.
DAVG/cmd	Average device latency per command, in milliseconds.
KAVG/cmd	Average ESXi VMkernel latency per command, in milliseconds.
GAVG/cmd	Average virtual machine operating system latency per command, in milliseconds.
QAVG/cmd	Average queue latency per command, in milliseconds.
DAVG/rd	Average device read latency per read operation, in milliseconds.

Column	Description
KAVG/rd	Average ESXi VMkernel read latency per read operation, in milliseconds.
GAVG/rd	Average guest operating system read latency per read operation, in milliseconds.
QAVG/rd	Average queue latency per read operation, in milliseconds.
DAVG/wr	Average device write latency per write operation, in milliseconds.
KAVG/wr	Average ESXi VMkernel write latency per write operation, in milliseconds.
GAVG/wr	Average guest operating system write latency per write operation, in milliseconds.
QAVG/wr	Average queue latency per write operation, in milliseconds.
FCMDS/s	Number of failed commands issued per second.
FREAD/s	Number of failed read commands issued per second.
FWRITE/s	Number of failed write commands issued per second.
FMBRD/s	Megabytes of failed read operations per second.
FMBWR/s	Megabytes of failed write operations per second.
FRESV/s	Number of failed SCSI reservations per second.
ABRTS/s	Number of commands cancelled per second.
RESETS/s	Number of commands reset per second.
PAECMD/s	The number of PAE (Physical Address Extension) commands per second.
PAECP/s	The number of PAE copies per second.
SPLTCMD/s	The number of split commands per second.
SPLTCP/s	The number of split copies per second.

The following table displays the interactive commands you can use with the storage adapter panel.

Table 542: Storage Adapter Panel Interactive Commands

Command	Description
e	Toggles whether storage adapter statistics appear expanded or unexpanded. Allows you to view storage resource utilization statistics broken down by individual paths belonging to an expanded storage adapter. You are prompted for the adapter name.
r	Sorts by READS/s column.
w	Sorts by WRITES/s column.
R	Sorts by MBREAD/s read column.
T	Sorts by MBWRTN/s written column.
N	Sorts first by ADAPTR column, then by PATH column. This is the default sort order.

Storage Device Panel

The storage device panel displays server-wide storage utilization statistics.

By default, the information is grouped per storage device. You can also group the statistics per path, per world, or per partition.

Table 543: Storage Device Panel Statistics

Column	Description
DEVICE	Name of the storage device.
PATH	Path name. This name is visible only if the corresponding device is expanded to paths. See the interactive command <code>p</code> in <i>Storage Device Panel Interactive Commands</i> .
WORLD	World ID. This ID is visible only if the corresponding device is expanded to worlds. See the interactive command <code>e</code> in <i>Storage Device Panel Interactive Commands</i> . The world statistics are per world per device.
PARTITION	Partition ID. This ID is visible only if the corresponding device is expanded to partitions. See interactive command <code>t</code> in <i>Storage Device Panel Interactive Commands</i> .
NPH	Number of paths.
NWD	Number of worlds.
NPN	Number of partitions.
SHARES	Number of shares. This statistic is applicable only to worlds.
BLKSZ	Block size in bytes.
NUMBLKS	Number of blocks of the device.
DQLEN	Current device queue depth of the storage device.
WQLEN	World queue depth. This is the maximum number of ESXi VMkernel active commands that the world is allowed to have. This is a per device maximum for the world. It is valid only if the corresponding device is expanded to worlds.
ACTV	Number of commands in the ESXi VMkernel that are currently active. This statistic applies to only worlds and devices.
QUED	Number of commands in the ESXi VMkernel that are currently queued. This statistic applies to only worlds and devices.
%USD	Percentage of the queue depth used by ESXi VMkernel active commands. This statistic applies to only worlds and devices.
LOAD	Ratio of ESXi VMkernel active commands plus ESXi VMkernel queued commands to queue depth. This statistic applies to only worlds and devices.
CMDS/s	Number of commands issued per second.
READS/s	Number of read commands issued per second.
WRITES/s	Number of write commands issued per second.
MBREAD/s	Megabytes read per second.
MBWRTN/s	Megabytes written per second.
DAVG/cmd	Average device latency per command in milliseconds.
KAVG/cmd	Average ESXi VMkernel latency per command in milliseconds.
GAVG/cmd	Average guest operating system latency per command in milliseconds.
QAVG/cmd	Average queue latency per command in milliseconds.
DAVG/rd	Average device read latency per read operation in milliseconds.
KAVG/rd	Average ESXi VMkernel read latency per read operation in milliseconds.
GAVG/rd	Average guest operating system read latency per read operation in milliseconds.
QAVG/rd	Average queue read latency per read operation in milliseconds.

Column	Description
DAVG/wr	Average device write latency per write operation in milliseconds.
KAVG/wr	Average ESXi VMkernel write latency per write operation in milliseconds.
GAVG/wr	Average guest operating system write latency per write operation in milliseconds.
QAVG/wr	Average queue write latency per write operation in milliseconds.
ABRTS/s	Number of commands cancelled per second.
RESETS/s	Number of commands reset per second.
PAECMD/s	Number of PAE commands per second. This statistic applies to only paths.
PAECP/s	Number of PAE copies per second. This statistic applies to only paths.
SPLTCMD/s	Number of split commands per second. This statistic applies to only paths.
SPLTCP/s	Number of split copies per second. This statistic applies to only paths.

The following table displays the interactive commands you can use with the storage device panel.

Table 544: Storage Device Panel Interactive Commands

Command	Description
e	Expand or roll up storage world statistics. This command allows you to view storage resource utilization statistics separated by individual worlds belonging to an expanded storage device. You are prompted for the device name. The statistics are per world per device.
P	Expand or roll up storage path statistics. This command allows you to view storage resource utilization statistics separated by individual paths belonging to an expanded storage device. You are prompted for the device name.
t	Expand or roll up storage partition statistics. This command allows you to view storage resource utilization statistics separated by individual partitions belonging to an expanded storage device. You are prompted for the device name.
r	Sort by READS/s column.
w	Sort by WRITES/s column.
R	Sort by MBREAD/s column.
T	Sort by MBWRTN column.
N	Sort first by DEVICE column, then by PATH, WORLD, and PARTITION column. This is the default sort order.
L	Changes the displayed length of the DEVICE column.

Virtual Machine Storage Panel

This panel displays virtual machine-centric storage statistics.

By default, statistics are aggregated on a per-resource-pool basis. One virtual machine has one corresponding resource pool, so the panel displays statistics on a per-virtual-machine basis. You can also view the statistics on per-VSCSI-device basis.

Table 545: Virtual Machine Storage Panel Statistics

Column	Description
ID	Resource pool ID or VSCSI ID of VSCSI device.
GID	Resource pool ID.
VMNAME	Name of the resource pool.
VSCSINAME	Name of the VSCSI device.
NDK	Number of VSCSI devices
CMDS/s	Number of commands issued per second.
READS/s	Number of read commands issued per second.
WRITES/s	Number of write commands issued per second.
MBREAD/s	Megabytes read per second.
MBWRTN/s	Megabytes written per second.
LAT/rd	Average latency (in milliseconds) per read.
LAT/wr	Average latency (in milliseconds) per write.

The following table lists the interactive commands you can use with the virtual machine storage panel.

Table 546: Virtual Machine Storage Panel Interactive Commands

Command	Description
e	Expand or roll up storage VSCSI statistics. Allows you to view storage resource utilization statistics broken down by individual VSCSI devices belonging to a group. You are prompted to enter the group ID. The statistics are per VSCSI device.
r	Sort by READS/s column.
w	Sort by WRITES/s column.
R	Sort by MBREAD/s column.
T	Sort by MBWRTN/s column.
N	Sort first by VMNAME column, and then by VSCSINAME column. It is the default sort order.

Network Panel

The Network panel displays server-wide network utilization statistics.

Statistics are arranged by port for each virtual network device configured. For physical network adapter statistics, see the row in the table that corresponds to the port to which the physical network adapter is connected. For statistics on a virtual network adapter configured in a particular virtual machine, see the row corresponding to the port to which the virtual network adapter is connected.

Table 547: Network Panel Statistics

Column	Description
PORT-ID	Virtual network device port ID.

Column	Description
UPLINK	Y means that the corresponding port is an uplink. N means it is not.
UP	Y means that the corresponding link is up. N means it is not.
SPEED	Link speed in Megabits per second.
FDUPLX	Y means the corresponding link is operating at full duplex. N means it is not.
USED-BY	Virtual network device port user.
DTYP	Virtual network device type. H means HUB and S means switch.
DNAME	Virtual network device name.
PKTTX/s	Number of packets transmitted per second.
PKTRX/s	Number of packets received per second.
MbTX/s	MegaBits transmitted per second.
MbRX/s	MegaBits received per second.
%DRPTX	Percentage of transmit packets dropped.
%DRPRX	Percentage of receive packets dropped.
TEAM-PNIC	Name of the physical NIC used for the team uplink.
PKTTXMUL/s	Number of multicast packets transmitted per second.
PKTRXMUL/s	Number of multicast packets received per second.
PKTTXBRD/s	Number of broadcast packets transmitted per second.
PKTRXBRD/s	Number of broadcast packets received per second.

The following table displays the interactive commands you can use with the network panel.

Table 548: Network Panel Interactive Commands

Command	Description
T	Sorts by Mb Tx column.
R	Sorts by Mb Rx column.
t	Sorts by Packets Tx column.
r	Sorts by Packets Rx column.
N	Sorts by PORT-ID column. This is the default sort order.
L	Changes the displayed length of the DNAME column.

Interrupt Panel

The interrupt panel displays information about the use of interrupt vectors.

Table 549: Interrupt Panel Statistics

Column	Description
VECTOR	Interrupt vector ID.
COUNT/s	Total number of interrupts per second. This value is cumulative of the count for every CPU.

Column	Description
COUNT_x	Interrupts per second on CPU x.
TIME/int	Average processing time per interrupt (in microseconds).
TIME_x	Average processing time per interrupt on CPU x (in microseconds).
DEVICES	Devices that use the interrupt vector. If the interrupt vector is not enabled for the device, its name is enclosed in angle brackets (< and >).

Using Batch Mode

Batch mode allows you to collect and save resource utilization statistics in a file.

After you prepare for batch mode, you can use `esxtop` or `resxtop` in this mode.

Prepare for Batch Mode

To run in batch mode, you must first prepare for batch mode.

1. Run `resxtop` (or `esxtop`) in interactive mode.
2. In each of the panels, select the columns you want.
3. Save this configuration to a file (by default `~/ .esxtop50rc`) using the `w` interactive command.

You can now use `resxtop` (or `esxtop`) in batch mode.

Use `esxtop` or `resxtop` in Batch Mode

After you have prepared for batch mode, you can use `esxtop` or `resxtop` in this mode.

1. Start `resxtop` (or `esxtop`) to redirect the output to a file.

For example:

```
esxtop -b > my_file.csv
```

The filename must have a `.csv` extension. The utility does not enforce this, but the post-processing tools require it.

2. Process statistics collected in batch mode using tools such as Microsoft Excel and Perfmon.

In batch mode, `resxtop` (or `esxtop`) does not accept interactive commands. In batch mode, the utility runs until it produces the number of iterations requested (see command-line option `n`, below, for more details), or until you end the process by pressing `Ctrl+c`.

Batch Mode Command-Line Options

You can use batch mode with command-line options.

Table 550: Command-Line Options in Batch Mode

Option	Description
<code>a</code>	Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default <code>~/ .esxtop50rc</code> configuration file or a user-defined configuration file.
<code>b</code>	Runs <code>resxtop</code> (or <code>esxtop</code>) in batch mode.

Option	Description
<code>c filename</code>	Load a user-defined configuration file. If the <code>-c</code> option is not used, the default configuration filename is <code>~/esxstop41rc</code> . Create your own configuration file, specifying a different filename, using the <code>W</code> single-key interactive command.
<code>d</code>	Specifies the delay between statistics snapshots. The default is five seconds. The minimum is two seconds. If a delay of less than two seconds is specified, the delay is set to two seconds.
<code>n</code>	Number of iterations. <code>resxstop</code> (or <code>esxstop</code>) collects and saves statistics this number of times, and then exits.
<code>server</code>	The name of the remote server host to connect to (required, <code>resxstop</code> only).
<code>vihost</code>	If you connect indirectly (through vCenter Server), this option should contain the name of the ESXi host you connect to. If you connect directly to the ESXi host, this option is not used. Note: The host name needs to be the same as what appears in the vSphere Client.
<code>portnumber</code>	The port number to connect to on the remote server. The default port is 443, and unless this is changed on the server, this option is not needed. (<code>resxstop</code> only)
<code>username</code>	The user name to be authenticated when connecting to the remote host. You are prompted by the remote server for a password, as well (<code>resxstop</code> only).

Using Replay Mode

In replay mode, `esxstop` replays resource utilization statistics collected using `vm-support`.

After you prepare for replay mode, you can use `esxstop` in this mode. See the `vm-support` man page.

In replay mode, `esxstop` accepts the same set of interactive commands as in interactive mode and runs until no more snapshots are collected by `vm-support` to be read or until the requested number of iterations are completed.

Prepare for Replay Mode

To run in replay mode, you must prepare for replay mode.

1. Run `vm-support` in snapshot mode in the ESXi Shell.

Use the following command.

```
vm-support -p -i interval -d duration
```

2. Unzip and untar the resulting tar file so that `esxstop` can use it in replay mode.

You can now use `esxstop` in replay mode.

Use esxstop in Replay Mode

You can use `esxstop` in replay mode.

Replay mode can be run to produce output in the same style as batch mode (see the command-line option `b`, below).

Note: Batch output from `esxstop` cannot be played back by `resxstop`.

Snapshots collected by `vm-support` can be replayed by `esxstop`. However, `vm-support` output generated by ESXi can be replayed only by `esxstop` running on the same version of ESXi.

To activate replay mode, enter the following at the command-line prompt.

```
esxstop -R vm-support_dir_path
```

Replay Mode Command-Line Options

You can use replay mode with command-line options.

The following table lists the command-line options available for `esxtop` replay mode.

Table 551: Command-Line Options in Replay Mode

Option	Description
R	Path to the vm-support collected snapshot's directory.
a	Show all statistics. This option overrides configuration file setups and shows all statistics. The configuration file can be the default <code>~/ .esxtop50rc</code> configuration file or a user-defined configuration file.
b	Runs <code>esxtop</code> in Batch mode.
c <i>filename</i>	Load a user-defined configuration file. If the <code>-c</code> option is not used, the default configuration filename is <code>~/ .esxtop50rc</code> . Create your own configuration file and specify a different filename using the <code>w</code> single-key interactive command.
d	Specifies the delay between panel updates. The default is five seconds. The minimum is two seconds. If a delay of less than two seconds is specified, the delay is set to two seconds.
n	Number of iterations <code>esxtop</code> updates the display this number of times and then exits.

Using the vimtop Plug-In to Monitor the Resource Use of Services

You can use the `vimtop` utility plug-in to monitor vSphere services that run in the vCenter Server.

`vimtop` is a tool similar to `esxtop`, which runs in the environment of the vCenter Server. By using the text-based interface of `vimtop` in the appliance shell, you can view overall information about the vCenter Server, and a list of vSphere services and their resource use.

Monitor Services By Using vimtop in Interactive Mode

You can use the `vimtop` plug-in to monitor services in real time.

The default view of the `vimtop` interactive mode consists of the overview tables and the main table. You can use single-key commands in interactive mode to switch the view from processes to disks or network.

1. From an SSH client application, log in to the vCenter Server shell.
2. Run the `vimtop` command to access the plug-in in interactive mode.

Interactive Mode Command-Line Options

You can use various command-line options when you run the `vimtop` command to enter the plug-in interactive mode.

Table 552: Interactive Mode Command-Line Options

Option	Description
-h	Prints help for the <code>vimtop</code> command-line options.
-v	Prints the <code>vimtop</code> version number.
-c <i>filename</i>	Loads a user-defined <code>vimtop</code> configuration file. If the <code>-c</code> option is not used, the default configuration file is <code>/root/vimtop/vimtop.xml</code> .

Option	Description
	You can create your own configuration file, specifying a different filename and path by using the <code>w</code> single-key interactive command.
<code>-n number</code>	Sets the number of performed iterations before the <code>vimtop</code> exits interactive mode. <code>vimtop</code> updates the display <i>number</i> number of times and exits. The default value is 10000.
<code>-p / -dseconds</code>	Sets the update period in seconds.

Interactive Mode Single-Key Commands for vimtop

When running in interactive mode, `vimtop` recognizes several single-key commands.

All interactive mode panels recognize the commands listed in the following table.

Table 553: Interactive Mode Single-Key Commands

Key Names	Description
<code>h</code>	Show a help menu for the current panel, giving a brief summary of commands, and the status of secure mode.
<code>i</code>	Show or hide the top line view of the overview panel of the <code>vimtop</code> plug-in.
<code>t</code>	Show or hide the Tasks section, which displays information in the overview panel about the tasks currently running on the vCenter Server instance.
<code>m</code>	Show or hide the Memory section in the overview panel.
<code>f</code>	Show or hide the CPU section which displays information in the overview panel about all available CPUs.
<code>g</code>	Show or hide the CPUs section which displays information in the overview panel about the top 4 physical CPUs.
<code>spacebar</code>	Immediately refreshes the current pane.
<code>p</code>	Pause the displayed information about the services resource use in the current panels.
<code>r</code>	Refresh the displayed information about the services resource use in the current panels.
<code>s</code>	Set refresh period.
<code>q</code>	Exit the interactive mode of the <code>vimtop</code> plug-in.
<code>k</code>	Displays the Disks view of the main panel.
<code>o</code>	Switch the main panel to Network view.
<code>Esc</code>	Clear selection or return to the Processes view of the main panel.
<code>Enter</code>	Select a service to view additional details.
<code>n</code>	Show or hide names of the headers in the main panel.
<code>u</code>	Show or hide the measurement units in the headers in the main panel.
<code>left, right arrows</code>	Select columns.
<code>up, down arrows</code>	Select rows.
<code><, ></code>	Move a selected column.
<code>Delete</code>	Remove selected column.
<code>c</code>	Add a column to the current view of the main panel. Use <code>spacebar</code> to add or remove columns from the displayed list.

Key Names	Description
a	Sort the selected column in ascending order.
d	Sort the selected column in descending order.
z	Clear the sort order for all columns.
l	Set width for the selected column.
x	Return the column widths to their default values.
+	Expand selected item.
-	Collapse selected item.
w	Write the current setup to a <code>vimtop</code> configuration file. The default file name is the one specified by <code>-c</code> option, or <code>/root/vimtop/vimtop.xml</code> if the <code>-c</code> option is not used. You can also specify a different file name on the prompt generated by the <code>w</code> command.

Monitoring Networked Devices with SNMP and vSphere

Simple Network Management Protocol (SNMP) is commonly used by management programs to monitor a variety of networked devices.

vSphere systems run SNMP agents, which can provide information to a management program in at least one of the following ways:

- In response to a `GET`, `GETBULK`, or `GETNEXT` operation, which is a specific request for information from the management system.
- By sending a notification which is an alert sent by the SNMP agent to notify the management system of a particular event or condition.

Management Information Base (MIB) files define the information that can be provided by managed devices. The MIB files define managed objects, described by object identifiers (OIDs) and variables arranged in a hierarchy.

vCenter Server and ESXi have SNMP agents. The agent provided with each product has different capabilities.

Using SNMP Traps with vCenter Server

The SNMP agent included with vCenter Server can be used to send traps when vCenter Server starts and when an alarm is triggered on vCenter Server. The vCenter Server SNMP agent functions only as a trap emitter and does not support other SNMP operations, such as receiving `GET`, `GETBULK`, and `GETNEXT` requests.

vCenter Server can send SNMPv1 traps to other management applications. You must configure your management server to interpret the SNMP traps sent by vCenter Server.

To use the vCenter Server SNMP traps, configure the SNMP settings on vCenter Server and your management client software to accept the traps from vCenter Server.

The traps sent by vCenter Server are defined in `VMWARE-VC-EVENT-MIB.mib`.

Configure SNMP Settings for vCenter Server

If you plan to use SNMP with vCenter Server, you must use the vSphere Client to configure the SNMP settings.

- Verify that the vSphere Client is connected to a vCenter Server instance.

- Verify that you have the domain name or IP address of the SNMP receiver, the port number of the receiver, and the community string.

1. In the vSphere Client, navigate to a vCenter Server instance.
2. Click the **Configure** tab.
3. Under Settings, click **General**.
4. On the vCenter Server Settings central pane, click **Edit**.
The **Edit vCenter Server Settings** wizard opens.
5. Click **SNMP receivers** to edit their settings.
6. Enter the following information for the primary receiver of the SNMP traps.

Option	Description
Primary Receiver URL	Enter the domain name or IP address of the receiver of SNMP traps.
Enable receiver	Select the check box to enable the SNMP receiver.
Receiver port	Enter the port number of the receiver to which the SNMP agent sends traps. If the port value is empty, vCenter Server uses port 162 by default.
Community string	Enter the community string that is used for authentication.

7. Optional: Enter information about other SNMP receivers in the **Receiver 2 URL**, **Receiver 3 URL**, and **Receiver 4 URL** options, and select **Enabled**.
8. Click **OK**.

The vCenter Server system is now ready to send traps to the management system you have specified. Configure your SNMP management software to receive and interpret data from the vCenter Server SNMP agent. See [Configure SNMP Management Client Software](#).

Configure SNMP for ESXi

ESXi includes an SNMP agent that can send notifications (traps and informs) and receive `GET`, `GETBULK`, and `GETNEXT` requests.

In ESXi 5.1 and later releases, the SNMP agent adds support for version 3 of the SNMP protocol, offering increased security and improved functionality, including the ability to send informs. You can use `esxcli` commands to enable and configure the SNMP agent. You configure the agent differently depending on whether you want to use SNMP v1/v2c or SNMP v3.

As an alternative to configuring SNMP manually using `esxcli` commands, you can use host profiles to configure SNMP for an ESXi host. See the *vSphere Host Profiles* documentation for more information.

Note: For information on configuring SNMP for ESXi 5.0 or earlier or ESX 4.1 or earlier, see the documentation for the appropriate product version.

Configure the SNMP Agent for Polling

If you configure the ESXi SNMP agent for polling, it can listen for and respond to requests from SNMP management client systems, such as `GET`, `GETNEXT` and `GETBULK` requests.

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

By default, the embedded SNMP agent listens on UDP port 161 for polling requests from management systems. You can use the `esxcli system snmp set` command with the `--port` option to configure an alternative port. To avoid conflicting with other services, use a UDP port that is not defined in `/etc/services`.

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

1. Run the `esxcli system snmp set` command with the `--port` option to configure the port.

For example, run the following command:

```
esxcli system snmp set --port port
```

Here, `port` is the port the SNMP agent uses to listen for polling requests.

Note: The port you specify must not be already in use by other services. Use IP addresses from the dynamic range, port 49152 and up.

2. Optional: If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

Configure ESXi for SNMPv1 and SNMPv2c

When you configure the ESXi SNMP agent for SNMPv1 and SNMPv2c, the agent supports sending notifications (traps and informs) and receiving `GET`, `GETBULK`, and `GETNEXT` requests.

In SNMPv1 and SNMPv2c, authentication is performed by using community strings. Community strings are namespaces which contain one or more managed objects. This form of authentication does not secure the communication between the SNMP agent and the management system. To secure the SNMP communication in your environment, use SNMPv3.

Configure SNMP Communities

To enable the ESXi SNMP agent to send and receive SNMP v1 and v2c messages, you must configure at least one community for the agent.

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

An SNMP community defines a group of devices and management systems. Only devices and management systems that are members of the same community can exchange SNMP messages. A device or management system can be a member of multiple communities.

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

Run the `esxcli system snmp set` command with the `--communities` option to configure an SNMP community.

For example, to configure `public`, `East`, and `West` network operation centers communities, run the following command:

```
esxcli system snmp set --communities public,eastnoc,westnoc
```

Each time you specify a community with this command, the settings you specify overwrite the previous configuration. To specify multiple communities, separate the community names with a comma.

Configure the SNMP Agent to Send SNMP v1 or v2c Notifications

You can use the ESXi SNMP agent to send virtual machine and environmental notifications to management systems.

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

To send SNMP v1/v2c notifications with the SNMP agent, you must configure the target (receiver) unicast address, community, and an optional port. If you do not specify a port, the SNMP agent sends traps to UDP port 162 on the target management system by default.

To configure SNMP v3 traps, see [Configure SNMP v3 Targets](#).

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

1. Run the `esxcli system snmp set` command with the `--targets` option:

```
esxcli system snmp set --targets target_address@port/community
```

Here, `target_address` is the address of the target system, `port` is the port number to send the notifications to, and `community` is the community name.

Each time you specify a target with this command, the settings you specify overwrite all previously specified settings. To specify multiple targets, separate them with a comma.

For example, run the following command for configuring the targets 192.0.2.1@163/westnoc and 2001:db8::1@163/eastnoc:

```
esxcli system snmp set --targets 192.0.2.1@163/westnoc,2001:db8::1@163/eastnoc
```

2. Optional: If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

3. Optional: Send a test trap to verify that the agent is configured correctly by running the `esxcli system snmp test` command.

The agent sends a `warmStart` trap to the configured target.

Configure ESXi for SNMP v3

When you configure the ESXi SNMP agent for SNMPv3, the agent supports sending notifications (traps and informs) and receiving `GET`, `GETBULK`, and `GETNEXT` requests. SNMPv3 also provides stronger security than SNMPv1 or SNMPv2c, including key authentication and encryption.

Inform is a notification that the sender resends up to three times or until the receiver acknowledges the notification.

Configure the SNMP Engine ID

Every SNMP v3 agent has an engine ID which serves as a unique identifier for the agent. The engine ID is used with a hashing function to generate keys for authentication and encryption of SNMP v3 messages.

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

If you do not specify an engine ID, when you enable the SNMP agent, an engine ID is automatically generated.

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

Run the `esxcli system snmp set` command with the `--engineid` option to configure the SNMP engine ID.

For example, run the following command:

```
esxcli system snmp set --engineid id
```

Here, `id` is the engine ID and it must be a hexadecimal string between 5 and 32 characters long.

Configure SNMP Authentication and Privacy Protocols

SNMPv3 optionally supports authentication and privacy protocols.

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

Authentication is used to ensure the identity of users. Privacy allows for encryption of SNMP v3 messages to ensure confidentiality of data. These protocols provide a higher level of security than is available in SNMPv1 and SNMPv2c, which use community strings for security.

Both authentication and privacy are optional. However, you must enable authentication to enable privacy.

The SNMPv3 authentication and privacy protocols are licensed vSphere features and might not be available in some vSphere editions.

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

1. Optional: Run the `esxcli system snmp set` command with the `--authentication` option to configure authentication. For example, run the following command:

```
esxcli system snmp set --authentication protocol
```

Here, `protocol` must be either `none` (for no authentication) or `SHA1`.

2. Optional: Run the `esxcli system snmp set` command with the `--privacy` option to configure privacy. For example, run the following command:

```
esxcli system snmp set --privacy protocol
```

Here, `protocol` must be either `none` (for no privacy) or `AES128`.

Configure SNMP Users

You can configure up to 5 users who can access SNMP v3 information. User names must be no more than 32 characters long.

- Verify that you have configured the authentication and privacy protocols before configuring users.
- Verify that you know the authentication and privacy passwords for each user you plan to configure. Passwords must be at least 7 characters long. Store these passwords in files on the host system.
- Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

While configuring a user, you generate authentication and privacy hash values based on the user's authentication and privacy passwords and the SNMP agent's engine ID. If you change the engine ID, the authentication protocol, or the privacy protocol after configuring users, the users are no longer valid and must be reconfigured.

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

1. If you are using authentication or privacy, get the authentication and privacy hash values for the user by running the `esxcli system snmp hash` command with the `--auth-hash` and `--priv-hash` flags.

For example, run the following command:

```
esxcli system snmp hash --auth-hash secret1 --priv-hash secret2
```

Here, `secret1` is the path to the file containing the user's authentication password and `secret2` is the path to the file containing the user's privacy password.

Alternatively, you can pass the `--raw-secret` flag and specify the passwords directly on the command line.

For example, you can run the following command:

```
esxcli system snmp hash --auth-hash authsecret --priv-hash privsecret --raw-secret
```

The produced output might be the following:

```
Authhash: 08248c6eb8b333e75a29ca0af06b224faa7d22d6
Privhash: 232ba5cbe8c55b8f979455d3c9ca8b48812adb97
```

The authentication and privacy hash values are displayed.

2. Configure the user by running the `esxcli system snmp set` command with the `--user` flag.

For example, you can run the following command:

```
esxcli system snmp set --users userid/authhash/privhash/security
```

The command accepts the following parameters:

Parameter	Description
<code>userid</code>	The user name.
<code>authhash</code>	The authentication hash value.
<code>privhash</code>	The privacy hash value.
<code>security</code>	The level of security enabled for that user, which can be <code>auth</code> (for authentication only), <code>priv</code> (for authentication and privacy), or <code>none</code> (for no authentication or privacy).

For example, run the following command to configure user1 for access with authentication and privacy:

```
esxcli system snmp set --users user1/08248c6eb8b333e75a29ca0af06b224faa7d22d6/
232ba5cbe8c55b8f979455d3c9ca8b48812adb97/priv
```

You must run the following command to configure user2 for access with no authentication or privacy:

```
esxcli system snmp set --users user2/-/-/none
```

3. Optional: Test the user configuration by running the following command:

```
esxcli system snmp test --userusername--auth-hashsecret1--priv-hashsecret2
```

If the configuration is correct, this command returns the following message: "User `username` validated correctly using engine id and security level: `protocols`". Here, `protocols` indicates the security protocols configured.

Configure SNMP v3 Targets

Configure SNMP v3 targets to allow the ESXi SNMP agent to send SNMP v3 traps and informs.

- Ensure that the users who access the traps or informs are configured as SNMP users for both the ESXi SNMP agent and the target management system.
- If you are configuring informs, you need the engine ID for the SNMP agent on the remote system that receives the inform message.
- Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

SNMP v3 allows for sending both traps and informs. An inform message is a type of a message that the sender resends a maximum of three times. The sender waits for 5 seconds between each attempt, unless the message is acknowledged by the receiver.

You can configure a maximum of three SNMP v3 targets, in addition to a maximum of three SNMP v1/v2c targets.

To configure a target, you must specify a hostname or IP address of the system that receives the traps or informs, a user name, a security level, and whether to send traps or informs. The security level can be either `none` (for no security), `auth` (for authentication only), or `priv` (for authentication and privacy).

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

- Optional: If you are configuring informs, configure the remote users by running the `esxcli system snmp set` command with the `--remote-users` option.

For example, run the following command:

```
esxcli system snmp set --remote-users userid/auth-protocol/auth-hash/priv-protocol/priv-hash/engine-id
```

The command accepts the following parameters:

Parameter	Description
<code>userid</code>	The user name.
<code>auth-protocol</code>	The authentication protocol, <code>none</code> (for no authentication) or <code>SHA1</code> .
<code>auth-hash</code>	The authentication hash or – if the authentication protocol is <code>none</code> .
<code>priv-protocol</code>	The privacy protocol, <code>AES128</code> , or <code>none</code> .
<code>priv-hash</code>	The privacy hash, or – if the privacy protocol is <code>none</code> .
<code>engine-id</code>	The engine ID of the SNMP agent on the remote system that receives the inform message.

- Run the `esxcli system snmp set` command with the `--v3targets` option.

For example, run the following command:

```
esxcli system snmp set --v3targets hostname@port/userid/secLevel/message-type
```

The parameters of the command are as follows.

Parameter	Description
<code>hostname</code>	The host name or IP address of the management system that receives the traps or informs.
<code>port</code>	The port on the management system that receives the traps or informs. If you do not specify a port, the default port, 162, is used.
<code>userid</code>	The user name.
<code>secLevel</code>	The level of authentication and privacy you have configured. Use <code>auth</code> if you have configured authentication only, <code>priv</code> if you have configured both authentication and privacy, and <code>none</code> if you have configured neither.
<code>message-type</code>	The type of the messages received by the management system. Use <code>trap</code> or <code>inform</code> .

- Optional: If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

- Optional: Send a test notification to verify that the agent is configured correctly by running the `esxcli system snmp test` command.

The agent sends a `warmStart` notification to the configured target.

Configure the Source of Hardware Events Received by the SNMP Agent

You can configure the ESXi SNMP agent to receive hardware events either from IPMI sensors or CIM indications.

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

IPMI sensors are used for hardware monitoring in ESX/ESXi 4.x and earlier. The conversion of CIM indications to SNMP notifications is available in ESXi 5.0 and later.

If you run ESXCLI commands through ESXCLI, you must supply connection options that specify the target host and login credentials. If you use ESXCLI commands directly on a host using the ESXi Shell, you can use the commands as given without specifying connection options. For more information on connection options see *ESXCLI Concepts and Examples*.

1. Run the `esxcli system snmp set --hwsrc source` command to configure the source for hardware events. Here, *source* is `sensors` or `indications`, for hardware event received from IPMI sensors or CIM indications respectively.
2. Optional: If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

Configure the SNMP Agent to Filter Notifications

You can configure the ESXi SNMP agent to filter out notifications if you don't want your SNMP management software to receive those notifications.

Configure the ESXi SNMP agent by using the ESXCLI commands. See *Getting Started with ESXCLI* for more information on how to use ESXCLI.

1. Run the `esxcli system snmp set` command to filter notifications:

```
esxcli system snmp set --notraps oid_list
```

Here, *oid_list* is a list of OIDs for the notifications to filter, separated by commas. This list replaces any OIDs that were previously specified using this command.

For example, to filter out `coldStart` (OID 1.3.6.1.4.1.6876.4.1.1.0) and `warmStart` (OID 1.3.6.1.4.1.6876.4.1.1.1) traps, run the following command:

```
esxcli system snmp set --notraps 1.3.6.1.4.1.6876.4.1.1.0,1.3.6.1.4.1.6876.4.1.1.1
```
2. Optional: If the ESXi SNMP agent is not enabled, run the following command:

```
esxcli system snmp set --enable true
```

The traps identified by the specified OIDs are filtered out of the output of the SNMP agent, and are not sent to SNMP management software.

To clear all notification filters, run the `esxcli system snmp set --notraps reset` command.

Configure SNMP Management Client Software

After you have configured a vCenter Server instance or an ESXi host to send traps, configure your management client software to receive and interpret those traps.

To complete this task, download the VMware MIB files from <https://kb.vmware.com/s/article/1013445>.

To configure your management client software, specify the communities for the managed device, configure the port settings, and load the VMware MIB files. See the documentation for your management system for specific instructions for these steps.

1. In your management software, specify the vCenter Server instance or ESXi host as an SNMP-based managed device.
2. If you are using SNMPv1 or SNMPv2c, set up appropriate community names in the management software.
These names must correspond to the communities set for the SNMP agent on the vCenter Server instance or ESXi host.
3. If you are using SNMPv3, configure users and authentication and privacy protocols to match those configured on the ESXi host.
4. If you configured the SNMP agent to send traps to a port on the management system other than the default UDP port 162, configure the management client software to listen on the port you configured.
5. Load the VMware MIBs into the management software so you can view the symbolic names for vCenter Server or the host variables.

To prevent lookup errors, load these MIB files in the following order before loading other MIB files:

1. VMWARE-ROOT-MIB.mib
2. VMWARE-TC-MIB.mib
3. VMWARE-PRODUCTS-MIB.mib

The management software can now receive and interpret traps from vCenter Server or ESXi hosts.

SNMP Diagnostics

You can use SNMP tools to diagnose configuration problems.

- Run the `esxcli system snmp test` command from the ESXCLI set to prompt the SNMP agent to send a test warmStart trap.
- Run the `esxcli system snmp get` command to display the current configuration of the SNMP agent.
- The `SNMPv2-MIB.mib` file provides several counters to aid in debugging SNMP problems. See [SNMPv2 Diagnostic Counters](#).
- The `VMWARE-ESX-AGENTCAP-MIB.mib` file defines the capabilities of the VMware SNMP agents by product version. Use this file to determine if the SNMP functionality that you want to use is supported.

Monitor Guest Operating Systems with SNMP

You can use SNMP to monitor guest operating systems or applications running in virtual machines.

The virtual machine uses its own virtual hardware devices. Do not install agents in the virtual machines that are intended to monitor physical hardware.

- Install the SNMP agents you normally would use for that purpose in the guest operating systems.

VMware MIB Files

VMware MIB files contain information provided by ESXi hosts and vCenter Server to SNMP management software.

Management Information Base (MIB) files define the information that can be provided by managed devices. The MIB files define managed objects, described by object identifiers (OIDs) and variables arranged in a hierarchy. SMI structure of management information (RFC 2578) is the syntax used to write MIB files for specific product and function. These MIB files are versioned independently of product and can be used to identify event types and event data related information.

You can download these MIB files from <https://kb.vmware.com/s/article/1013445>.

Table 1. VMware MIB Files lists the MIB files provided by VMware and describes the information provided by each file.

Table 554: VMware MIB Files

MIB File	Description
VMWARE-ROOT-MIB.mib	Contains VMware's enterprise OID and top-level OID assignments.
VMWARE-ESX-AGENTCAP-MIB.mib	Defines the capabilities of the VMware agents by product versions. This file is optional and might not be supported by all management systems.
VMWARE-CIMOM-MIB.mib	Defines variables and trap types used to report on the state of the CIM Object Management subsystem.
VMWARE-ENV-MIB.mib	Defines variables and trap types used to report on the state of physical hardware components of the host computer. Enables conversion of CIM indications to SNMP traps.
VMWARE-OBSOLETE-MIB.mib	For use with versions of ESX/ESXi prior to 4.0. Defines OIDs that have been made obsolete to maintain backward compatibility with earlier versions of ESX/ESXi. Includes variables formerly defined in the files VMWARE-TRAPS-MIB.mib and VMWARE-VMKERNEL-MIB.mib.
VMWARE-PRODUCTS-MIB.mib	Defines OIDs to uniquely identify each SNMP agent on each VMware platform by name, version, and build platform.
VMWARE-RESOURCES-MIB.mib	Defines variables used to report information on resource usage of the VMkernel, including physical memory, CPU, and disk utilization.
VMWARE-SYSTEM-MIB.mib	The VMWARE-SYSTEM-MIB.mib file is obsolete. Use the SNMPv2-MIB to obtain information from sysDescr.0 and sysObjec ID.0.
VMWARE-TC-MIB.mib	Defines common textual conventions used by VMware MIB files.
VMWARE-VC-EVENT-MIB.mib	Defines traps sent by vCenter Server. Load this file if you use vCenter Server to send traps.
VMWARE-VMINFO-MIB.mib	Defines variables for reporting information about virtual machines, including virtual machine traps.

Table 2. Other MIB Files lists MIB files included in the VMware MIB files package that are not created by VMware. These can be used with the VMware MIB files to provide additional information.

Table 555: Other MIB Files

MIB File	Description
ENTITY-MIB.mib	Allows the description of relationships between physical entities and logical entities managed by the same SNMP agent. See RFC 4133 for more information.
HOST-RESOURCES-MIB.mib	Defines objects that are useful for managing host computers.
HOST-RESOURCES-TYPES.mib	Defines storage, device, and filesystem types for use with HOST-RESOURCES-MIB.mib.
IEEE8021-BRIDGE-MIB	Defines objects for managing devices that support IEEE 802.1D.
IEEE8023-LAG-MIB	Defines objects for managing devices that support IEEE 802.3ad link aggregation.
IEEE8021-Q-BRIDGE-MIB	Defines objects for managing Virtual Bridged Local Area Networks.
IF-MIB.mib	Defines attributes related to physical NICs on the host system.

MIB File	Description
IP-MIB.mib	Defines objects for managing implementations of the Internet Protocol (IP) in an IP version-independent manner.
IP-FORWARD-MIB.mib	Defines objects for managing IP forwarding.
LLDP-V2-MIB.mib	Defines objects for managing devices using Linked Layer Discovery Protocol (LLDP).
SNMPv2-CONF.mib	Defines conformance groups for MIBs.
SNMPv2-MIB.mib	Defines the SNMP version 2 MIB objects.
SNMPv2-SMI.mib	Defines the Structure of Management Information for SNMP version 2.
SNMPv2-TC.mib	Defines textual conventions for SNMP version 2.
TCP-MIB.mib	Defines objects for managing devices using the TCP protocol.
UDP-MIB.mib	Defines objects for managing devices using the UDP protocol.

SNMPv2 Diagnostic Counters

The `SNMPv2-MIB.mib` file provides a number of counters to aid in debugging SNMP problems.

[Table 556: Diagnostic Counters from SNMPv2-MIB](#) lists some of these diagnostic counters.

Table 556: Diagnostic Counters from SNMPv2-MIB

Variable	ID Mapping	Description
<code>snmpInPkts</code>	<code>snmp 1</code>	The total number of messages delivered to the SNMP entity from the transport service.
<code>snmpInBadVersions</code>	<code>snmp 3</code>	The total number of SNMP messages that were delivered to the SNMP entity and were for an unsupported SNMP version.
<code>snmpInBadCommunityNames</code>	<code>snmp 4</code>	The total number of community-based SNMP messages delivered to the SNMP entity that used an invalid SNMP community name.
<code>snmpInBadCommunityUses</code>	<code>snmp 5</code>	The total number of community-based SNMP messages delivered to the SNMP entity that represented an SNMP operation that was not allowed for the community named in the message.
<code>snmpInASNParseErrs</code>	<code>snmp 6</code>	The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
<code>snmpEnableAuthenTraps</code>	<code>snmp 30</code>	Indicates whether the SNMP entity is permitted to generate authenticationFailure traps. The value of this object overrides any configuration information. It therefore provides a means of disabling all authenticationFailure traps.
<code>snmpSilentDrops</code>	<code>snmp 31</code>	The total number of Confirmed Class PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate Response Class PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.

Variable	ID Mapping	Description
snmpProxyDrops	snmp 32	The total number of Confirmed Class PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in a manner other than a time-out such that no Response Class PDU could be returned.

System Log Files

In addition to lists of events and alarms, vSphere components generate assorted logs.

These logs contain additional information about activities in your vSphere environment.

View System Logs on an ESXi Host

You can use the Direct Console User Interface (DCUI) to view system logs on an ESXi host. These logs provide information about system operational events.

1. From the ESXi Host Client, click **Manage > Services**. Confirm that the DCUI service is running.
2. Open an SSH client (such as PuTTY) and connect to the ESXi host.
3. Log in with the root account.
4. Enter `DCUI` to start the Direct Console User Interface.
5. Click inside the console window and press F2 to customize the system.
6. Type the password for the root user and press Enter.
7. Use the arrow keys to select **View System Logs**.
8. Press a corresponding number key to view a log.
vCenter Server agent (vpxa) logs appear if the host is managed by vCenter Server.
9. Press Enter or the spacebar to scroll through the messages.
10. Optional: Perform a regular expression search.
 - a) Press the slash key (/).
 - b) Type the text to find.
 - c) Press Enter

The found text is highlighted on the screen.
11. Press q to return to the direct console.

Configure Syslog on ESXi Hosts

You can use the vSphere Client, the VMware Host Client, or the `esxcli system syslog` command to configure the syslog service.

The syslog service receives, categorizes, and stores log messages for analyses that help you take preventive action in your environment.

System Logs

VMware technical support might request several log files to help resolve any issues you might have with the product. This section describes the types and locations of log files found on ESXi hosts and vSphere Client.

ESXi System Logs

You might need the ESXi system log files to resolve technical issues.

The ESXi system logs can be found in the `/var/run/log` directory.

vSphere Client Logs

You might need the vSphere Client system log files to resolve technical issues.

When you use a vCenter Server instance, the vSphere Client system logs can be found in the location listed in the table.

Table 557: Location of vSphere Client Logs

vCenter Server System	Location
vCenter Server	<code>/var/log/vmware/vsphere-ui/logs</code>

The main vSphere Client log file is `vsphere_client_virgo.log`.

Export System Log Files

When the vSphere Client is connected to vCenter Server, you can select hosts from which to download system log files.

To save diagnostic data for ESXi hosts and vCenter Server, the vSphere Client must be connected to the vCenter Server system.

Required privileges:

- To view diagnostic data: **Read-Only User**
- To manage diagnostic data: **Global > Diagnostics**

- In the inventory, right click a vCenter Server instance.
Export System Logs option is available for Hosts and Virtual Machines.
- Click **Export System Logs**.
- If you are connected to vCenter Server, select the object for which you want to export data.
Selecting an object selects all of its child objects.
- If you are connected to vCenter Server, select **Include information from vCenter Server and vSphere UI logs** to download vCenter Server and vSphere Client log files and host log files, and click **Next**.
- If the selected host supports manifest driven exports of system log files, select the system log files to collect. Select the specific system log files to download.
If the host does not support manifest exports of log files, all system log files are exported.
- Select **Gather performance data** to include performance data information in the log files.
You can update the duration and interval time you want the data collected.
- Optional: Select to apply a password for encrypted core dumps to the support package.
You can make that password available to your support representative in a secure channel.
If only some of the host in your environment use encryption, some of the files in the package are encrypted.

8. Click **Finish**.

9. Specify the location to which to save the log files.

The host or vCenter Server generates a `.zip` file containing the log files.

The **Recent Tasks** panel shows the Generate diagnostic bundles task in progress.

The Downloading Log Bundles dialog box appears when the Generating Diagnostic Bundle task is finished. The download status of each bundle appears in the dialog box.

Some network errors can cause download failures. When you select an individual download in the dialog box, the error message for that operation appears under the name and location of the log bundle file.

10. Click **Save**.

11. Verify the information in the Summary and click **Finish** to download the log files.

Diagnostic bundles containing log files for the specified objects are downloaded to the location specified.

[Upload Logs Package to a VMware Service Request](#).

ESXi Log Files

Log files are an important component of troubleshooting attacks and obtaining information about breaches. All ESXi hosts run a syslog service, which logs messages from the VMkernel and other system components to local files or to a remote host.

To increase the security of the host, take the following measures.

- Configure persistent logging to a datastore. By default, the logs on ESXi hosts are stored in the in-memory file system. Therefore, they are lost when you reboot the host, and only 24 hours of log data is stored. When you enable persistent logging, you have a dedicated activity record for the host.
- Remote logging to a central host allows you to gather log files on a central host. From that host, you can monitor all hosts with a single tool, do aggregate analysis, and search log data. This approach facilitates monitoring and reveals information about coordinated attacks on multiple hosts.
- Configure the remote secure syslog on ESXi hosts by using ESXCLI or PowerCLI, or by using an API client.
- Query the syslog configuration to make sure that the syslog server and port are valid.

See the *vSphere Monitoring and Performance* documentation for information about syslog setup, and for additional information on ESXi log files.

Upload Logs Package to a VMware Service Request

If you already have a Service Request ID with VMware, you can use the vSphere Client to upload the system log bundles directly to your service request.

Request a Service Request ID from VMware Technical Support.

1. In the vSphere Client, navigate to **Administration**.
2. Click **Support**, and click **Upload File to Service Request**.

An Upload File to Service Request dialog box opens.

3. Enter your Service Request ID with VMware.
4. Click **Choose File**, and select the log bundle you want to attach to your service request with VMware, and click **OK**.
5. If you protected your support package with a password, provide the password to VMware Technical Support by using a secure channel.

The log bundle is sent to your service request.

Manage the Amount of Virtual Machine Log Data

Virtual machines can write support and troubleshooting information into a virtual machine log file stored on a VMFS volume.

The default settings for virtual machines are appropriate for most situations.

If your environment relies heavily on using vMotion, or if the defaults do not seem suitable for other reasons, you can modify the logging settings for virtual machine guest operating systems.

A new log file is ALWAYS created during virtual machine power-on, resume, migration with vMotion, and so on.

You can further adjust the amount of log data generated by a virtual machine by making adjustments to two parameters:

- `vmx.log.keepOld`: The number of log files retained per virtual machine is controlled by the `vmx.log.keepOld` property.
Note: VMware recommends saving a minimum of 10 log files. If you need log files for a longer time span, you can set `vmx.log.keepOld` to 20 or more.
- `vmx.log.rotateSize`: The maximum size of each log file, in bytes, is controlled by `vmx.log.rotateSize` property.
 To ensure that a new log file is created more frequently, you can limit the maximum size of the log files with this configuration parameter. Each time an entry is written to the log, the size of the log is checked. If `vmx.log.rotateSize` is set to a non default value, and the size is over the limit, the next entry is written to a new log file. If the maximum number of log files exists, the oldest log file is deleted.
Note: If this parameter is used, VMware recommends setting it to a minimum of 2000000 bytes.

The total amount of storage space consumed by all the log files is the product of these two parameters.

These settings may be set on an individual virtual machine basis or for all virtual machines on an ESXi host.

The following procedure discusses limiting the amount of virtual machine log data for all virtual machines on an ESXi host:

- To place limits on the amount of virtual machine log data, edit `/etc/vmware/config`
- If the `vmx.log.keepOld` and `vmx.log.rotateSize` properties are not defined in the file, add them.

For example, to retain ten log files, each one no larger than 2000000 bytes, add the following to `/etc/vmware/config`:

```
vmx.log.keepOld = "10"
```

```
vmx.log.rotateSize = "2000000"
```

The maximum amount of storage space consumed by each virtual machine is $10 * 2000000$ bytes.

Note:

You can use a PowerCLI script to change these properties for all the virtual machines on a host.

Handle Virtual Machine Log Throttling

You can deal with the log throttling of a virtual machine.

By default, virtual machine logs monitor the logging rate. If the logging rate exceeds an established rate, logging will be throttled (logging will be suspended for a time before resuming). This prevents problematic logging from consuming all of the available log file storage space.

Virtual machine log throttle is identified by the presence of `<< file Throttled>>` in virtual machine log files.

Note: If throttling occurs, you may be asked to temporarily disable log throttling for debugging purposes. Once the debugging has completed, it is very important to re-enable throttling.

The virtual machine log throttling rate is controlled by the `vmx.log.throttleBytesPerSec` property.

This setting may be set on an individual virtual machine basis or for all virtual machines on an ESXi host.

The following procedure discusses disabling virtual machine throttling for all virtual machines on an ESXi host:

- To disable virtual machine log throttling, edit `/etc/vmware/config`.
- If the `vmx.log.throttleBytesPerSec` property is not defined in the file, add it.

For example:

```
vmx.log.throttleBytesPerSec = "-1"
```

Note: You can use a PowerCLI script to change these properties for all the virtual machines on a host.

Collecting Log Files

VMware technical support might request several files to help resolve technical issues. The following sections describe script processes for generating and collecting some of these files.

Set Verbose Logging

You can set the verbose log file specification

You can set verbose logging for vpxd logs.

1. Select a vCenter Server instance.
2. Click **Configure** tab.
3. Click **General** and in the vCenter Server Settings, click **EDIT**.
4. In the Edit vCenter general settings window, Click **Logging Settings** and select **verbose** from the Log level drop down menu.
5. Click **SAVE**.

Collect vSphere Log Files

You can collect vSphere log files in to a single location.

You can download the log bundle from vSphere Client connected to a vCenter Server system.

1. In vSphere Client menu, select **Administration > Deployment > System Configuration**.
2. Select a vCenter Server node and click **Export Support Bundle**.
3. Select the support bundle type:
 - Full Support Bundle: Contains information on all services.
 - Minimal Support Bundle: Contains basic product and configuration information.
 - Custom Support Bundle: Contains customized information. Select any of the manifest properties to download the corresponding log files.

4. Click **EXPORT**.
5. Optional: You can also download the vCenter log bundle using another method:
 1. Right click vCenter Server instance and click **Actions > Export System Logs....**
 2. Click **Select hosts** to put selected host logs into an export bundle. Check **Include vCenter Server and vSphere UI Client logs** to optionally include vCenter Server and vSphere UI logs in the bundle. Click **NEXT**.
 3. Click **Select logs** to select specific system logs to export. Click **EXPORT LOGS**.

Note: The log bundle is generated as a .zip file. By default, the vpxd logs within the bundle are compressed as .tgz files. You must use `gunzip` to uncompress these files.

Collect ESXi Log Files

You can collect and package all relevant ESXi system and configuration information, as well as ESXi log files. This information can be used to analyze the problems.

- Run the following script on the ESXi Shell: `/usr/bin/vm-support`
The resulting file has the following format: `esx-date-unique-xnumber.tgz`

ESXi Log File Locations

ESXi records host activity in log files, using a syslog facility.

Table 558: ESXi Log File Locations

Component	Location	Purpose
Authentication	<code>/var/log/auth.log</code>	Contains all events related to authentication for the local system.
ESXi host agent log	<code>/var/log/hostd.log</code>	Contains information about the agent that manages and configures the ESXi host and its virtual machines.
Shell log	<code>/var/log/shell.log</code>	Contains a record of all commands typed into the ESXi Shell and shell events (for example, when the shell was enabled).
System messages	<code>/var/log/syslog.log</code>	Contains all general log messages and can be used for troubleshooting. This information was formerly located in the messages log file.
vCenter Server agent log	<code>/var/log/vpxa.log</code>	Contains information about the agent that communicates with vCenter Server (if the host is managed by vCenter Server).
Virtual machines	The same directory as the affected virtual machine's configuration files, named <code>vmware.log</code> and <code>vmware*.log</code> . For example, <code>/vmfs/volumes/datastore/virtual machine/vmware.log</code>	Contains virtual machine power events, system failure information, tools status and activity, time sync, virtual hardware changes, vMotion migrations, machine clones, and so on.
VMkernel	<code>/var/log/vmkernel.log</code>	Records activities related to virtual machines and ESXi.
VMkernel summary	<code>/var/log/vmksummary.log</code>	Used to determine uptime and availability statistics for ESXi (comma separated).

Component	Location	Purpose
VMkernel warnings	/var/log/vmkwarning.log	Records activities related to virtual machines and ESXi.
Quick Boot	/var/log/loadESX.log	Contains all events related to restarting an ESXi host through Quick Boot.
Trusted infrastructure agent	/var/run/log/kmxa.log	Records activities related to the Client Service on the ESXi Trusted Host.
Key Provider Service	/var/run/log/kmxd.log	Records activities related to the vSphere Trust Authority Key Provider Service.
Attestation Service	/var/run/log/attestd.log	Records activities related to the vSphere Trust Authority Attestation Service.
ESX Token Service	/var/run/log/esxtokend.log	Records activities related to the vSphere Trust Authority ESX Token Service.
ESX API Forwarder	/var/run/log/esxapiadapter.log	Records activities related to the vSphere Trust Authority API forwarder.

Configure Log Filtering on ESXi Hosts

The log filtering capability lets you modify the logging policy of the syslog service that is running on an ESXi host.

You can create log filters to reduce the number of repetitive entries in the ESXi logs and to denylist specific log events entirely.

Install ESXCLI. See *Getting Started with ESXCLI*. For troubleshooting, run `esxcli` commands in the ESXi Shell.

Starting with vSphere 7.0 Update 2, you can add log filters and enable log filtering by using ESXCLI. A log filter, once established, remains in place until it is removed, even across ESXi reboots.

Log filters affect all log events that are processed by the ESXi host `vm syslogd` service, whether they are recorded to a log directory or to a remote syslog server.

You must enable the log filtering capability and reload the syslog daemon to activate the log filters on the ESXi host.

ESXCLI commands to configure log filters follow this pattern: `esxcli system syslog config logfilter {cmd} [cmd options]`.

For example, to get the list of available log filters, run the following command: `[root@xxx-xx-dhcp-xx-xx:~] esxcli system syslog config logfilter list`.

Use the `set` command to activate or deactivate log filtering: `[root@xxx-xx-dhcp-xx-xx:~] esxcli system syslog config logfilter set`.

Use the `add` command to add a log filter and the `remove` command to remove a log filter.

Use the `get` command to determine if log filtering is enabled.

A log filter is specified by three components and uses the following syntax: `numLogs | ident | logRegexp`.

Parameter	Description
numLogs	Specifies the number of matches of the logRegexp Python regular expression that will be allowed before filtering begins.

Parameter	Description
<code>ident</code>	The <code>ident</code> string is how an application identifies itself to the syslog facility. The <code>logRegex</code> filter must be associated with the same application. You can find the <code>ident</code> string of an application by inspecting the log files in <code>/var/run/log</code> . The third field of each log file begins with the <code>ident</code> string and ends with <code>].</code>
<code>logRegex</code>	Python regular expression that identifies the messages which you want to filter out.

For example, to filter out all messages from the `hostd` daemon that contain the word "mark" after the tenth occurrence, use the following command: `esxcli system syslog config logfilter add --filter="10|Hostd|mark"`.

To remove the log filter, use the command `esxcli system syslog config logfilter remove --filter="10|Hostd|mark"`.

For more information, see [ESXi Syslog Options](#).

Turn Off Compression for vpxd Log Files

By default, vCenter Server `vpxd` log files are rolled up and compressed into `.gz` files. You can turn off this setting to leave the `vpxd` logs uncompressed.

1. Log in to the vCenter Server using the vSphere Client.
2. Select **Administration > vCenter Server Settings**.
3. Select **Advanced Settings**.
4. In the **Key** text box, type `log.compressOnRoll`.
5. In the **Value** text box, type `false`.
6. Click **Add**, and click **OK**.

ESXi VMkernel Files

If the VMkernel fails, an error message appears and then the virtual machine reboots. If you specified a VMware core dump partition when you configured your virtual machine, the VMkernel also generates a core dump and error log.

More serious problems in the VMkernel can freeze the machine without an error message or core dump.

vSphere Single Host Management - VMware Host Client

vSphere Single Host Management - VMware Host Client provides information about managing single hosts with the VMware Host Client.

The VMware Host Client can be used to conduct emergency management when vCenter Server is unavailable. You can use the VMware Host Client to perform administrative tasks, basic troubleshooting tasks, and advanced administrative tasks.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we have updated this guide to remove instances of non-inclusive language.

Intended Audience

This information is intended for anyone who wants to use the VMware Host Client to manage single ESXi hosts. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

What is VMware Host Client

The VMware Host Client is an HTML5-based client that is used to connect to and manage single ESXi hosts.

You can use the VMware Host Client to:

- Perform administrative and basic troubleshooting tasks, and advanced administrative tasks on your target ESXi host.
- Conduct an emergency management when vCenter Server is not available.

It is important to know that the VMware Host Client is different from the vSphere Client. You use the vSphere Client to connect to vCenter Server and manage multiple ESXi hosts, whereas you use the VMware Host Client to manage a single ESXi host.

VMware Host Client functions include, but are not limited to the following operations:

- Basic virtualization operations, such as deploying and configuring virtual machines of various complexity
- Creating and managing networking and datastores
- Advanced tuning of host level options to improve performance

VMware Host Client System Requirements

Make sure that your browser supports the VMware Host Client.

VMware Host Client supports the following guest operating systems and Web browser versions.

Supported Browsers	Mac OS	Windows 32-bit and 64-bit	Linux
Google Chrome	89+	89+	75+
Mozilla Firefox	80+	80+	60+
Microsoft Edge	90+	90+	N/A
Safari	9.0+	N/A	N/A

Use the VMware Host Client

You can use the VMware Host Client to conduct an emergency management when vCenter Server is temporarily unavailable.

Start the VMware Host Client and Log In

You can use the VMware Host Client to manage single ESXi hosts and perform various administrative and troubleshooting tasks on your virtual machines.

To log in to the ESXi host, perform the following steps.

1. In a Web browser enter the target host name or IP address using the form `https://host-name/ui` or `https://host-IP-address/ui`.
A log in screen appears.
2. Enter your user name and your password.
3. Click **Login** to continue.
4. Review the VMware Customer Experience Improvement Program (CEIP) page and choose whether you want to join the program.
To learn about the program and how to configure it at any time, see [Leave and Rejoin the Customer Experience Improvement Program in the VMware Host Client](#).
5. Click **OK**.

Log Out of the VMware Host Client

When you no longer need to view or manage your target ESXi host, log out of the VMware Host Client.

Note: Closing a VMware Host Client session does not stop the host.

To log out of the ESXi host, click the user name at the top of the VMware Host Client window and select **Log out** from the drop-down menu.

You are now logged out of the VMware Host Client. Your target ESXi host continues to run all its normal activities.


How to Customize Your VMware Host Client UI Theme

With vSphere 8.0, you can customize the branding of the VMware Host Client UI and the way it looks and displays content.

You can choose among three preset themes - light, dark, and classic, and apply a theme to the VMware Host Client, based on your preferences.

1. On the VMware Host Client toolbar, click **Help** and click **About**.
The **About** window appears.

2. From the **UI Preferences Theme** drop-down menu, select the theme that you want to apply.
3. To change the theme name and modify up to 10 parameters of the theme you select, click the **Customize** button.



- a) In the **Theme name** field, enter a custom name of the theme.
- b) To select a custom color for each parameter, click the colored box in front of each parameter, select a color, and click **Enable**.
- a) To revert to the default palette, click the **Reset** button.

Configure a Login Banner for the UI Login Screen of the VMware Host Client

To display legal warnings or official announcements, you can configure a login page banner by using a limited form of markdown syntax.

~~By modifying the text file `/etc/vmware/welcome` directly on the host, you can change the login banner content that appears on the right of the user name and password login fields.~~

Note: A markdown parser is applied to the content block and certain character sequences, such as #, ` , * , might unintentionally trigger markdown formatting rules.

You can apply the following limited set of markdown directives in the `welcome` file.

Layout Concept	Markdown Code Syntax	Output
Heading labels	Starting on a new line, enter between 1 and 6 hash mark symbols. Example <code># My Title.</code>	Generates a large HTML <code><H1></code> tag for "My Title". My Title
Horizontal rule	Starting on a new line, enter a series of at least 3 dash characters only. Example <code>-----</code>	Generates a <code><hr /></code> rule tag in HTML.
Literal or code block	- Starting on a new line, enter 3 backtick characters only. - Add the source material on subsequent lines. - To close the source, on a new line enter 3 backtick characters. Example <pre> <code>''' My content - - - *Login Secure* >_ Read the policy '''</code> </pre>	Displays the text block between the backtick lines with no formatting or interpretation in a monospace font. <pre> My content - - - *Login Secure* >_ Read the policy </pre> Note: If the content is accidentally formatted by the markdown parser, wrap content with a pair of backtick characters lines. Since the spaces characters are preserved, ASCII art can be used as a monospace font is used.
Bold text	Wrap a string of text with double asterisk characters on both sides. Example <code>**important message**.</code> Note: The markdown double underscore character syntax is omitted to avoid conflicts with URLs.	<code> important message</code>
Italic text	Wrap a string of text with single asterisk characters on both sides. Example <code>*A named document*.</code> Note: The markdown underscore character syntax is omitted to avoid conflicts with URLs.	<code><i> A named document</i></code>
Hyperlink	To link an absolute URL, use the markdown syntax of square brackets enclosing the link text followed by wrapped parentheses for the URL. Example <code>[My link] (https://www.example.com?search=virtual)</code>	Generates a normal hyperlink anchor tag with clickable text. <code>My link</code>

Supported variables

You can insert the following variables anywhere within the text file.

Variable Concept	Meta Tag Variable Code	Output
Fully qualified domain name of the current host or IP address	{hostname}	Displays the full name of the current host. For example, sample.host.com
ESXi version as a dotted numeric format	{esxversion}	Displays, for example, 7.0.0
ESXi full product name, version, and build number	{esxproduct}	Displays, for example, VMware ESXi 7.0.0 build-16324942
Current date on the user's machine	{client-current-date}	Displays, for example, Tuesday, August 30, 2022 Note: This is locale specific.
Current time on the user's machine	{client-current-time}	Displays, for example, 08:00 AM Note: This is locale specific.

Advanced tags

The advanced tags offer visual and behavioral changes depending on the rules that you apply to the login page. Inset these tags at the very end of the text file.

UI concept	Meta tag variable code	Output
Custom image	{logo src="https://site/logo.png" width="100" height="100" align="center"}My Secured ESXi Server Tooltip{/logo}	Displays a 100x100 px image logo.png that is centered horizontally above the message block. An accessible tooltip title of My Secured ESXi Server Tooltip is added to the image. Note: Make sure to use the width, height, and align attributes although they are optional. All web image formats are supported.
User agreement check box	{accept}Please accept the terms{/accept}	Displays a check box with the "Please accept the terms" label at the bottom of the message content.
Enforced acceptance error message	{mustaccept}You must agree before logging into the system{/mustaccept}	Adds form validation to require the user to select the check box before log in. If the user does not select the check box, the "You must agree before logging into the system" message appears above the login button.

Examples

Simple markdown

A one-line markdown for a simple text only message

```
## Warning: Authorized Users Only
```



vmware®

ESXi Host Client™

root

Password

LOGIN

War

Advanced markdown

An advanced markdown sample for a fictitious cloud storage company Vaulted with logo, links and a required accept check box on the form.

```
## Warning: Authorized Users Only
```

```
The information on this host is the property of ***Vaulted Storage*** *(sample organization)* and is protected under sovereign **intellectual property rights**.
```

```
You must be assigned an account on this computer to access information and are only allowed to access information defined by the system administrators.
```

```
*All activities* are monitored and **trespassing violators** will be reported to a federal law enforcement agency.
```

```
### Policy bulletins
```

```
*Please refer to the helpful links* below on end user protection guidelines.
```

```
* [Privacy addendum](https://en.wikipedia.org/wiki/Computer_security)
```

```
* [Terms of Use](https://en.wikipedia.org/wiki/Terms_of_service#:~:text=Terms%20of%20service%20(also%20known,to%20use%20the%20offered%20service.)
```

```
...
```

```

      ,-----,
    ,-----,
  ,-----,
,-----,
+-----+
| .----- . | | | |
| |         | | |
| |  ACCESS  | | |
| |  RESTRICTED! | | |
| |  login as: _ | | |
| |         | | |
| | `-----' | |,"
+-----+
  /_)_____(_/ //'
      /_____/
/  oooooooooooooooooo .o.  oooo /,
/  ==oooooooooooooooo==.o.  ooo= //
/_==_=====___==_ooo_ooo=_/'
`-----'

```

```
-----
Host: {hostname}
```

```
ESXi Server Version: {esxversion}
```

```
Product: {esxproduct}
-----
```

```
...
```

```
{logo align="center" width="200" height="200" src="https://i.postimg.cc/y6wZXTFm/vaulted-logo-white-text.png"}Vaulted Enterprise Storage{/logo}
```

```
{accept}I accept the terms of use policy{/accept}
```

```
{mustaccept}User must check terms of use to login. LOG OFF immediately if you do not agree to the conditions  
stated in the warning.{/mustaccept}
```



vmware®

ESXi Host Client™

LOGIN

Warning: Authorization

The information on this host is the property of VMware. All activities are monitored and trespassing is prohibited.

You must be assigned an account on this host to access system administrators.

All activities are monitored and trespassing is prohibited.

Policy bulletins

Please refer to the helpful links below.

- * [Privacy addendum](#)
- * [Terms of Use](#)

ACCESS
RESTRICTED!
login as: _

3486

Leave and Rejoin the Customer Experience Improvement Program in the VMware Host Client

You can participate in the Customer Experience Improvement Program (CEIP) to provide anonymous feedback or information to VMware for quality, reliability, and functionality improvements of VMware products and services.

You can choose to leave the Customer Experience Improvement Program (CEIP), or rejoin the CEIP at any time.

Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

1. To leave and rejoin the CEIP, click the user name at the top of the VMware Host Client page.
2. Select **Settings** > **Send usage statistics**, to leave or rejoin the CEIP.

Host Management with the VMware Host Client

With the VMware Host Client, you can manage single ESXi hosts during vCenter Server upgrades or when vCenter Server stops responding or becomes unavailable.

The VMware Host Client has a crucial set of troubleshooting functions, which allow you to perform tasks on the ESXi host that you are logged in to if vCenter Server is unavailable. These functions include but are not limited to configuring advanced host settings, licensing, managing certificates, using the ESXi Shell, enabling Lockdown mode, and so on.

Manage System Settings in the VMware Host Client

With the VMware Host Client, you can manage advanced host settings, assign or remove licenses to your host, configure start and stop policies for host services, and manage time and date configuration for the host.

Manage Advanced Settings in the VMware Host Client

You can change the settings of a host by using the VMware Host Client.

CAUTION: Changing advanced options is considered unsupported unless VMware technical support or a KB article instruct you to do so. In all other cases, changing these options is considered unsupported. In most cases, the default settings produce the optimum result.

1. Click **Manage** in the inventory and click **System**.
2. Click **Advanced settings**.
3. Right-click the appropriate item from the list and select **Edit option** from the drop-down menu.
The **Edit option** dialog box is displayed.
4. Edit the value and click **Save** to apply your changes.
5. Optional: Right-click the appropriate item from the list and select **Reset to default** to go back to the original settings of the item.

Create an Initial Welcome Message for the Direct Console User Interface and the VMware Host Client

By using the VMware Host Client, you can create a welcome message that appears on the initial screen of the Direct Console User Interface (DCUI) and on the login window of the VMware Host Client. You can also create a welcome message that appears after a user logs into the VMware Host Client and decide whether to display the welcome message.

1. Click **Manage** in the VMware Host Client inventory and click **Advanced Settings**.

Option	Action
Create a welcome message that appears before you log in to DCUI and VMware Host Client	1. Enter <code>Annotations.WelcomeMessage</code> in the Search text box and click the Search icon.

Option	Action
	- Right-click <code>Annotations.WelcomeMessage</code> and select Edit option from the drop-down menu. The Edit option dialog box opens. - In the New value text box, enter the welcome message. To set the default message, leave the New value text box blank.
Create a welcome message that appears after you log in to the VMware Host Client	- Enter <code>UserVars.HostClientWelcomeMessage</code> in the Search text box and click the Search icon. - Right-click <code>UserVars.HostClientWelcomeMessage</code> and select Edit option from the drop-down menu. The Edit option dialog box opens. - In the New value text box, enter the welcome message. To set the default message, leave the New value text box blank.
Activate or deactivate the appearance of the welcome message after you log in to the VMware Host Client	- Enter <code>UserVars.HostClientEnableMOTDNotification</code> in the Search text box and click the Search icon. - Right-click <code>UserVars.HostClientEnableMOTDNotification</code> and select Edit option from the drop-down menu. The Edit option dialog box opens. - In the New value text box, enter the new value. A value of zero (0) deactivates the appearance of the welcome message. A value of one (1) activates the appearance of the welcome message.

2. Click **Save**.

3. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

Configure the VMware Host Client User Interface Session Timeout

In VMware Host Client, the User Interface session automatically times out every 15 minutes and then you must log back in to the VMware Host Client.

You can increase the default inactivity timeout by changing an advanced configuration parameter. The default value is 900 seconds.

Configure the User Interface session timeout.

Option	Action
From the VMware Host Client Advanced Settings	- Click Manage in the VMware Host Client inventory and click Advanced Settings. - Enter <code>UserVars.HostClientSessionTimeout</code> in the Search text box and click the Search icon. - Right-click <code>UserVars.HostClientSessionTimeout</code> and select Edit option from the drop-down menu. The Edit option dialog box opens. - In the New value text box, enter the timeout setting in seconds. Note: A value of zero (0) deactivates the timeout. - Click Save. - To reset the key setting to default, right-click the appropriate key from the list and select Reset to default.
From the User Settings drop-down menu	- Click the user name at the top of the VMware Host Client window and select Settings > Application timeout > . - To specify the inactivity timeout, select the time. - To deactivate the inactivity timeout, select Off.

Configure the SOAP Session Timeout in the VMware Host Client

In VMware Host Client you can configure the SOAP session timeout.

1. Click **Manage** in the VMware Host Client inventory and click **Advanced Settings**.
2. Enter `Config.HostAgent.vmacore.soap.sessionTimeout` in the **Search** text box and click the **Search** icon.
3. Right-click `Config.HostAgent.vmacore.soap.sessionTimeout` and select **Edit option** from the drop-down menu.
The **Edit option** dialog box opens.
4. In the **New value** text box, enter the timeout setting in seconds.
A value of zero (0) deactivates the timeout.
5. Click **Save**.
6. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

Configure the Passwords and Account Lockout Policy in the VMware Host Client

For ESXi hosts, you must use a password with predefined requirements. You can change the required password length, character class requirements, or allow passphrases, all using the `Security.PasswordQualityControl` advanced option. You can also set the number of passwords to remember for each user using the `Security.PasswordHistory` advanced option. This setting prevents duplicate or similar passwords. The `Security.PasswordMaxDays` advanced option allows you to set up the maximum number of days between password changes.

Note: Always perform additional testing after you change the default password settings.

If you attempt to log in with incorrect credentials, the account lockout policy specifies when and for how long the system locks your account.

ESXi Passwords

ESXi enforces password requirements for access.

- By default, when you create a password, you must include a mix of characters from any three of the following four character classes: lowercase letters, uppercase letters, numbers, and special characters such as underscore or dash.
- By default, the password must contain a length of at least 7 characters and a maximum of 40 characters.
- Passwords must not contain a dictionary word or part of a dictionary word.
- Passwords must not contain the user name or parts of the user name.

Note:

An uppercase character that begins a password does not count toward the number of character classes used. A number that ends a password does not count toward the number of character classes used.

The following password candidates illustrate potential passwords if the option is set as follows:

```
retry=3 min=disabled,disabled,disabled,7,7
```

With this setting, a user is prompted up to three times (`retry=3`) for a new password that is not sufficiently strong or if the password was not entered correctly twice. Passwords with one or two character classes and password phrases are not allowed, because

Example of ESXi Passwords

the first three items are deactivated. Passwords from three and four character classes require 7 characters.

The following password candidates meet the password requirements:

- xQaTEhb!: Contains eight characters from three character classes.
- xQaT3#A: Contains seven characters from four character classes.

The following password candidates do not meet the password requirements:

- Xqat3hi: Begins with an uppercase character, reducing the effective number of character classes to two. The minimum number of required character classes is three.
- xQaTEh2: Ends with a number, reducing the effective number of character classes to two. The minimum number of required character classes is three.

Password Quality Control

You can control the quality of passwords by using the Security.PasswordQualityControl advanced option.

Security.PasswordQualityControl consists of several settings that follow the pattern:

```
retry=N min=N0,N1,N2,N3,N4 max=N passphrase=N similar=permit|deny
```

Password Quality Control Settings	Description	Default
retry=N	The number of times the user must provide a new password if the password is incorrect or not sufficiently strong.	retry=3
min=N0,N1,N2,N3,N4	Character class and the passphrase minimum length requirement. • N0 is minimum length of passwords from a single character class. • N1 is minimum length of passwords from two character classes. • N2 is minimum length for a passphrase.	min=disabled,disabled,disabled,7,7

Password Quality Control Settings	Description	Default
	- N3 is minimum length for three character classes. - N4 is minimum length for four character classes. You can use disabled to disallow a password with the specified number of character classes.	
max=N	The maximum allowed password length.	max=40
passphrase=N	The number of words required for a passphrase. To make sure that the passphrase is recognized, do not set N2 from the min setting to disabled.	passphrase=3
similar=permit deny	Indicates whether a password is allowed to be similar to the old password. To use this setting, make sure that you set the Security.PasswordHistory option to a non-zero value. Starting with vSphere 8.0 Update 1, the default value is 5.	similar=deny

ESXi Passphrase

Instead of a password, you can use a passphrase. Passphrases are deactivated by default. You can change the default setting by using the Security.PasswordQualityControl advanced option.

For example, you can change the option to the following.

```
retry=3 min=disabled,disabled,16,7,7
```

This example allows passphrases of at least 16 characters. The passphrase must consist of at least 3 words, separated by spaces.

Example Password History and Rotation Policy

To remember history of 6 passwords, set the Security.PasswordHistory option to 6.

ESXi Account Lockout Policy

To enforce a 90 day password rotation policy, set the `Security.PasswordMaxDays` option to 90.

Users are locked out after a preset number of consecutive failed attempts. By default, users are locked out after 5 consecutive failed attempts in 3 minutes and a locked account is unlocked automatically after 15 minutes by default. You can change the maximum allowed failed attempts and the period of time in which the user account is locked out by using the `Security.AccountLockFailures` and `Security.AccountUnlockTime` advanced options.

To configure the administrator passwords and account lockout behaviour, perform the following steps.

1. Click **Manage** in the VMware Host Client inventory and click **Advanced Settings**.

Option	Action
Configure the required password length, character class requirement, or allow passphrases	1. Enter <code>Security.PasswordQualityControl</code> in the Search text box and click the Search icon. 2. Right-click <code>Security.PasswordQualityControl</code> and select Edit option from the drop-down menu.
Configure the number of passwords to remember for each user	1. Enter <code>Security.PasswordHistory</code> in the Search text box and click the Search icon. 2. Right-click <code>Security.PasswordHistory</code> and select Edit option from the drop-down menu. Note: Zero (0) deactivates password history.
Configure the maximum number of days between password changes	1. Enter <code>Security.PasswordMaxDays</code> in the Search text box and click the Search icon. 2. Right-click <code>Security.PasswordMaxDays</code> and select Edit option from the drop-down menu.
Configure the number of failed login attempts allowed before lockout	1. Enter <code>Security.AccountLockFailures</code> in the Search text box and click the Search icon. 2. Right-click <code>Security.AccountLockFailures</code> and select Edit option from the drop-down menu. Note: Zero (0) deactivates account locking.
Configure the period of time in which the user's account is locked out	1. Enter <code>Security.AccountUnlockTime</code> in the Search text box and click the Search icon. 2. Right-click <code>Security.AccountUnlockTime</code> and select Edit option from the drop-down menu.

The **Edit option** dialog box opens.

2. In the **New value** text box, enter the new setting.
3. Click **Save**.
4. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

Configure Syslog in the VMware Host Client

To configure the syslog service, you can use the VMware Host Client.

1. Click **Manage** in the VMware Host Client inventory and click **Advanced Settings**.
2. In the **Search** text box, enter the name of the setting that you want to change and click the **Search** icon.

Option	Description
Syslog.global.LogHost	Remote host to which syslog messages are forwarded and the port on which the remote host receives syslog messages. You can include the protocol and the port, for example, <code>protocol://hostName1:port</code> where <code>protocol</code> can be <code>udp</code> , <code>tcp</code> , or <code>ssl</code> . You can use only port 514 for UDP. The <code>ssl</code> protocol uses TLS 1.2. For example: <code>ssl://hostName1:1514</code> . The value of <code>port</code> can be any decimal number between 1 and 65535. While no hard limit to the number of remote hosts to receive syslog messages exists, it is recommended to keep the number of remote hosts to five or less.
Syslog.global.logCheckSSLCerts	Enforce checking of the SSL certificates when you log in to a remote host.
Syslog.global.defaultRotate	Maximum number of archives to keep. You can set this number globally and for individual subloggers.
Syslog.global.defaultSize	Default size of the log, in KB, before the system rotates logs. You can set this number globally and for individual subloggers.
Syslog.global.LogDir	Directory where logs are stored. The directory can be on mounted NFS or VMFS volumes. Only the <code>/scratch</code> directory on the local file system is persistent across reboots. Specify the directory as <code>[datastoreName] path_to_file</code> , where the path is relative to the root of the volume backing the datastore. For example, the path <code>[storage1] /systemlogs</code> maps to the path <code>/vmfs/volumes/storage1/systemlogs</code> .
Syslog.global.logDirUnique	Selecting this option creates a subdirectory with the name of the ESXi host under the directory specified by Syslog.global.LogDir . A unique directory is useful if the same NFS directory is used by multiple ESXi hosts.

3. Right-click the setting name and select **Edit option** from the drop-down menu.
The **Edit option** dialog box opens.
4. To perform the SSL certificates check when you log in to a remote host, click **True** from the **New value**.
5. Click **Save**.
6. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

Configure Advanced TLS/SSL Key Options

You can configure the security protocols and cryptographic algorithms that are used to encrypt communications with the ESXi host.

For more information, see the VMware KB article at <https://kb.vmware.com/s/article/79476>.

The Transport Layer Security (TLS) key secures communication with the host using the TLS protocol. Upon first boot, the ESXi host generates the TLS key as a 2048-bit RSA key. Currently, ESXi does not implement automatic generation of ECDSA keys for TLS. The TLS private key is not intended to be serviced by the administrator.

The SSH key secures communication with the ESXi host using the SSH protocol. Upon first boot, the system generates the SSH key as a 2048-bit RSA key. The SSH server is deactivated by default. SSH access is intended primarily for troubleshooting purposes. The SSH key is not intended to be serviced by the administrator. Logging in through SSH requires administrative privileges equivalent to full host control. To enable SSH access, see [Enable the Secure Shell \(SSH\) in the VMware Host Client](#).

You can configure the following ESXi host security key settings.

Note: The UserVars.ESXiVPsAllowedCiphers security key setting only affects I/O filters.

Key	Default	Description
UserVars.ESXiVPsAllowedCiphers	!aNULL:kECDH+AESGCM:EC DH+AESGCM:RSA+AESGCM:kECDH+AES:EC DH+AES:RSA+AES	The default cipher control string.
Config.HostAgent.ssl.keyStore.allowAny	False	You can add any certificate to the ESXi CA trust store.
Config.HostAgent.ssl.keyStore.allowSelfSigned	False	You can add non-CA self-signed certificates to the ESXi CA trust store, that is, certificates that do not have the CA bit set.
Config.HostAgent.ssl.keyStore.discardLeaf	True	Discards leaf certificates added to the ESXi CA trust store.

To configure the ESXi security key settings:

1. Click **Manage** in the VMware Host Client inventory and click **Advanced Settings**.
2. Enter the security key in the **Search** text box and click the **Search** icon.
3. Right-click the security key and select **Edit option** from the drop-down menu.
The **Edit option** dialog box opens.
4. In the **New value** field enter the new value and click **Save**.
5. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

Configure Userworld Memory Zeroing

With VMware Host Client, you can use the advanced option Mem.MemEagerZero to determine how pages are zeroed out for virtual machines and user space applications.

To zero all pages when they allocated to virtual machines and user space applications, set Mem.MemEagerZero to one (1). If the memory is not reused, this setting prevents exposing the information from a virtual machine or user space applications to other clients while preserving the previous content in memory.

When you set Mem.MemEagerZero to 1, pages are zeroed when a user space application exits. For virtual machines, such pages are zeroed if:

- The virtual machine is powered off.
- The virtual machine pages are migrated.
- The ESXi host reclaims virtual machines memory.

Note: For virtual machines, you can obtain this behaviour by setting the sched.mem.eagerZero advanced option to `TRUE`.

For information about setting the advanced virtual machine options, see the *vSphere Resource Management* documentation.

To configure the userworld memory zeroing, perform the following steps.

1. Click **Manage** in the VMware Host Client inventory and click **Advanced Settings**.
2. Enter `Mem.MemEagerZero` in the **Search** text box and click the **Search** icon.
3. Right-click `Mem.MemEagerZero` and select **Edit option** from the drop-down menu.
The **Edit option** dialog box opens.
4. In the **New value** text box, enter the new value.
The default value is zero (0).
5. Click **Save**.
6. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

Change Autostart Configuration in the VMware Host Client

Configure autostart options for the ESXi host to set up when the host starts and stops.

Important:

The global `Autostart` setting affects the `Autostart` setting of a virtual machine.

1. Click **Manage** in the inventory and click **System**.
2. Click **Autostart**.
3. Click **Edit settings**.
4. Select **Yes** to enable changing the autostart configuration.

Option	Description
Start delay	After you start the ESXi host, it starts powering on the virtual machines that are configured for automatic startup. After the ESXi host powers on the first virtual machine, the host waits for the specified delay time and then powers on the next virtual machine.
Stop delay	Stop delay is the maximum time the ESXi host waits for a shutdown command to complete. The order in which virtual machines are shut down is the reverse of their startup order. After the ESXi host shuts down the first virtual machine within the time that you specify, the host shuts down the next virtual machine. If a virtual machine does not shut down within the specified delay time, the host runs a power off command and then starts shutting down the next virtual machine. The ESXi host shuts down only after all virtual machines are shut down.
Stop action	Select a shut down action that is applicable to the virtual machines on the host when the host shuts down. • System default • Power off • Suspend • Shut down
Wait for heartbeat	Select Yes to enable the Wait for heartbeat option. You can use this option if the guest operating system of the virtual machine has VMware Tools installed. After the ESXi host powers on the first virtual machine, the host immediately powers on the next virtual machine. The startup order in which

Option	Description
	virtual machines are powered on continues after the virtual machine receives the first heartbeat.

If you set a delay option to -1 , the system uses the default option.

- Click **Save**.

Edit the Time Configuration of anESXi Host in theVMware Host Client

By using theVMware Host Client, you can configure the time settings of a host manually or can synchronize the time and date of the host with an NTP or a PTP server. NTP provides millisecond timing accuracy and PTP maintains microsecond timing accuracy.

The NTP service on the host periodically takes the time and date from the NTP server. You can use the **Start**, **Stop**, or **Restart** buttons to change the status of the NTP service on the host at any time regardless of the selected startup policy for the NTP service.

PTP provisions precise time synchronization for the virtual machines within a network. To change the PTP service on the host at any time, you can use the **Start**, **Stop**, or **Restart** buttons. Starting or stopping the PTP service automatically enables or disables PTP. To apply the change when you enable or disable PTP manually, start or stop the PTP service.

For more information about services, see [Manage Services in the VMware Host Client](#).

Note: The NTP and PTP services cannot run simultaneously.

- Click **Manage** in theVMware Host Client inventory.
- On the **System** tab, click **Time & date**.
- Set the time and date for the host.

Option	Action
Manually configure the date and time on this host	- Click Edit NTP Settings. The Edit NTP Settings dialog box appears. - Set the time and date for the host manually. - Click Save.
Use Network Time Protocol (Enable NTP client)	- Click Edit NTP Settings. The Edit NTP Settings dialog box appears. - Select the Use Network Time Protocol radio button. - In the NTP Servers text box, enter the IP addresses or host names of the NTP servers that you want to use. - From the NTP Service Startup Policy drop-down menu, select an option for starting and stopping the NTP service on the host. Start and stop with port usage. Starts or stops the NTP service when the NTP client port is enabled or disabled for access in the security profile of the host. Start and stop with host. Starts and stops the NTP service when the host powers on and shuts down. Start and stop manually. Enables manual starting and stopping of the NTP service. If you select the Start and stop manually policy, the status of the NTP service changes only when you use the UI controls. - Click Save.
Use Precision Time Protocol (Enable PTP client)	- Click Edit PTP Settings. - Select the Enable check box. - From the Network interface drop-down menu, select a network interface. The IPv4 and Subnet mask appear.

Option	Action
	4. Click Save .

Manage Hardware for an ESXi Host by Using the VMware Host Client

When you log in to an ESXi host by using the VMware Host Client, you can manage PCI devices and configure power management settings.

Host Power Management Policies in ESXi

You can apply several power management features in ESXi that the host hardware provides to adjust the balance between performance and power. You can control how ESXi uses these features by selecting a power management policy.

Selecting a high-performance policy provides more absolute performance, but at lower efficiency and performance per watt. Low-power policies provide less absolute performance, but at higher efficiency.

You can select a policy for the host that you manage by using the VMware Host Client. If you do not select a policy, ESXi uses Balanced by default.

Table 559: CPU Power Management Policies

Power Management Policy	Description
High Performance	Do not use any power management features.
Balanced (Default)	Reduce energy consumption with minimal performance compromise
Low Power	Reduce energy consumption at the risk of lower performance
Custom	User-defined power management policy. Advanced configuration becomes available.

When a CPU runs at lower frequency, it can also run at lower voltage, which saves power. This type of power management is typically called Dynamic Voltage and Frequency Scaling (DVFS). ESXi attempts to adjust CPU frequencies so that virtual machine performance is not affected.

When a CPU is idle, ESXi can apply deep halt states, also known as C-states. The deeper the C-state, the less power the CPU uses, but it also takes longer for the CPU to start running again. When a CPU becomes idle, ESXi applies an algorithm to predict the idle state duration and chooses an appropriate C-state to enter. In power management policies that do not use deep C-states, ESXi uses only the shallowest halt state for idle CPUs, C1.

Change Power Management Policies in the VMware Host Client

Change the power management policies of the host that you are managing to control the energy consumption of your host.

1. Click **Manage** in the VMware Host Client inventory and click **Hardware**.
2. Click **Power Management** and click **Change policy**.
The available power management policies are displayed.

3. Select the policy that you want to apply and click **OK**.

Change the Hardware Label in the VMware Host Client

In the VMware Host Client, you can change the hardware label of all available PCI passthrough devices on a virtual machine. You use hardware labels to restrict the virtual machine placement to specific hardware instances. You can add all available devices with the same hardware label or with a blank hardware label to a virtual machine.

1. Click **Manage** in the VMware Host Client inventory.
2. On the **Hardware** tab, click **PCI Devices**.
3. Select an available device from the list and click **Hardware label**.
The toggle passthrough must be active for the selected device.
The **Edit Hardware Label** dialog box appears.
4. Edit the hardware label and click **Save** to apply the changes.

The new hardware label appears in the hardware label column.

Licensing for ESXi Hosts

ESXi hosts are licensed with vSphere licenses. Each vSphere license has a certain capacity that you can use to license multiple physical CPUs on ESXi hosts.

There are four main licensing models for vSphere:

- Per core licensing with a minimum of 16 cores licensed per CPU.
- Per virtual machine licensing.
- vSphere+ subscription capacity based licensing.
- Per CPU licensing that covers one CPU with up to 32 cores.

To license an ESXi host, you must assign to it a vSphere license that meets the following prerequisites:

- The license must have sufficient capacity depending on the licensing model.
- The license must support all the features that the host uses. For example, if the host is associated with a vSphere Distributed Switch, the license that you assign must support the vSphere Distributed Switch feature.

If you attempt to assign a license that has insufficient capacity or does not support the features that the host uses, the license assignment fails.

Per Core Licensing Model for vSphere

The per core licensing model is subscription based. To calculate the capacity you need for your environment, you need the total number of the physical CPU cores for each CPU on all ESXi hosts in your environment. Each core requires a single license, and the minimum license capacity you can purchase is 16 cores per CPU.

For example, if you have 1 ESXi host in your inventory, with 1 CPU, and 8 CPU cores per CPU, you must purchase the subscription capacity of 16 cores per CPU because it is the minimum license capacity.

Number of ESXi Hosts	Number of CPUs	Cores per CPU	Number of Core Licenses
1	1	8	16
2	2	8	64
2	2	16	64

For more information about calculating the number of licenses you need for your environment, see the VMware knowledge base article at <https://kb.vmware.com/s/article/95927>.

Per Virtual Machine Licensing Model for vSphere

The license use equals the total number of powered on desktop virtual machines running on the hosts that are assigned such a license.

Subscription Based Licensing Model for vSphere

If you currently use the vSphere+ workload platform, you have a subscription capacity based model for management of vSphere. For more information about your subscription model, see the *VMware vSphere +* documentation.

Per CPU Licensing Model for vSphere

If you use the per CPU licensing model, one PU license covers one CPU with up to 32 cores. If a CPU has more than 32 cores, you need additional CPU licenses.

Number of CPUs	Cores per CPU	Number of CPU Licenses
1	1-32	1
2	1-32	2
1	33-64	2
2	33-64	4

When you assign a vSphere license to a host, the amount of capacity consumed is determined by the number of physical CPUs on the host and the number of physical cores in each physical CPU.

If you use the licensing model with up to 32 cores, you can assign a vSphere license for 10 32-core CPUs to any of the following combinations of hosts:

- Five 2-CPU hosts with 32 cores per CPU
- Five 1-CPU hosts with 64 cores per CPU
- Two 2-CPU hosts with 48 cores per CPU and two 1-CPU hosts with 20 cores per CPU

Dual-core and quad-core CPUs, such as Intel CPUs that combine two or four independent CPUs on a single chip, count as one CPU.

Evaluation Mode for ESXi Hosts

After you install ESXi, it operates in evaluation mode for up to 60 consecutive days. Evaluation mode provides all features of vSphere.

After you assign a license to an ESXi host, at any time before the evaluation period expires, you can set the host back to evaluation mode to explore the entire set of features available for the remaining evaluation period.

For example, if you use an ESXi host in evaluation mode for 20 days, then assign a vSphere Standard license to the host, and 5 days later set the host back to evaluation mode, you can explore the entire set of features available for the host for the remaining 35 days of the evaluation period.

License and Evaluation Period Expiry for ESXi Hosts

For ESXi hosts, license or evaluation period expiry leads to disconnection from vCenter Server. All powered on virtual machines continue to work, but you cannot power on virtual machines after they are powered off. You cannot change the current configuration of the features that are in use. You cannot use the features that remained unused before the license expiration.

Note:

When there are expiring licenses, a notification appears 90 days before the license expiration.

Licensing ESXi Hosts After an Upgrade

If you upgrade an ESXi host to a version that starts with the same number, you do not need to replace the existing license with a new one. For example, if you upgrade a host from ESXi 8.0 to 8.1, you can use the same license for the host.

If you upgrade an ESXi host to a major version that starts with a different number, the evaluation period restarts and you must assign a new license. For example, if you upgrade an ESXi host from 7.x to 8.x, you must license the host with a vSphere 8 license.

View Licensing Information About the VMware Host Client Environment

You can view the available licenses in the VMware Host Client along with their expiration dates, license key, and various features. You can also view the available products and assets.

Click **Manage** in the VMware Host Client inventory and click **Licensing**.

You can view the license key, the expiration date, and all the available features and assets.

Assign a License Key to an ESXi Host in the VMware Host Client

By using the VMware Host Client, you can assign an existing or new license key to an ESXi host.

Verify that you have the **Global > Licenses** privilege.

Note: If you use vCenter Server to manage your ESXi host, you can only change your licenses from the vSphere Client.

1. Click **Manage** in the VMware Host Client inventory and click **Licensing**.
2. Click **Assign license**, enter a license key in the form `XXXXX-XXXXX-XXXXX-XXXXX-XXXXX`, and click **Check license**.
3. Click **Assign license** to save your changes.

Remove a License from an ESXi Host in the VMware Host Client

To remain in compliance with the licensing models of products that you use with vSphere, you must remove all unassigned licenses from the inventory.

If you have divided, combined, or upgraded licenses in Customer Connect, you must remove the old licenses.

For example, suppose that you have upgraded a vSphere license from 6.5 to 6.7 in Customer Connect. You assign the license to ESXi 6.7 hosts. After assigning the new vSphere 6.7 licenses, you must remove the old vSphere 6.5 license from the inventory.

1. Click **Manage** in the VMware Host Client inventory and click **Licensing**.
2. Click **Remove license**, and click **OK**.

Manage Services in the VMware Host Client

In the VMware Host Client, you can start, stop, and restart services that are running on the host that you are logged in to, and you can configure the host service policy.

You can restart services when you change host configurations or in case of suspected functional or performance issues.

1. Click **Manage** in the VMware Host Client inventory and click **Services**.
2. From the **Services** list, select a service.
3. From the **Actions** drop-down menu, select an operation.
 - **Restart**
 - **Start**
 - **Stop**
4. Optional: From the **Actions** drop-down menu, select **Policy** and select an option for the service from the menu.
 - **Start and stop with firewall ports**
 - **Start and stop with host**
 - **Start and stop manually**

Managing Security and Users for an ESXi Host by Using the VMware Host Client

The ESXi hypervisor architecture has many built-in security features that you can configure to improve security.

By using the VMware Host Client, you can configure features, such as active directory, and you can also manage certificates.

Manage Host Authentication by Using the VMware Host Client

When you log in to an ESXi host by using the VMware Host Client, you can check whether active directory and smart card authentication are enabled, and you can also join the host to a directory service domain.

Join an ESXi Host to a Directory Service Domain by Using the VMware Host Client

To use a directory service for your host, you must join the host to the directory service domain.

You can enter the domain name in one of two ways:

- `name.tld` (for example, `domain.com`): The account is created under the default container.
- `name.tld/container/path` (for example, `domain.com/OU1/OU2`): The account is created under a particular organizational unit (OU).

To use the vSphere Authentication Proxy service, see *vSphere Security*.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Authentication** and click **Join domain**.
3. Enter a domain name.

Use the form `name.tld` or `name.tld/container/path`.

4. Enter the user name and password of a directory service user account that has permissions to join the host to the domain and click **Join domain**.
5. Optional: If you intend to use an authentication proxy, enter the proxy server IP address and click **Join domain**.

Using Active Directory to Manage ESXi Users

You can configure ESXi to use a directory service such as Active Directory to manage users.

Creating local user accounts on each host presents challenges with having to synchronize account names and passwords across multiple hosts. Join ESXi hosts to an Active Directory domain to eliminate the need to create and maintain local user accounts. Using Active Directory for user authentication simplifies the ESXi host configuration and reduces the risk for configuration issues that could lead to unauthorized access.

When you use Active Directory, users supply their Active Directory credentials and the domain name of the Active Directory server when adding a host to a domain.

Using vSphere Authentication Proxy

You can add ESXi hosts to an Active Directory domain by using vSphere Authentication Proxy instead of adding the hosts explicitly to the Active Directory domain.

You only have to set up the host so it knows about the domain name of the Active Directory server and about the IP address of vSphere Authentication Proxy. When vSphere Authentication Proxy is enabled, it automatically adds hosts that are being provisioned with Auto Deploy to the Active Directory domain. You can also use vSphere Authentication Proxy with hosts that are not provisioned by using Auto Deploy.

Auto Deploy

If you are provisioning hosts with Auto Deploy, you can set up a reference host that points to Authentication Proxy. You then set up a rule that applies the reference host's profile to any ESXi host that is provisioned with Auto Deploy. vSphere Authentication Proxy stores the IP addresses of all hosts that Auto Deploy provisions using PXE in its access control list. When the host boots, it contacts vSphere Authentication Proxy, and vSphere Authentication Proxy joins those hosts, which are already in its access control list, to the Active Directory domain.

Even if you use vSphere Authentication Proxy in an environment that uses certificates that are provisioned by VMCA or third-party certificates, the process works seamlessly if you follow the instructions for using custom certificates with Auto Deploy.

You can set up other hosts to use vSphere Authentication Proxy if you want to make it possible for the host to join the domain without using Active Directory credentials. That means you do not need to transmit Active Directory credentials to the host, and you do not save Active Directory credentials in the host profile.

In that case, you add the host's IP address to the vSphere Authentication Proxy access control list, and vSphere Authentication Proxy authorizes the host based on its IP address by default. You can enable client authentication to have vSphere Authentication Proxy check the host's certificate.

Other ESXi Hosts

Note: IPv6 is supported across vSphere. It is not supported on vSphere with Tanzu.

Manage Host Certificates by Using the VMware Host Client

When you log in to an ESXi host by using the VMware Host Client, you can view the certificate details of your host, such as the issuer and the validity period, and you can also import new certificates.

View Certificate Details for an ESXi Host in the VMware Host Client

You can use the certificate information for debugging.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Certificates**.

You can view the following certificate details.

Field	Description
Issuer	The issuer of the certificate.
Not valid after	Date on which the certificate expires.
Not valid before	Date on which the certificate is generated.
Subject	The subject used during certificate generation.

Import a New Certificate for an ESXi Host in the VMware Host Client

You can import a certificate from a trusted certificate authority when you are logged in to an ESXi host with the VMware Host Client.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Certificates** and click **Import new certificate**.
3. Generate a certificate signing request:

Option	Description
Generate FQDN signing request	• Click Generate FQDN signing request, click the Copy to clipboard button, and click Close. • To generate the signed certificate, pass the certificate signing request to the certificate authority (CA). • In the Certificate text box, paste the generated signed certificate in PEM format and click Import.
Generate IP signing request	• Click Generate IP signing request, click the Copy to clipboard button, and click Close. • To generate the signed certificate, pass the certificate signing request to the CA. • In the Certificate text box, paste the generated signed certificate in PEM format and click Import.

You do not have to import the certificate immediately. To make sure that you can use the signed certificate, do not restart the host between generating the certificate signing request and importing the certificate.

The certificate signing request is then passed to the certificate authority to generate the official certificate.

An FQDN request has the fully qualified host name of the host in the resulting common name field of the certificate. The IP signing request has the current IP address of the host in the common name field.

Manage Users with the VMware Host Client

Manage users to control who is authorized to log in to ESXi.

Users and roles control who has access to the ESXi host components and what actions each user can perform.

In vSphere 5.1 and later, ESXi user management has the following caveats .

- The users created when you connect directly to an ESXi host are not the same as the vCenter Server users. When the host is managed by vCenter Server, vCenter Server ignores users created directly on the host.
- You cannot create ESXi users by using the vSphere Client. You must log in to the host directly with the VMware Host Client to create ESXi users .
- ESXi 5.1 and later does not support local groups. However, Active Directory groups are supported.

To prevent anonymous users, such as root, from accessing the host with the Direct Console User Interface (DCUI) or ESXi Shell, remove the user's administrator privileges on the root folder of the host. This applies to both local users and Active Directory users and groups.

Add an ESXi User in the VMware Host Client

Adding a user to the users table updates the internal user list that the host maintains.

For information about password requirements, see [Configure the Passwords and Account Lockout Policy in the VMware Host Client](#) or the *vSphere Security* documentation.

1. Log in to ESXi with the VMware Host Client.
You cannot create ESXi users with the vSphere Client. To create ESXi users, you must directly log in to the host with the VMware Host Client.
2. Click **Manage** in the inventory and click **Security & Users**.
3. Click **Users**.
4. Click **Add user**.
5. Enter a user name, and a password.

Note: Do not create a user named `ALL`. Privileges associated with the name `ALL` might not be available to all users in some situations. For example, if a user named `ALL` has Administrator privileges, a user with the **ReadOnly** privileges might be able to log in to the host remotely. This is not the intended behavior.

- Do not include any spaces in the user name.
 - Do not include any non-ASCII characters in the user name.
 - Create a password that meets the length and complexity requirements. The host checks for password compliance using the default authentication plug-in, `pam_passwdqc.so`. If the password is not compliant, an error message indicates password requirements.
6. To activate the local access to the ESXi Shell, select the **Enable Shell Access** check box.
 7. Click **Add**.

Update an ESXi User in the VMware Host Client

You can change the description and password for an ESXi user in the VMware Host Client.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Users**.
3. Select a user from the list and click **Edit user**.
4. Update the user details and click **Save**.

Remove a Local ESXi User from a Host in the VMware Host Client

You can remove a local ESXi user from the host.

CAUTION: Do not remove the root user.

If you remove a user from the host, they lose permissions to all objects on the host and cannot log in again.

Note: Users who are logged in and are removed from the domain keep their host permissions until you restart the host.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Users**.
3. Select the user that you want to remove from the list, click **Remove user**, and click **Yes**.
Do not remove the root user for any reason.

Manage ESXi Roles in the VMware Host Client

ESXi grants access to objects only to users who are assigned permissions for the object. When you assign a user permission for the object, you do so by pairing the user with a role.

A role is a predefined set of privileges. For more information about privileges, see the *vSphere Security* documentation.

ESXi hosts provide three default roles, and you cannot change the privileges associated with these roles. Each subsequent default role includes the privileges of the previous role. For example, the Administrator role inherits the privileges of the Read Only role. Roles that you create do not inherit privileges from any of the default roles.

You can create custom roles by using the role-editing functions in the VMware Host Client to create privilege sets that match your user needs. Also, the roles you create directly on a host are not accessible in vCenter Server. You can work with these roles only if you log in to the host directly from the VMware Host Client.

Note: When you add a custom role and do not assign any privileges to it, the role is created as a read-only role with the **System.Anonymous**, **System.View**, and **System.Read** system-defined privilege.

If you manage an ESXi host through vCenter Server, maintaining custom roles in the host and vCenter Server can result in confusion and misuse. In this type of configuration, maintain custom roles only in vCenter Server.

You can create host roles and set permissions through a direct connection to the ESXi host with the VMware Host Client.

Add a Role in the VMware Host Client

You can create roles to suit the access control needs of your environment.

Verify that you are logged in as a user with Administrator privileges, such as root or vpxuser.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Roles**.
3. Click **Add role**.
4. Enter a name for the new role.
5. Select privileges from the list to associate with the new role and click **Add**.

Update a Role in the VMware Host Client

When you edit a role, you can change the privileges selected for that role. When complete, these privileges are applied to any user or group that is assigned the edited role.

Verify that you are logged in as a user with Administrator privileges, such as root or vpxuser.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Roles**.
3. Select a role from the list and click **Edit role**.
4. Update the role details and click **Save**.

Remove a Role in the VMware Host Client

When you remove a role that is not assigned to any users or groups, the definition is removed from the list of roles. When you remove a role that is assigned to a user or group, you can remove assignments or replace them with an assignment to another role.

Verify that you are logged in as a user with Administrator privileges, such as root or vpxuser.

CAUTION: You must understand how users will be affected before removing all assignments or replacing them. Users who have no permissions granted to them cannot log in.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Roles**.
3. Select the name of the role that you want to remove from the list.
4. Click **Remove role**, select **Remove only if unused**, and click **Yes**.

Managing Hosts in vCenter Server

To monitor all hosts in your vSphere environment from one place and to simplify host configuration, connect the hosts to a vCenter Server system.

For information about configuration management of ESXi hosts, see the *vSphere Networking* documentation, the *vSphere Storage* documentation, and the *vSphere Security* documentation.

Update Your VMware Host Client Environment to the Latest Version

To evaluate whether you are using the latest version of the VMware Host Client, check what VIBs are installed to your environment and examine the VIBs version information. You can update your VMware Host Client environment by entering a URL or a datastore path to either a VIB or the `metadata.zip` file in an ESXi offline bundle package.

If you provide a VIB file, an existing VIB that is installed to your VMware Host Client environment is updated to the new VIB.

If you provide an offline bundle, you update the entire ESXi host to the version described by the `metadata.zip` file in the bundle. Make sure that the entire offline bundle is available through the URL or is uploaded to the datastore.

To update your environment to the latest version, perform the following tasks:

Task	Steps
Upload a VIB to a datastore	1. Click Storage from the VMware Host Client environment. 2. Select a datastore from the list and click Datastore browser. 3. To store the VIB, select a directory and click Upload. 4. Browse to and double-click the file.
Upload an offline bundle to a datastore	1. Download the ESXi offline bundle package. 2. Upload the ESXi offline bundle package to the ESXi host. You can either upload the offline bundle package by using the Datastore browser or by using SCP or WinSCP. 3. Extract the contents of the offline bundle on the ESXi host. For example, log in to the host by using SSH.

Task	Steps
	- Navigate to the directory where you uploaded the offline bundle. - Extract the contents by using the <code>unzip</code> command.
Update your environment	- Click Manage in the VMware Host Client and click Packages. - Click Install update and enter the URL or the datastore path to a VIB or a <code>metadata.zip</code> file in an offline bundle. - Click Update. CAUTION: If you update an ESXi host managed by vSphere Lifecycle Manager, the host might become non-compliant. - Click Refresh to make sure that the update is successful.

Unable to Connect from the VMware Host Client to an ESXi Host After Upgrading to a Newer Version of ESXi

After you upgrade your ESXi host to a newer version, your browser console might display an error message when you attempt to access the ESXi host by using the VMware Host Client, and your connection might fail.

After you upgrade your ESXi host to a newer version, attempting to navigate to `https://host-fqdn/ui` or `https://1.2.3.4/ui` might result in the following error:

```
503 Service Unavailable (Failed to connect to endpoint:
[N7Vmacore4Http16LocalServiceSpecE:0xffa014e8] _serverNamespace = /ui _isRedirect = false
_port = 8308)
```

A change to `/etc/vmware/rhttpproxy/endpoints.conf` remains after an upgrade and causes the `/ui` endpoint to override the VMware Host Client.

When the `/ticket` is missing from the `endpoint.conf` file on your 6.0 or later ESXi host, your in-browser virtual machine console displays a `Failed to connect` error message but the continues to work.

- Log in to your ESXi host either by using SSH or ESXi Shell.
If you use SSH, you might need to enable SSH first. You can enable SSH by using DCUI.
- Back up the `endpoints.conf` file.

```
cp /etc/vmware/rhttpproxy/endpoints.conf /tmp
```
- Open the `/etc/vmware/rhttpproxy/endpoints.conf` file in an editor and remove the following line.

```
/ui local 8308 redirect allow
```
- Restart the `rhttpproxy` configuration management server.

```
/etc/init.d/rhttpproxy restart
```


5. Access the VMware Host Client using the host's specified full name in the secure URL with `https://host-fqdn/ui` or a valid numeric IP address `https://1.2.3.4/ui`.

Switch to the vSphere Client

To access the full set of capabilities, and advanced administrative and troubleshooting functions of the ESXi host, connect the ESXi host to vCenter Server.

1. Right-click **Host** in the VMware Host Client inventory and select **Manage with vCenter Server** from the drop-down menu.
The vCenter Server login page opens in a new window.
2. Enter your credentials and click **Login**.

Disconnect an ESXi Host from vCenter Server by Using the VMware Host Client

If you no longer want to use the advanced set of capabilities available through vCenter Server for host management, or if vCenter Server has failed and you must perform emergency operations on the host, you can disconnect your ESXi host from vCenter Server.

Disconnecting an ESXi host might take up to several minutes.

1. Right-click **Host** in the VMware Host Client inventory and select **Disconnect from vCenter Server** from the pop-up menu.
Note: Disconnecting a host signals vCenter Server that this host is not responding.
2. Click **Disconnect from vCenter Server**.

Reboot or Shut Down an ESXi Host in the VMware Host Client

You can power off or restart any ESXi host by using the VMware Host Client. Powering off a managed host disconnects it from vCenter Server, but does not remove it from the inventory.

To be able to reboot or shut down a host, you need these privileges.

- **Host > Configuration > Maintenance**
- **Global > Log event**

Always perform the following tasks before you reboot or shut down a host:

- Power off all virtual machines on the host.
- Place the host in maintenance mode.

1. Right-click the host, select **Shut down host** or **Reboot host**.

Note: If the host is not in maintenance mode, shutting down or rebooting it does not stop the virtual machines that are running on this host safely and unsaved data may be lost. If the host is part of a vSAN cluster, you might lose access to the vSAN data on the host.

2. Click **Shut down** or **Reboot** to complete the procedure.

Using the ESXi Shell

The ESXi Shell provides essential maintenance commands and is deactivated by default on ESXi hosts. You can activate local and remote access to the shell if necessary. To reduce the risk of unauthorized access, activate the ESXi Shell for troubleshooting only.

The ESXi Shell is independent of lockdown mode. Even if the host is running in lockdown mode, you can still log in to the ESXi Shell if it is activated.

The applicable services are as follows.

ESXi Shell SSH

Activate this service to access the ESXi Shell locally.
Activate this service to access the ESXi Shell remotely by using SSH.

The root user and users with the Administrator role can access the ESXi Shell. Users who are in the Active Directory group ESX Admins are automatically assigned the Administrator role. By default, only the root user can run system commands (such as `vmware -v`) by using the ESXi Shell.

Note: Do not activate the ESXi Shell unless you actually need access.

Enable the Secure Shell (SSH) in the VMware Host Client

Enable the Secure Shell (SSH) to access the ESXi Shell remotely by using SSH.

1. To enable or deactivate the Secure Shell (SSH), right-click **Host** in the VMware Host Client inventory.
2. Select **Services** from the drop-down menu.
3. To enable the Secure Shell (SSH), select **Enable Secure Shell (SSH)**.
4. To enable the ESXi Shell, select **Enable ESXi shell**.

Enable the ESXi Console Shell in the VMware Host Client

When you enable this service while running in lockdown mode, you can log in locally to the direct console user interface as the root user and deactivate lockdown mode. You can then access the host using a direct connection to the VMware Host Client or by enabling the ESXi Shell.

1. To activate or deactivate the Console Shell, right-click **Host** in the VMware Host Client inventory.
2. Select **Services** from the drop-down menu and select **Console Shell**.
3. Select a task to perform.
 - If the Console Shell is activated, click **Disable** to deactivate it.
 - If the Console Shell is deactivated, click **Enable** to activate it.

Create a Timeout for ESXi Shell Availability in the VMware Host Client

The ESXi Shell is deactivated by default. To increase security when you enable the shell, you can set an availability timeout for the ESXi Shell.

The availability timeout defines how long both local and remote shell logins are allowed before the ability to log in through the shell is deactivated. When the availability timeout expires, any existing shell sessions remains, but new shell sessions are not allowed.

1. Click **Manage** in the VMware Host Client inventory.
2. On the **System** tab, select **Advanced settings**.
3. Enter `UserVars.ESXiShellTimeOut` in the **Search** text box and click the **Search** icon.
4. Select `UserVars.ESXiShellTimeOut` and click **Edit option**.
The **Edit option** dialog box opens.
5. In the **New value** text box, enter the timeout setting.
A value of zero (0) deactivates the timeout.

6. Click **Save**.

You must restart the SSH service and the ESXi Shell service for the timeout to take effect.

7. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

Create a Timeout for Idle ESXi Shell Sessions in the VMware Host Client

If you enable the ESXi Shell on a host, but forget to log out of the session, the idle session remains connected indefinitely. The open connection increases the potential for someone to gain privileged access to the ESXi host. Prevent this by setting a timeout for idle sessions.

The idle timeout is the amount of time that can elapse before you are logged out of an idle interactive session.

1. Click **Manage** in the VMware Host Client inventory.
2. On the **System** tab, click **Advanced settings**.
3. Enter `UserVars.ESXiShellInteractiveTimeout` in the **Search** text box and click the **Search** icon.
4. Select `UserVars.ESXiShellInteractiveTimeout` and click **Edit option**.
The **Edit option** dialog box opens.
5. In the **New value** text box, enter the timeout setting.
A value of zero (0) deactivates the timeout.
6. Click **Save**.
The timeout takes effect only for newly logged in sessions.
7. Optional: To reset the key setting to default, right-click the appropriate key from the list and select **Reset to default**.

If the session is idle, users are logged out after the timeout period elapses.

Place an ESXi Host in Maintenance Mode in the VMware Host Client

You place an ESXi host in maintenance mode when you need to upgrade it, for example, to install more memory. A host enters or leaves maintenance mode only as the result of a user request.

Before you place a host in maintenance mode, power off all virtual machines that are running on that host or migrate them to another host either manually or automatically by DRS.

The host is in a state of **Entering Maintenance Mode** until all running virtual machines are powered off or migrated to different hosts. You cannot power off virtual machines or migrate virtual machines to a host that is entering or in maintenance mode.

To place a host in maintenance mode, all virtual machines that are running on the host must be powered off or migrated to different hosts. If you attempt to place a host that has running virtual machines on it in maintenance mode, DRS must power off or migrate the running virtual machines for the task to complete. If a time out occurs before the virtual machines are powered off or migrated, an error message appears.

When all virtual machines on the host are inactive, the host's icon displays **under maintenance** and the host's Summary panel indicates the new state. While in maintenance mode, the host does not allow you to deploy or power on a virtual machine.

1. Right-click the host and select **Enter maintenance mode**.
The **Confirm maintenance mode change** window appears.
2. Click **Yes**.

The host is in maintenance mode until you select **Exit maintenance mode**.

Manage Permissions in the VMware Host Client

In the VMware Host Client, permissions refer to the access roles assigned to users for various objects, such as virtual machines or ESXi hosts. The roles allow users to perform various tasks on the assigned objects.

For example, to configure memory for the host, a user must be granted a role that includes the **Host > Configuration > Memory Configuration** privilege. By assigning different roles to users for different objects, you can control the tasks that users can perform by using the VMware Host Client.

When connecting directly to a host with the VMware Host Client, the root and vpxuser user accounts have the same access rights as any user assigned the Administrator role on all objects.

All other users initially have no permissions on any object, which means the users cannot view or perform tasks on these objects. A user with Administrator privileges must assign permissions to these users to allow them to perform tasks.

Many tasks require permissions on more than one object. The following rules can help you determine which roles to assign to users to allow particular tasks:

- Any task that consumes hard disk space, such as creating a virtual disk or taking a snapshot, requires the **Datastore > Allocate Space** privilege on the target datastore and the privilege to perform the operation itself.
- Each host and cluster has its own implicit resource pool that contains all the resources of that host or cluster. Deploying a virtual machine directly to a host or cluster requires the **Resource > Assign Virtual Machine to Resource Pool** privilege.

The list of privileges is the same for both ESXi and vCenter Server.

You can create roles and set permissions through a direct connection to the ESXi host.

Permission Validation

vCenter Server and ESXi hosts that use Active Directory regularly validate users and groups against the Windows Active Directory domain. Validation occurs whenever the host system starts and at regular intervals specified in the vCenter Server settings.

For example, if user Smith was assigned permissions and in the domain the user's name was changed to Smith2, the host concludes that Smith no longer exists and removes permissions for that user when the next validation occurs.

Similarly, if user Smith is removed from the domain, all permissions are removed when the next validation occurs. If a new user Smith is added to the domain before the next validation occurs, the new user Smith receives all the permissions the old user Smith was assigned.

Assign Permissions to a User for an ESXi Host in the VMware Host Client

To perform particular activities on an ESXi host, a user must have permissions that are associated with a particular role. In the VMware Host Client, you can assign roles to users and give the users the permissions necessary to perform various tasks on the host.

- Right-click **Host** in the VMware Host Client inventory and click **Permissions**.
The **Manage Permissions** window appears.
- Click **Add user**.
- From the **Select a user** text box and select the user that you want to assign a role to.
- Click the arrow next to the **Select a role** text box and select a role from the list.
- Optional: Select **Propagate to all children** or **Add as group**.

If you set a permission at a vCenter Server level and propagate it to the children objects, the permission applies to data centers, folders, clusters, hosts, virtual machines, and other objects in the vCenter Server instance.

6. Click **Add user** and click **Close**.

Remove Permissions for a User in the VMware Host Client

Removing a permission for a user does not remove the user from the list of users available. It also does not remove the role from the list of available items. It removes the user and role pair from the selected inventory object.

1. Right-click **Host** in the VMware Host Client inventory and click **Permissions**.
The **Manage permissions** window appears.
2. Select a user from the list and click **Remove user**.
3. Click **Close**.

Assign User Permissions for a Virtual Machine in the VMware Host Client

Assign a role to a particular user to give that user permissions to perform specific tasks on a virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine from the list and select **Permissions**.
The **Manage permissions** window appears.
3. Click **Add user**.
4. Click the arrow next to the **Select a user** text box and select the user that you want to assign a role for.
5. Click the arrow next to the **Select a role** text box and select a role from the list.
6. Optional: Select **Propagate to all children**.
If you set a permission at a vCenter Server level and propagate it to the children objects, the permission applies to data centers, folders, clusters, hosts, virtual machines, and similar objects in the vCenter Server instance.
7. Click **Add user** and click **Close**.

Remove Permissions for a Virtual Machine in the VMware Host Client

To make a user unable to perform tasks on a particular virtual machine, remove the permissions of the user for that virtual machine.


Removing a permission for a user does not remove the user from the list of users available. It also does not remove the role from the list of available items. It removes the user and role pair from the selected inventory object.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine from the list and select **Permissions**.
The **Manage permissions** window appears.
3. Select a user from the list and click **Remove user**.
4. Click **Close**.

Generate a Support Bundle in the VMware Host Client

You can generate a support bundle for the ESXi host that you are logged in on. The support bundle contains the log files and system information that you can use to diagnose and resolve problems.

1. Right-click **Host** in the VMware Host Client inventory and select **Generate support bundle** from the drop-down menu.



A dialog box that contains a link to download the bundle pops up when the support bundle is created.

2. Optional: Click **Monitor** in the VMware Host Client inventory, click **Tasks**, and click a log bundle from the list. You can view the link to the log bundle under the table.

Lockdown Mode in the VMware Host Client

To increase the security of your ESXi hosts, you can put them in lockdown mode. In lockdown mode, operations must be performed through vCenter Server by default.

Normal Lockdown Mode and Strict Lockdown Mode

With vSphere 6.0 and later, you can select normal lockdown mode or strict lockdown mode.

Normal Lockdown Mode

In normal lockdown mode, the DCUI service remains active. If the connection to the vCenter Server system is lost, and access through the vSphere Client is unavailable, privileged accounts can log in to the ESXi host's Direct Console Interface and exit lockdown mode. Only the following accounts can access the Direct Console User Interface:

- Accounts in the Exception User list for lockdown mode who have administrative privileges on the host. The Exception Users list is meant for service accounts that perform specific tasks. Adding ESXi administrators to this list defeats the purpose of lockdown mode.

- Users defined in the DCUI.Access advanced option for the host. This option is for emergency access to the Direct Console Interface in case the connection to vCenter Server is lost. These users do not require administrative privileges on the host.

Strict Lockdown Mode

In strict lockdown mode, the DCUI service is stopped. If the connection to vCenter Server is lost and the vSphere Client is no longer available, the ESXi host becomes unavailable, unless the ESXi Shell and SSH services are enabled and Exception Users are defined. If you cannot restore the connection to the vCenter Server system, you must reinstall the host.

Lockdown Mode and the ESXi Shell and SSH Services

Strict lockdown mode stops the DCUI service. However, the ESXi Shell and SSH services are independent of lockdown mode. For lockdown mode to be an effective security measure, ensure that ESXi Shell and SSH services are also deactivated. These services are deactivated by default.

When a host is in lockdown mode, users on the Exception Users list can access the host from the ESXi Shell and through SSH if they have the Administrator role on the host. This access is possible even in strict lockdown mode. Leaving the ESXi Shell service and the SSH service deactivated is the most secure option.

Note: The Exception Users list is meant for service accounts that perform specific tasks such as host backups, and not for administrators. Adding administrator users to the Exception Users list defeats the purpose of lockdown mode.

Put an ESXi Host in Normal Lockdown Mode by Using the VMware Host Client

You can use the VMware Host Client to enter normal lockdown mode.

1. Right-click **Host** in the VMware Host Client inventory, select **Lockdown mode** from the drop-down menu, and select **Enter normal lockdown**.
A warning message appears.
2. Click **Enter normal lockdown**.

Put an ESXi Host in Strict Lockdown Mode by Using the VMware Host Client

You can use the VMware Host Client to enter strict lockdown mode.

1. Right-click **Host** in the VMware Host Client inventory, select **Lockdown mode** from the drop-down menu, and select **Enter strict lockdown**.
The warning message appears.
2. Click **Enter strict lockdown**.

Exit Lockdown Mode by Using the VMware Host Client

If you have entered normal or strict lockdown mode on an ESXi host, you can exit lockdown by using the VMware Host Client.

Right-click **Host** in the VMware Host Client inventory, select **Lockdown mode** from the drop-down menu, and select **Exit lockdown**.

Specify Lockdown Mode Exception Users in the VMware Host Client

With vSphere 6.0 and later, you can add users to the exception users list by using the VMware Host Client. These users do not lose their permissions when the host enters lockdown mode. You can add service accounts, such as a backup agent to the exception users list.

Exception users are host local users or Active Directory users with privileges defined locally for the ESXi host. They are not members of an Active Directory group and are not vCenter Server users. These users are allowed to perform operations on the host based on their privileges. That means, for example, that a read-only user cannot deactivate lockdown mode on a host.

Note: The exception users list is useful for service accounts that perform specific tasks, such as host backups, and not for administrators. Adding administrator users to the exception users list defeats the purpose of lockdown mode.

1. Click **Manage** in the inventory and click **Security & Users**.
2. Click **Lockdown mode**.
3. Click **Add user exception**, enter the name of the user, and click **Add exception**.
4. Optional: Select a name from the exception users list, click **Remove user exception**, and click **Confirm**.

Manage CPU Resources by Using the VMware Host Client

When you connect to an ESXi host with the VMware Host Client, you have access to a limited number of resource management settings.

View Processor Information by Using the VMware Host Client

In the VMware Host Client, you can access information about the current CPU configuration of the ESXi host that you are logged in to.

1. Click **Host** in the VMware Host Client inventory.
2. Expand **Hardware** and expand **CPU**.

You can view the information about the number and type of physical processors, and the number of logical processors.

Assign a Virtual Machine to a Specific Processor in the VMware Host Client

By using CPU affinity, you can assign a virtual machine to a specific processor. This way, you can assign a virtual machine only to a particular available processor in multiprocessor systems.

Power off the virtual machine.

1. Right-click the virtual machine in the VMware Host Client inventory and select **Edit settings**.
2. Under **Virtual Hardware**, expand **CPU**.
3. Under **Scheduling Affinity**, select physical processor affinity for the virtual machine.
Use a hyphen to indicate ranges and a comma to separate values.
For example, 0, 2, 4-7 would indicate processors 0, 2, 4, 5, 6, and 7.

4. Click **Save** to apply your changes.

Monitor an ESXi Host in the VMware Host Client

When you connect to a host using the VMware Host Client, you can monitor the host health status, and view performance charts, events, tasks, system logs, and notifications.

View Charts in the VMware Host Client

When you are logged in to the VMware Host Client, you can view information about resource usage of the ESXi host that you are managing in line chart form.

To reduce memory consumption, the VMware Host Client only contains statistics for the last hour.

1. Click **Monitor** in the VMware Host Client and click **Performance**.
2. Optional: To view the host usage for the last hour, select an option from the drop-down menu.
 - To view the percentage of CPU that the host used during the last hour, select **CPU**.
 - To view the percentage of memory that the host consumed during the last hour, select **Memory**.
 - To view the percentage of network that the host consumed during the last hour, select **Network**.
 - To view the disk usage that the host consumed during the last hour, select **Disk**.

Monitor Hardware Health Status in the VMware Host Client

When you are logged in to the VMware Host Client, you can monitor the health status of the ESXi host hardware.

Note: Hardware health status is only available when the underlying hardware supports it.

1. Click **Monitor** in the VMware Host Client inventory and click **Hardware**.
2. Select the type of information to view.
3. Optional: Use the filter controls above the list to filter the list.
4. Optional: Click a column heading to sort the list.


View Events in the VMware Host Client

Events are records of user actions or system actions that occur on an ESXi host. When you are logged in the VMware Host Client, you can view all events associated with the host that you are managing.

Required privilege: **Read-only**.

Click **Monitor** in the VMware Host Client inventory and click **Events**.

- a) Optional: Select an event to see event details.
- b) Optional: Use the filter controls above the list to filter the list.



- c) Optional: Click a column heading to sort the list.

View Tasks in the VMware Host Client

When you are logged in to the VMware Host Client, you can view tasks that are related to the ESXi host. You can view information about task initiator, task state, task result, task description, and so on.

Click **Monitor** in the VMware Host Client inventory and click **Tasks**.

- a) Optional: Select a task to see task details.
- b) Optional: Use the filter controls above the list to filter the list.
- c) Optional: Click a column heading to sort the list.

View System Logs in the VMware Host Client

When you are logged in to an ESXi host with the VMware Host Client, you can view log entries to get information such as who generated an event, when the event was created, and the type of event.

1. Click **Monitor** in the VMware Host Client inventory and click **Logs**.
The list of logs is displayed.

2. Optional: Click a log to view log details.
3. Optional: Right-click a log and select one of the following options:
 - **Open in new window**
 - **Generate support bundle**

View Notifications in the VMware Host Client

When you are logged in the VMware Host Client, you can view host notifications and recommendations for related tasks that you should perform.

1. Click **Monitor** in the VMware Host Client inventory and click **Notifications**.
2. Select a notification from the list to view the recommended action.

A message with a recommended action and a description is displayed under the notifications list.

Virtual Machine Management with the VMware Host Client

You can configure virtual machines to perform the same tasks as physical computers. Unlike physical computers, virtual machines support special features.

You can use the VMware Host Client to create, register, and manage virtual machines, and to conduct daily administrative and troubleshooting tasks.

Create a Virtual Machine in the VMware Host Client

Virtual machines are the key component of a virtual infrastructure. When you create a virtual machine, you can add it to the host inventory, associate it with a particular datastore, and select an operating system and virtual hardware options.

Verify that you have the **VirtualMachine > Inventory > Create** privileges.

Depending on the properties of the virtual machine you want to create, you might need the following additional privileges:

- **VirtualMachine > Config > AddExistingDisk** if including a virtual disk device that refers to an existing virtual disk file (not RDM).
- **VirtualMachine > Config > AddNewDisk** if including a virtual disk device that creates a new virtual disk file (not RDM).
- **VirtualMachine > Config > RawDevice** if including a raw device mapping (RDM) or SCSI passthrough device.
- **VirtualMachine > Config > HostUSBDevice** if including a VirtualUSB device backed by a host USB device.
- **VirtualMachine > Config > AdvancedConfig** if setting values in `ConfigSpec.extraConfig`.
- **VirtualMachine > Config > SwapPlacement** if setting swapPlacement.
- **Datastore > AllocateSpace** required on all datastores where the virtual machine and its virtual disks are created.
- **Network > Assign** required on the network which is assigned to the new virtual machine that is being created.

After you power on the virtual machine, it consumes resources dynamically as the workload increases, or returns resources dynamically as the workload decreases.

Every virtual machine has virtual devices that provide the same function as physical hardware. A virtual machine gets CPU and memory, access to storage, and network connectivity from the host it runs on.

1. Right-click **Host** in the VMware Host Client inventory and select **Create/Register VM**.
The **New Virtual Machine** wizard opens.

2. Select **Create a new virtual machine** and click **Next**.
3. On the **Select a name and guest OS** page, provide a unique name for the virtual machine and configure the guest operating system.
 - a) In the **Name** text box, enter a name for the virtual machine.
 - b) From the **Compatibility** drop-down menu, select the virtual machine compatibility.
 - c) From the **Guest OS family** drop-down menu, select the guest operating system.
 - d) From the **Guest OS version** drop-down menu, select the guest operating system version.
 - e) To enable VBS on the virtual machine, select the **Enable Windows Virtualization Based Security** check box and click **Next**.

Note: The **Enable Windows Virtualization Based Security** option only appears for the latest Windows OS versions, for example, Windows 10 and Windows Server 2016, and if the virtual machine compatibility is ESXi 6.7 and later.

When you enable this option, hardware virtualization, IOMMU, EFI, and secure boot become available to the guest operating system. You must also enable **Virtualization Based Security** within the guest operating system of this virtual machine.

4. Click **Next**.
5. On the **Select storage** page, select the storage type for the virtual machine and a datastore where to store the virtual machine files.
 - a) To save all the virtual machine disks and configuration files on a standard datastore, click the **Standard** button.
 - b) To save the virtual machine hard disks on the host-local PMem datastore, click the **Persistent Memory** button.
 - c) Select a datastore from the list and click **Next**.

Note: You cannot store the configuration files on a PMem datastore. If you choose to use PMem, you must select a regular datastore for the configuration files of the virtual machine.

6. On the **Customize settings** page, configure the virtual machine hardware and options and click **Next**.
For information about virtual machine options and virtual disk configuration, including instructions for adding different types of devices, see *vSphere Virtual Machine Administration*.
 - a) On the **Customize settings** page, click **Virtual Hardware** and add a new virtual hardware device.
 - Click the **Add hard disk** icon to add a new virtual hard disk.
Note: You can add a standard or a persistent memory hard disk to the virtual machine. The persistent memory hard disk is stored on the host-local PMem datastore.
 - Click the **Add network adapter** icon to add a NIC to the virtual machine.
 - Click the **Add other device** icon to choose other type of device to add to the virtual machine.
Note: If the virtual machine uses PMem storage, the hard disks that are stored on a PMem datastore and the NVDIMM devices that you add to the virtual machine all share the same PMem resources. So, you must adjust

the size of the newly added devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.

- b) Optional: To view and configure device settings, expand any device.

Option	Description
CPU	The CPU or processor is the portion of a computer system that carries out the instructions of a computer program and is the primary element carrying out the computer's functions. CPUs contain cores. The number of virtual CPUs that are available to a virtual machine depends on the number of licensed CPUs on the host, and the number of CPUs supported by the guest operating system. To use the VMware multicore virtual CPUs feature, you must comply with the requirements of the guest operating system EULA.
Memory	You can add, change, or configure virtual machine memory resources or options to enhance virtual machine performance. You can set most of the memory parameters during virtual machine creation or after the guest operating system is installed. The memory resource settings for a virtual machine determine how much of the host's memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine.
Hard disk	You can add large-capacity virtual disks to virtual machines and add more space to existing disks, even when the virtual machine is running. You can set most of the virtual disk parameters during virtual machine creation or after you install the guest operating system.
SCSI controller	Storage controllers appear to a virtual machine as different types of SCSI controllers, including BusLogic Parallel, LSI Logic Parallel, LSI Logic SAS, and VMware Paravirtual SCSI. You can set the type of SCSI bus sharing for a virtual machine and indicate whether the SCSI bus is shared. Depending on the type of sharing, virtual machines can access the same virtual disk simultaneously on the same server or on any server. You can change the SCSI controller configuration for a virtual machine on an ESXi host only.
SATA controller	If a virtual machine has multiple hard disks or CD/DVD-ROM devices, you can add up to three additional SATA controllers to assign the devices to. When you spread the devices among several controllers, you can improve performance and avoid data traffic congestion. You can also add additional controllers if you exceed the thirty-device limit for a single controller. You can boot virtual machines from SATA controllers and use them for large-capacity virtual hard disks.
Network adapter	When you configure a virtual machine, you can add network adapters (NICs) and specify the adapter type. The type of network adapters that are available depend on the following factors: • The virtual machine compatibility, which depends on the host that created or most recently updated it. • Whether the virtual machine compatibility has been updated to the latest version for the current host. • The guest operating system.
CD/DVD drive	You can configure DVD or CD devices to connect to client devices, host devices, or Datastore ISO files.

Option	Description
Video Card	You can select the default settings or specify custom settings. You can specify the number of displays, the total video memory, and enable 3D support for guest operating systems on which VMware supports 3D.
PCI device	You can configure PCI devices on an ESXi host to make them available for passthrough. You can also change the hardware label to restrict the virtual machine placement to specific hardware instances.
Dynamic PCI device	PCI passthrough devices are automatically grouped by their vendor and model name. You can configure the desired devices by the vendor and model name instead of selecting a physical PCI device by the hardware address. You can add all available devices with the same hardware label or with a blank hardware label to a virtual machine. When you power on a virtual machine, specific physical PCI passthrough devices with matching vendor and model names attach to the virtual machine.
Security devices	You can configure Virtual Intel® Software Guard Extensions (vSGX) for virtual machines and provide additional security to your workloads. You can activate or deactivate vSGX when you create a virtual machine or edit an existing virtual machine.

c) Optional: To remove a device, click the delete (✖) icon next to the device.

This option appears only for virtual hardware that you can remove safely.

d) Optional: To customize virtual machine options, click the **VM Options** button.

7. On the **Ready to complete** page, review the details and click **Finish**.

Deploy a Virtual Machine from an OVF or OVA File in the VMware Host Client

When you connect to an ESXi host by using the VMware Host Client, you can deploy virtual machines from OVF, VMDK, and OVA files.

OVF and OVA Limitations for the VMware Host Client

Several limitations are applicable to this deployment method.

OVA Limitations

You can upload OVA files by using either a Web browser or a client. The memory requirements are significant and might cause the Web browser to stop responding or make the system unstable. The size of the OVA file that can be uploaded depends on how much memory is available on your system. VMware tests show that Google Chrome can upload OVA files of about 1 gigabyte. Mozilla Firefox can extract larger OVA files, but might become unresponsive.

To deploy a large OVA file, VMware recommends to first extract the OVA on your system by running the command `tar -xvf <file.ova>`. Then you can provide the deployment wizard with the OVF and VMDKs as separate files.

OVF Limitations


The size of OVF files that a Web browser can upload are also limited. Different Web browsers have different file size limits. Mozilla Firefox has a 4 GB limit. Google Chrome can handle larger files and there is no documented limit.

Deploy a Virtual Machine from an OVF or OVA File in the VMware Host Client

Use the **New Virtual Machine** wizard to deploy virtual machines from OVF and VMDK files, or from OVA files.

OVA deployment is limited to files under 1 GB in size due to Web browser limitations. To deploy an OVA file greater than 1 GB, extract the OVA file using tar and provide the OVF and VMDK files separately.

1. Right-click **Host** in the VMware Host Client inventory and select **Create/Register VM**.



2. On the **Select creation type** page, select **Deploy a virtual machine from an OVF or OVA file** and click **Next**.
3. On the **Select OVF and VMDK files** page, provide a unique name for the virtual machine.

Note: Virtual machine names can contain up to 80 characters.

4. To select an OVF and VMDK, or an OVA file to deploy, click the blue pane.
Your local system storage opens.
5. Select the file that you want to deploy your virtual machine from and click **Open**.
The file you selected appears in the blue pane.
6. Click **Next**.
7. On the **Select storage** page, select the storage type for the virtual machine.
 - a) To save all the virtual machine disks and configuration files on a standard datastore, click the **Standard** button.
 - b) To save the virtual machine hard disks on the host-local PMem datastore, click the **Persistent Memory** button.
 - c) Select a datastore from the list and click **Next**.

Important: The configuration files cannot be stored on a PMem datastore. If you choose to use PMem, you must select a regular datastore for the configuration files of the virtual machine.

8. On the **Deployment options** page, select the network mappings, disk provisioning, and whether you want the virtual machine to power on after deployment.
9. Click **Next**.
10. On the **Ready to complete** page, review the details and click **Finish**.

The virtual machine appears in the VMware Host Client inventory under **Virtual Machines**.

Register an Existing Virtual Machine in the VMware Host Client

If you unregister a virtual machine from a host, but you do not delete the virtual machine from the datastore, you can re-register the virtual machine by using the VMware Host Client. Re-registering a virtual machine makes it appear in the inventory.

Use the datastore browser to select either a datastore, a directory, or a `.vmx` file to add to the list of virtual machines that you register. Selecting a datastore or directory searches for all `.vmx` files in that location. You can browse more than once to append virtual machines to the list.

1. Right-click **Host** in the VMware Host Client inventory and select **Create/Register VM**.
The **New virtual machine** wizard opens.
2. On the **Select creation type** page, select **Register an existing virtual machine** and click **Next**.
3. On the **Select VMs for registration** page, click **Select one or more virtual machines, a datastore, or a directory**, locate the virtual machine that you want to register, and click **Select**.
4. To remove a virtual machine from the list, select the name of the file and click **Remove selected**.
5. To clear your selection and start again, click **Remove all**.
6. Click **Next**.
7. On the **Ready to complete** page, review the details and click **Finish**.

Use Consoles in the VMware Host Client

You can access a virtual machine through a browser console or through VMware Remote Console (VMRC) in the VMware Host Client and perform different tasks on the virtual machine.

Using Browser Console

Note: The browser console is not supported for any version of ESXi prior to 6.0. You must use VMRC in order to access the browser console.

You can use a browser console to gain access to the guest operating system without installing additional software. For additional console functionalities, such as attaching local hardware, install VMware Remote Console.

Note: Currently browser consoles support only US, Japanese and German keyboard layouts. You must select the desired keyboard layout before opening the console.

Using VMware Remote Console


provides access to virtual machines on remote hosts and performs console and device operations, such as configuring operating system settings and monitoring the virtual machine console for *VMware vSphere*. You can perform a variety of tasks on the virtual machine, such as restarting and shutting down the virtual machine guest operating system, resuming and suspending the virtual machine, configuring VMware Tools updates, configuring and managing the virtual machine and different devices, and so on. VMRC can also modify virtual machine settings such as RAM, CPU cores, and disks. VMware Workstation™, VMware Fusion™ or VMware Player™ work as VMRC clients so you do not need to download and install VMRC if you have any of the three installed on your system.

For a full set of console features, download and install VMRC.

Install the VMware Remote Console Application in the VMware Host Client

The (VMRC) is a standalone console application that enables you to connect to client devices and launch virtual machine consoles on remote hosts.

1. Click **Virtual Machines** in the VMware Host Client inventory.
The list of virtual machines available on the host is displayed.
2. Select a virtual machine from the list and from the **Actions** menu, select **Console > Download VMRC**.



The VMware Customer Connect website opens.

3. Download the VMRC installer at [Broadcom Support](#).

Note: You must have a profile at [Broadcom Support](#) to download the VMRC installer.

Launch Remote Console for a Virtual Machine in the VMware Host Client

You can access virtual machines in the VMware Host Client by using . You can launch one or more consoles to access several remote virtual machines at the same time.

Verify that is installed on your local system.

1. Click **Virtual Machines** in the VMware Host Client inventory and select a virtual machine from the list.
2. From the **Actions** menu, select **Console > Launch remote console**.
The opens as a standalone application for the selected virtual machine.

Open a Virtual Machine Console in the VMware Host Client

With the VMware Host Client, you can access the desktop of a virtual machine by launching a console to the virtual machine. From the console, you can perform tasks in the virtual machine, such as configuring operating system settings, running applications, monitoring performance, and so on.

1. Click **Virtual Machines** in the VMware Host Client inventory.
2. Select a powered-on virtual machine from the list.
3. From the **Actions** menu, select **Console** and select whether to open the console in a pop-up window, new window, or a new tab.

Manage a Guest Operating System in the VMware Host Client

With the VMware Host Client, you can manage the guest operating system of a virtual machine. You can install and upgrade VMware Tools, shut down, reboot, and change the configured guest operating system.

Shut Down or Restart a Guest Operating System by Using the VMware Host Client

Install VMware Tools on a virtual machine to be able to shut down and restart the guest operating system on that virtual machine.

Click **Virtual Machines** in the VMware Host Client inventory, select a virtual machine, and select your task.

- To shut down a virtual machine, right-click the virtual machine and select **Guest OS > Shut down**.
- To restart a virtual machine, right-click the virtual machine and select **Guest OS > Restart**.

Change the Guest Operating System in the VMware Host Client

When you change the guest operating system type in the virtual machine settings, you change the setting for the guest operating system in the configuration file of the virtual machine. To change the guest operating system itself, you must install the new operating system in the virtual machine.

Power off the virtual machine.

When you set the guest operating system type for a new virtual machine, vCenter Server applies configuration defaults based on the type of guest operating system. Changing the guest operating system type setting affects the available ranges and recommendations of the virtual machine settings.

1. In the VMware Host Client inventory, right-click the virtual machine and select **Edit Settings**.
2. Click the **VM Options** tab and expand **General Options**.
3. Select a guest operating system type and version.
If you choose a Windows OS version that supports VBS and if the virtual machine's compatibility is ESXi 6.7 and later, the VBS row appears on the **VM Options** tab.
4. Optional: Click the **Enable Virtualization Based Security** to enable VBS.
Important: Enabling VBS requires that you use EFI to boot the virtual machine. Changing the firmware might make the guest OS unbootable.
5. Click **Save** to apply your changes.

The virtual machine configuration parameters for the guest operating system are changed. You can now install the guest operating system.

Introduction to VMware Tools

VMware Tools is a set of services and modules that enable several features in VMware products for better management of guests operating systems and seamless user interactions with them.

VMware Tools has the ability to:

- Pass messages from the host operating system to the guest operating system.
- Customize guest operating systems as a part of the vCenter Server and other VMware products.
- Run scripts that help automate guest operating system operations. The scripts run when the power state of the virtual machine changes.
- Synchronize the time in the guest operating system with the time on the host operating system

VMware Tools Lifecycle Management provides a simplified and scalable approach for installation and upgrade of VMware Tools. It includes a number of feature enhancements, driver-related enhancements, and support for new guest operating systems.

You must run the latest version of VMware Tools or use open-vm-tools distributed with the Linux OS distribution. Although a guest operating system can run without VMware Tools, you must always run the latest version of VMware Tools in your guest operating systems to access the latest features and updates.

You can configure your virtual machine to automatically check and apply VMware Tools upgrades each time you power on your virtual machines.

For information about enabling automatic upgrade of VMware Tools on your virtual machines, see *vSphere Virtual Machine Administration Guide*.

Installing VMware Tools

Although you can use guest operating systems without VMware Tools, many VMware features are not available unless you install VMware Tools. VMware Tools enhances the performance of the guest operating system of your virtual machines.

Installing VMware Tools is part of the process of creating new virtual machines. It is important to upgrade VMware Tools as updates become available. For information about creating virtual machines, see the *VMware Tools User Guide*.

The installers for VMware Tools are ISO image files. An ISO image file looks like a CD-ROM to your guest operating system. Each type of guest operating system, including Windows, Linux, Solaris, FreeBSD, and NetWare, has an ISO image file. When you install or upgrade VMware Tools, the first virtual CD-ROM disk drive of the virtual machine temporarily connects to the VMware Tools ISO file of your guest operating system.

For information about installing or upgrading VMware Tools in Windows virtual machines, Linux virtual machines, Mac OS X virtual machines, Solaris virtual machines, NetWare virtual machines, or FreeBSD virtual machines, see the *VMware Tools User Guide*.

Install VMware Tools from the VMware Host Client

VMware Tools is a suite of utilities that you install in the operating system of a virtual machine. VMware Tools enhances the performance and management of the virtual machine.

You can install VMware Tools in one or more virtual machines by using the VMware Host Client.

1. Click **Virtual Machines** in the inventory.
2. Select a virtual machine from the list.

The virtual machine must be powered on to install VMware Tools.

3. Click **Actions**, select **Guest OS** from the drop-down menu, and select **Install VMware Tools**.

Upgrading VMware Tools

You can upgrade VMware Tools manually or you can configure virtual machines to check for newer versions of VMware Tools and install them.

The guest operating system checks the version of VMware Tools when you power on a virtual machine. The status bar of your virtual machine displays a message when a new version is available.

For vSphere virtual machines, when the installed version of VMware Tools is outdated, the status bar displays the message:

A newer version of Tools is available for this VM

In Windows virtual machines, you can set VMware Tools to notify you when an upgrade is available. If this notification option is enabled, the VMware Tools icon in the Windows taskbar includes a yellow caution icon when a VMware Tools upgrade is available.

To install a VMware Tools upgrade, you can use the same procedure that you used for installing VMware Tools the first time. Upgrading VMware Tools means installing a new version.

For Windows and Linux guest operating systems, you can configure the virtual machine to automatically upgrade VMware Tools. Although the version check is performed when you power on the virtual machine, in case of Windows guest operating systems, the automatic upgrade occurs when you power off or restart the virtual machine. The status bar displays the message `Installing VMware Tools ...` when an upgrade is in progress. The procedure is mentioned below.

Note:

Upgrading VMware Tools on Windows guest operation systems automatically installs the SVGA graphics drivers. The SVGA graphics driver allows the sleep mode available in guest OS power settings to adjust the sleep options. For example, you can use the sleep mode setting **Change when the computer sleeps** to configure your guest OS to automatically go to sleep mode after a certain time or prevent your guest OS from automatically switching to sleep mode after being idle for some time.

Upgrade VMware Tools in the VMware Host Client

You can upgrade VMware Tools on a virtual machine by using the VMware Host Client.

Power on the virtual machine.

1. Click **Virtual Machines** in the VMware Host Client inventory.
2. Select a virtual machine from the list.
3. Click **Actions**, select **Guest OS** from the drop-down menu, and select **Upgrade VMware Tools**.

Configuring a Virtual Machine in the VMware Host Client

You can add or configure most virtual machine properties during the virtual machine creation process or after you create the virtual machine and install the guest operating system.

You can configure three types of virtual machine properties.

Hardware

View existing hardware configuration and add or remove hardware.

Options

View and configure a number of virtual machine properties such as power management interaction between the guest operating system and the virtual machine, and VMware Tools settings.

Resources

Configure CPUs, CPU hyperthreading sources, memory, and disks.

Check the Virtual Machine Compatibility in the VMware Host Client

You can check the virtual machine compatibility by looking at the virtual machine summary page.

1. Click **Virtual Machines** in the VMware Host Client inventory.
2. Click a virtual machine from the list.

The virtual machine compatibility appears under the virtual machine name.

Change the Name of a Virtual Machine in the VMware Host Client

You can change the name of a virtual machine after you finish the creation process. Changing the name does not change the name of any virtual machine files or the name of the directory in which the files reside.

Power off the virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Click **VM Options**.
4. In the **VM Name** text box, enter a new name for the virtual machine.
5. Click **Save**.

View the Location of the Virtual Machine Configuration File in the VMware Host Client

You can view the location of the configuration files and the working files of a virtual machine by using the VMware Host Client.

Power off the virtual machine.

This information is useful when you configure backup systems.

1. Click **Virtual Machines** in the inventory.
2. Right-click the virtual machine and click **Edit Settings**.
3. Click the **VM Options** tab and expand **General Options**.
4. Record the location of the configuration files and the working files.
5. Click **Cancel** to exit the screen.

Configure the Virtual Machine Power States in the VMware Host Client

Changing virtual machine power states is useful when you do maintenance on the host. You can use the system default settings for the virtual machine power controls, or you can configure the controls to interact with the guest operating system.

- Verify that you have privileges to perform the intended power operation on the virtual machine.
- To set optional power functions, install VMware Tools on the virtual machine.
- Power off the virtual machine before editing the VMware Tools options.

For example, you can configure the **Power off** control to power off the virtual machine or shut down the guest operating system.

You can modify many virtual machine configurations while the virtual machine is running, but you might need to change the virtual machine power state for some configurations.

You cannot configure a **Power on** () action. This action powers on a virtual machine that is stopped, or starts a virtual machine and runs a script if the virtual machine is suspended and VMware Tools is installed and available. If VMware Tools is not installed, it starts the suspended virtual machine and does not run a script.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine from the list and select **Edit settings** from the drop-down menu.
3. On the **VM Options** tab, expand **VMware Tools**.
4. Select an option for the virtual machine **Power off** () control from the drop-down menu.

Option	Description
Power Off	Immediately stops the virtual machine. A Power Off action shuts down the guest operating system or powers off the virtual machine. A message indicates that the guest operating system might not shut down properly. Use this power off option only when necessary.
Shut Down Guest	Uses VMware Tools to initiate an orderly system shut down of the virtual machine. Soft power operations are possible only if the tools are installed in the guest operating system.
System Default	Follows system settings. The current value of the system settings appears in parentheses.

5. Select an option for the **Suspend** () control from the drop-down menu.

Option	Description
Suspend	Pauses all virtual machine activity. When VMware Tools is installed and available, a suspend action runs a script and suspends the virtual machine. If VMware Tools is not installed, a Suspend action suspends the virtual machine without running a script.
Put Guest on Standby	Puts the guest operating system on standby. This option stops all processes, but all virtual devices remain connected to the virtual machine.
System Default	Follows system settings. The current value of the system settings appears in parentheses.

6. Select an option for the **Reset** () control from the drop-down menu.

Option	Description
Reset	Shuts down and restarts the guest operating system without powering off the virtual machine. If VMware Tools is not installed, a Reset action resets the virtual machine.
Restart Guest	Uses VMware Tools to initiate an orderly restart. Soft power operations are possible only if the tools are installed in the guest operating system.
Default	Follows system settings. The current value of the system settings appears in parentheses.

- Click **Save**.

Edit the Configuration File Parameters in the VMware Host Client

To fix certain problems with your system, VMware documentation or a VMware Technical Support representative might instruct you to change or add virtual machine configuration parameters.

Power off the virtual machine.

Important: Changing or adding parameters when a system does not have problems might lead to decreased system performance and instability.

The following conditions apply:

- To change a parameter, you must change the existing value for the keyword/value pair. For example, if the existing pair is keyword/value, and you change it to keyword/value2, the new keyword is value2.
- You cannot delete a configuration parameter entry.

CAUTION: You must assign a value to configuration parameter keywords. If you do not assign a value, the keyword might receive a value of 0 or false, which might result in a virtual machine that cannot power on.

- Click **Virtual Machines** in the inventory.
- Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
- On the **VM Options** tab, expand **Advanced**.
- In the Configuration Parameters row, click **Edit Configuration**.
The **Configuration Parameters** dialog box opens.
- Optional: To add a parameter, click **Add Parameter** and enter a name and value for the parameter.
- Optional: To change a parameter, type a new value in the **Value** text box for that parameter.
- Click **OK** to save your changes and exit the **Configuration Parameters** dialog box.
- Click **Save**.

Configure Autostart Options for a Virtual Machine in the VMware Host Client

You can configure auto start options for a virtual machine to set up the virtual machine to start before or after the other virtual machines on the host.

- Click **Virtual Machines** in the inventory.
- Right-click a virtual machine from the list.
- To configure any autostart option, select **Autostart > Configure**.
- To set up the default timing and startup order for virtual machines, select an option from the pop-up menu.

Option	Description
Start delay	Select the time in seconds. After the ESXi host powers on the first virtual machine, the host waits for the specified delay time and powers on the next virtual machine.
Stop delay	Select the time in seconds. Stop delay is the maximum time the ESXi host waits for a shutdown command to complete. The order in which virtual machines are shut down is the reverse of their startup order. After the ESXi host shuts down the first virtual machine within

Option	Description
	the time you specify, the host shuts down the next virtual machine.
Stop action	From the Stop action drop-down menu, select a shut down action that is applicable to the virtual machines on the host when the host shuts down. • System default • Power off • Suspend • Shut down
Wait for heartbeat	• Select Yes to enable the Wait for heartbeat option. You can use this option if the virtual machine has VMware Tools installed. After the ESXi host powers on the first virtual machine, the host immediately powers on the next virtual machine. The startup order in which virtual machines are powered on continues after the virtual machine receives the first heartbeat.

5. Click **Save**.

Upgrade Virtual Machine Compatibility by Using the VMware Host Client

The virtual machine compatibility determines the virtual hardware available to the virtual machine, which corresponds to the physical hardware available on the host machine.

- Create a backup or snapshot of the virtual machines. See [Manage Virtual Machines With Snapshots](#).
- Upgrade VMware Tools. On virtual machines that run Microsoft Windows, if you upgrade the compatibility before you upgrade VMware Tools, the virtual machine might lose its network settings.
- Verify that all `.vmdk` files are available to the ESXi host on a VMFS3, VMFS5, or NFS datastore.
- Verify that the virtual machine is stored on VMFS3, VMFS5, or NFS datastores.
- Verify that the compatibility settings for the virtual machines are not set to the latest supported version.
- Determine the ESXi versions that you want the virtual machines to be compatible with. See *vSphere Virtual Machine Administration*.

You can upgrade the compatibility level to make a virtual machine compatible with the latest version of ESXi running on the host.

For information about virtual machine hardware versions and compatibility, see *vSphere Virtual Machine Administration*.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine from the list and select **Upgrade VM Compatibility** from the pop-up menu.
3. Select the latest supported version and click **Upgrade**.

Virtual CPU Configuration and Limitations

You can set most of the CPU parameters when you create virtual machines or after you install the guest operating system. Some actions require that you power off the virtual machine before you change the settings.

VMware uses the following terminology. Understanding these terms can help you plan your strategy for CPU resource allocation.

CPU	The CPU, or processor, is the component of a computer system that performs the tasks required for computer applications to
------------	--

CPU Socket

run. The CPU is the primary element that performs the computer functions. CPUs contain cores.

Core

A CPU socket is a physical connector on a computer motherboard that connects to a single physical CPU. Some motherboards have multiple sockets and can connect multiple multicore processors (CPUs).

Resource sharing

A core contains a unit containing an L1 cache and functional units needed to run applications. Cores can independently run applications or threads. One or more cores can exist on a single CPU.

Resource allocation

Shares specify the relative priority or importance of a virtual machine or resource pool. If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when the two virtual machines are competing for resources.

You can change CPU resource allocation settings, such as shares, reservation, and limit, when available resource capacity does not meet demands. For example, if at year end, the workload on accounting increases, you can increase the accounting resource pool reserve.

vSphere Virtual Symmetric Multiprocessing (Virtual SMP)

Virtual SMP or vSphere Virtual Symmetric Multiprocessing is a feature that enables a single virtual machine to have multiple processors.

Virtual CPU Limitations

The maximum number of virtual CPUs that you can assign to a virtual machine is 768. The number of virtual CPUs depends on the number of logical CPUs on the host, and the type of guest operating system that is installed on the virtual machine.

Be aware of the following limitations:

- A virtual machine cannot have more virtual CPUs than the number of logical cores on the host. The number of logical cores is equal to the number of physical cores if hyperthreading is deactivated or two times that number if hyperthreading is activated.
- If a running virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a running virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.
- The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.
- Not every guest operating system supports Virtual SMP, and guest operating systems that support this functionality might support fewer processors than are available on the host. For information about Virtual SMP support, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Configuring Multicore Virtual CPUs

VMware multicore virtual CPU support lets you control the number of cores per virtual socket in a virtual machine. This capability lets operating systems with socket restrictions use more of the host CPU cores, which increases overall performance.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

Using multicore virtual CPUs can be useful when you run operating systems or applications that can take advantage of only a limited number of CPU sockets.

You can configure a virtual machine with ESXi 7.0 Update 1 and later compatibility to have up to 768 virtual CPUs. A virtual machine cannot have more virtual CPUs than the actual number of logical CPUs on the host. The number of logical CPUs means the number of physical processor cores or two times that number if hyperthreading is enabled. For example, if a host has 128 logical CPUs, you can configure the virtual machine for 128 virtual CPUs.

You configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single-core CPU, dual-core CPU, tri-core CPU, and so on. Your selection determines the number of sockets that the virtual machine has.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

For more information about multicore CPUs, see the *vSphere Resource Management* documentation.

Virtual CPU Limitations

The maximum number of virtual CPUs that you can assign to a virtual machine is 768. The number of virtual CPUs depends on the number of logical CPUs on the host, and the type of guest operating system that is installed on the virtual machine.

Be aware of the following limitations:

- A virtual machine cannot have more virtual CPUs than the number of logical cores on the host. The number of logical cores is equal to the number of physical cores if hyperthreading is deactivated or two times that number if hyperthreading is activated.
- If a running virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a running virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.
- The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.
- Not every guest operating system supports Virtual SMP, and guest operating systems that support this functionality might support fewer processors than are available on the host. For information about Virtual SMP support, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Configuring Multicore Virtual CPUs

VMware multicore virtual CPU support lets you control the number of cores per virtual socket in a virtual machine. This capability lets operating systems with socket restrictions use more of the host CPU cores, which increases overall performance.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

Using multicore virtual CPUs can be useful when you run operating systems or applications that can take advantage of only a limited number of CPU sockets.

You can configure a virtual machine with ESXi 7.0 Update 1 and later compatibility to have up to 768 virtual CPUs. A virtual machine cannot have more virtual CPUs than the actual number of logical CPUs on the host. The number of logical CPUs means the number of physical processor cores or two times that number if hyperthreading is enabled. For example, if a host has 128 logical CPUs, you can configure the virtual machine for 128 virtual CPUs.

You configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether

you want a single-core CPU, dual-core CPU, tri-core CPU, and so on. Your selection determines the number of sockets that the virtual machine has.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

For more information about multicore CPUs, see the *vSphere Resource Management* documentation.

Change the Number of Virtual CPUs

A virtual machine with ESXi 7.0 Update 1 and later compatibility can have up to 768 virtual CPUs. You can change the number of virtual CPUs while your virtual machine is powered off. If virtual CPU hot add is enabled, you can increase the number of virtual CPUs while the virtual machine is running.

- If CPU hot add is not enabled, power off the virtual machine before adding virtual CPUs.
- To hot add multicore CPUs, verify that the virtual machine is compatible with ESXi 5.0 and later.
- Verify that you have the **Virtual Machine > Configuration > Change CPU Count** privilege.

Virtual CPU hot add is supported for virtual machines with multicore CPU support and ESXi 5.0 and later compatibility. When the virtual machine is powered on and CPU hot add is enabled, you can hot add virtual CPUs to the running virtual machine. You can add only multiples of the number of cores per socket.

If a virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **CPU**.
3. From the **CPU** drop-down menu, select the number of cores.
4. From the **Cores Per Socket** drop-down menu, select the number of cores per socket and click **OK**.

Allocate CPU Resources in the VMware Host Client

To manage workload demands, you can change the amount of CPU resources allocated to a virtual machine by using the shares, reservations, and limits settings.

Power off the virtual machine.

A virtual machine has the following user-defined settings that affect its CPU resource allocation.

Limit	Places a limit on the consumption of CPU time for a virtual machine. This value is expressed in MHz or GHz.
Reservation	Specifies the guaranteed minimum allocation for a virtual machine. The reservation is expressed in MHz or GHz.
Shares	Each virtual machine is granted CPU shares. The more shares a virtual machine has, the more often it receives a time slice of a

CPU when there is no CPU idle time. Shares represent a relative metric for allocating CPU capacity.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **CPU**, and allocate CPU capacity for the virtual machine.

Option	Description
Reservation	Guaranteed CPU allocation for this virtual machine.
Limit	Upper limit for this virtual machine's CPU allocation. Select Unlimited to specify no upper limit.
Shares	CPU shares for this virtual machine in relation to the parent's total. Sibling virtual machines share resources according to their relative share values bounded by the reservation and limit. Select Low , Normal , or High , which specify share values respectively in a 1:2:4 ratio. Select Custom to give each virtual machine a specific number of shares, which express a proportional weight.

4. Click **Save**.

Virtual Memory Configuration

You can add, change, or configure virtual machine memory resources or options to improve virtual machine performance. You can set most of the memory parameters during virtual machine creation or after the guest operating system is installed.

Some actions require that you power off the virtual machine before changing the settings.

The memory resource settings for a virtual machine determine how much of the host's memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size. ESXi hosts limit the memory resource use to the maximum amount useful for the virtual machine, so that you can accept the default of Unlimited memory resources.

Change the Memory Configuration

You can reconfigure the amount of memory allocated to a virtual machine to improve performance.

Verify that you have the **Virtual machine > Change Configuration > Change Memory** privilege on the virtual machine.

Minimum memory size is 4 MB for virtual machines that use BIOS firmware. Virtual machines that use EFI firmware require at least 96 MB of RAM or they cannot power on.

Maximum memory size for virtual machines that use BIOS firmware is 24560 GB. You must use EFI firmware for virtual machines with memory size greater than 6128 GB.

Maximum memory size for a virtual machine depends on the physical memory of the host and the virtual machine compatibility settings.

If the virtual machine memory is greater than the host memory size, swapping occurs, which can have a severe effect on virtual machine performance. The maximum for best performance represents the threshold above which the physical memory of the host is insufficient to run the virtual machine at full speed. This value fluctuates as conditions on the host change, for example, as virtual machines are powered on or off.

The memory size must be a multiple of 4 MB.

Table 560: Maximum Virtual Machine Memory

Introduced in Host Version	Virtual Machine Compatibility	Maximum Memory Size
9.0	9.0 and later	24560 GB
8.0 Update 3	8.0 Update 3 and later	24560 GB
8.0 Update 2	8.0 Update 2 and later	24560 GB
8.0 Update 1	8.0 Update 1 and later	24560 GB
8.0	8.0 and later	24560 GB
7.0 Update 3	7.0 Update 3 and later	24560 GB
7.0 Update 2	7.0 Update 2 and later	24560 GB
7.0 Update 1	7.0 Update 1 and later	24560 GB
7.0	7.0 and later	6128 GB
6.7 Update 2	6.7 Update 2 and later	6128 GB
6.7	6.7 and later	6128 GB
6.5	6.5 and later	6128 GB

The ESXi host version indicates when support began for the increased memory size. For example, the memory size of a virtual machine with 6.0 and later compatibility running on 6.5 is restricted to 4080 GB.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Memory** and change the memory configuration.
 - a) In the **Memory** text box, enter the amount of RAM to assign to the virtual machine.
 - b) Select whether the memory is specified in MB, GB or TB.
3. Click **OK**.

Allocate Memory Resources to a Virtual Machine in the VMware Host Client

You can change the amount of memory resources allocated to a virtual machine by using the shares, reservations, and limits settings. The host determines the appropriate amount of physical RAM to allocate to virtual machines based on these settings. You can assign a high or low shares value to a virtual machine, depending on its load and status.

Power off the virtual machine.

The following user-defined settings affect the memory resource allocation of a virtual machine.

Limit	Places a limit on the consumption of memory for a virtual machine. This value is expressed in megabytes.
Reservation	Specifies the guaranteed minimum allocation for a virtual machine. The reservation is expressed in megabytes. If the reservation cannot be met, the virtual machine will not power on.
Shares	Each virtual machine is granted a number of memory shares. The more shares a virtual machine has, the greater share of host memory it receives. Shares represent a relative metric for allocating memory capacity. For more information about share values, see the <i>vSphere Resource Management</i> documentation.

You cannot assign a reservation to a virtual machine that is larger than the virtual machine configured memory. If you give a virtual machine a large reservation and reduce the virtual machine configured memory size, the reservation is reduced to match the new configured memory size.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **Memory**, and allocate the memory capacity for the virtual machine.

Option	Description
Reservation	Guaranteed memory allocation for this virtual machine.
Limit	Upper limit for this virtual machine's memory allocation. Select Unlimited to specify no upper limit.
Shares	The values Low , Normal , High , and Custom are compared to the sum of all shares of all virtual machines on the server.

4. Click **Save**.

Change Memory Hot Add Settings

Memory hot add lets you add memory resources for a virtual machine while that virtual machine is powered on.

- Power off the virtual machine.
- Verify that the virtual machine has a guest operating system that supports memory hot add functionality.
- Verify that the virtual machine compatibility is ESXi 4.x and later.
- Verify that VMware Tools is installed.

Enabling memory hot add produces some memory overhead on the ESXi host for the virtual machine.

Note: If the ESXi host is of version 7.0 Update 2 and earlier, hot-adding memory to a virtual machine with NVIDIA vGPU requires that the ESXi host has a free vGPU slot. Starting with vSphere 7.0 Update 3, the source host does not require to have a free vGPU slot.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Memory**, and select **Enable** to enable adding memory to the virtual machine while it is powered on.
3. Click **OK**.

You can now add memory to a virtual machine, even if the virtual machine is turned on.

Add an NVDIMM device to a VM in the VMware Host Client

Add a virtual NVDIMM device to a virtual machine to enable it to use non-volatile, or persistent, computer memory. Non-volatile memory (NVM), or persistent memory (PMem), combines the high data transfer rates of the volatile memory with the persistence and resiliency of traditional storage. The virtual NVDIMM device is a virtual NVM device that can retain stored data through reboots or power source failures.

- Verify that the guest OS of the virtual machine supports PMem.
- Verify that the virtual hardware version is 14 or later.
- Verify that you have the **Datastore > Allocate space** privilege.
- Verify that the host or the cluster on which the virtual machine resides has available PMem resources.

Virtual machines consume the PMem resource of the host through a virtual non-volatile dual in-line memory module (NVDIMM) or through a virtual persistent memory disk.

For more information about persistent memory, see [Manage Persistent Memory](#)

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Under the **Virtual Hardware** tab, click **Add other device** and select **NVDIMM** from the drop-down menu.
The NVDIMM device appears in the Virtual Hardware devices list. Each virtual machine can have a maximum of 64 NVDIMM devices.
4. Configure the newly added NVDIMM device.
 - a) In the Virtual Hardware devices list, expand **New NVDIMM**.
 - b) Enter the size of the new NVDIMM device.

Note: You can change the size of the NVDIMM device at a later time. The virtual machine must be powered off.

 - c) Select the NVDIMM controller location or leave the default one.
5. Click **Save** to close the wizard.

Network Virtual Machine Configuration

ESXi networking features enable communication between virtual machines on the same host, between virtual machines on different hosts, and between other virtual and physical machines.

The networking features also allow management of ESXi hosts and enable communication between VMkernel services, such as NFS, iSCSI, or vSphere vMotion, and the physical network. When you configure networking for a virtual machine, you select or change an adapter type, a network connection, and whether to connect the network when the virtual machine powers on.

Network Adapter Basics

When you configure a virtual machine, you can add network adapters (NICs) and specify the adapter type.

Network Adapter Types

The type of network adapters that are available depend on the following factors:

- The virtual machine compatibility, which depends on the host that created or most recently updated it.
- Whether the virtual machine compatibility has been updated to the latest version for the current host.
- The guest operating system.

Supported NICs currently differ between an on-premises environment and VMware Cloud on AWS. The following NIC types are supported in an on-premises deployment:

E1000E	Emulated version of the Intel 82574 Gigabit Ethernet NIC. E1000E is the default adapter for Windows 8 and Windows Server 2012.
E1000	Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in most newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later.
Flexible	Identifies itself as a Vlan adapter when a virtual machine boots, but initializes itself and functions as either a Vlan or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlan adapter to the higher performance VMXNET adapter.
Vlan	Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in 32-bit legacy guest

VMXNET

operating systems. A virtual machine configured with this network adapter can use its network immediately.

Optimized for performance in a virtual machine and has no physical counterpart. Because operating system vendors do not provide built-in drivers for this card, you must install VMware Tools to have a driver for the VMXNET network adapter available.

VMXNET 2 (Enhanced)

Based on the VMXNET adapter but provides high-performance features commonly used on modern networks, such as jumbo frames and hardware offloads. VMXNET 2 (Enhanced) is available only for some guest operating systems on ESX/ ESXi 3.5 and later.

VMXNET 3

A paravirtualized NIC designed for performance. VMXNET 3 offers all the features available in VMXNET 2 and adds several new features, such as multiqueue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery. VMXNET 3 is not related to VMXNET or VMXNET 2.

PVRDMA

A paravirtualized NIC that supports remote direct memory access (RDMA) between virtual machines through the OFED verbs API. All virtual machines must have a PVRDMA device and should be connected to a distributed switch. PVRDMA supports VMware vSphere vMotion and snapshot technology. It is available in virtual machines with hardware version 13 and guest operating system Linux kernel 4.6 and later.

For information about assigning an PVRDMA network adapter to a virtual machine, see the *vSphere Networking* documentation.

SR-IOV passthrough

Representation of a virtual function (VF) on a physical NIC with SR-IOV support. The virtual machine and the physical adapter exchange data without using the VMkernel as an intermediary. This adapter type is suitable for virtual machines where latency might cause failure or that require more CPU resources.

SR-IOV passthrough is available in ESXi 6.0 and later for guest operating systems Red Hat Enterprise Linux 6 and later, and Windows Server 2008 R2 with SP2. An operating system release might contain a default VF driver for certain NICs, while for others you must download and install it from a location provided by the vendor of the NIC or of the host.

For information about assigning an SR-IOV passthrough network adapter to a virtual machine, see the *vSphere Networking* documentation.

For network adapter compatibility considerations, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Legacy Network Adapters and ESXi Virtual Hardware Versions

The default network adapter types for all legacy virtual machines depend on the adapters available and compatible to the guest operating system and the version of virtual hardware on which the virtual machine was created.

If you do not upgrade a virtual machine to use a virtual hardware version, your adapter settings remain unchanged. If you upgrade your virtual machine to take advantage of newer virtual hardware, your default adapter settings will likely change to be compatible with the guest operating system and upgraded host hardware.

To verify the network adapters that are available to your supported guest operating system for a particular version of vSphere ESXi, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Network Adapters and Legacy Virtual Machines

Legacy virtual machines are virtual machines that are supported by the product in use, but are not current for that product.

The default network adapter types for all legacy virtual machines depend on the adapters available and compatible to the guest operating system and the version of virtual hardware on which the virtual machine was created.

If you do not upgrade a virtual machine to correspond with an upgrade to a newer version of an ESXi host, your adapter settings remain unchanged. If you upgrade your virtual machine to take advantage of newer virtual hardware, your default adapter settings will likely change to be compatible with the guest operating system and upgraded host hardware.

To verify the network adapters that are available to your supported guest operating system for a particular version of vSphere ESXi, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Change the Configuration of the Virtual Network Adapter in the VMware Host Client

You can configure the power-on connection setting, the MAC address, and the network connection of the virtual network adapter of a virtual machine.

Required privileges:

- **Virtual Machine > Configuration > Modify device settings** for editing the MAC address and network.
- **Virtual Machine > Interaction > Device connection** for changing **Connect** and **Connect at power on**.
- **Network > Assign network**

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Click the **Virtual Hardware** tab and select the appropriate Network Adapter (NIC) from the hardware list.
4. Optional: To connect the virtual NIC when the virtual machine is powered on, select **Connect at power on**.
5. Optional: Select the adapter type from the **Adapter Type** drop-down menu.
6. Select an option for MAC address configuration.

Option	Description
Automatic	vSphere assigns a MAC address automatically.
Manual	Enter the MAC address to use.

7. Click **Save**.

Add a Network Adapter to a Virtual Machine in the VMware Host Client

When you add a network adapter (NIC) to a virtual machine, you must select the adapter type, the network connection, and whether the device connects when the virtual machine is powered on.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Click the **Virtual Hardware** tab and click **Add network adapter**.
4. In the network connection panel, select either a network with a specific label or a legacy network.
5. Optional: To configure the virtual NIC to connect when the virtual machine is powered on, select **Connect at power on**.
6. Click **Save**.

Virtual Disk Configuration

You can add large-capacity virtual disks to virtual machines and add more space to existing disks, even when the virtual machine is running.

You can set most of the virtual disk parameters during virtual machine creation or after you install the guest operating system.

You can store virtual machine data in a new virtual disk, an existing virtual disk, or a mapped SAN LUN. A virtual disk appears as a single hard disk to the guest operating system. The virtual disk is composed of one or more files on the host file system. You can copy or move virtual disks on the same hosts or between hosts.

For virtual machines running on an ESXi host, you can store virtual machine data directly on a SAN LUN instead of using a virtual disk file. This option is useful if in your virtual machines you run applications that must detect the physical characteristics of the storage device. Mapping a SAN LUN also allows you to use existing SAN commands to manage storage for the disk.

When you map a LUN to a VMFS volume, vCenter Server or the ESXi host creates a raw device mapping (RDM) file that points to the raw LUN. Encapsulating disk information in a file allows vCenter Server or the ESXi host to lock the LUN so that only one virtual machine can write to it. This file has a `.vmdk` extension, but the file contains only disk information that describes the mapping to the LUN on the ESXi system. The actual data is stored on the LUN. You cannot deploy a virtual machine from a template and store its data on a LUN. You can store its data only in a virtual disk file.

The amount of free space in the datastore is always changing. Ensure that you leave sufficient space for virtual machine creation and other virtual machine operations, such as growth of sparse files, snapshots, and so on. To review space utilization for the datastore by file type, see the *vSphere Monitoring and Performance* documentation.

Thin provisioning lets you create sparse files with blocks that are allocated upon first access, which allows the datastore to be over-provisioned. The sparse files can continue growing and fill the datastore. If the datastore runs out of disk space while the virtual machine is running, it can cause the virtual machine to stop functioning.

About Virtual Disk Provisioning Policies

When you perform certain virtual machine management operations, you can specify a provisioning policy for the virtual disk file. The operations include creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine.

NFS datastores with Hardware Acceleration and VMFS datastores support the following disk provisioning policies. On NFS datastores that do not support Hardware Acceleration, only thin format is available.

You can use Storage vMotion or cross-host Storage vMotion to transform virtual disks from one format to another.

Thick Provision Lazy Zeroed

Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand later on first write from the virtual machine. Virtual machines do not read stale data from the physical device.

Thick Provision Eager Zeroed

A type of thick virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take longer to create virtual disks in this format than to create other types of disks. Increasing the size of an Eager Zeroed Thick virtual disk causes a significant stun time for the virtual machine.

Thin Provision

Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the virtual disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations. If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire datastore space provisioned to it. Thin provisioning is the fastest method to create a virtual disk because it creates a disk with just the header information. It does not allocate or zero out storage blocks. Storage blocks are allocated and zeroed out when they are first accessed.

Note: If a virtual disk supports clustering solutions such as Fault Tolerance, do not make the disk thin.

Change the Virtual Disk Configuration in theVMware Host Client

If you run out of disk space, you can increase the size of the disk. You can change the virtual device node and the persistence mode of virtual disk configuration of a virtual machine.

Power off the virtual machine.

Verify that you have the following privileges:

- **Virtual machine > Configuration > Modify device settings** on the virtual machine.

- **Virtual machine > Configuration > Extend virtual disk** on the virtual machine.
- **Datastore > Allocate space** on the datastore.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand the hard disk to view all disk options.
4. Optional: To change the size of the disk, enter a new value in the text box and select the units from the drop-down menu.
5. Optional: To change the way that disks are affected by snapshots, select a disk mode from the **Disk Mode** drop-down menu.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent-Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.
Independent-Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset the virtual machine.

6. Click **Save**.

Add a New Standard Hard Disk to a Virtual Machine in the VMware Host Client

You can add a virtual hard disk to an existing virtual machine, or you can add a hard disk when you customize the virtual machine hardware during the virtual machine creation process. For example, you might need to provide additional disk space for an existing virtual machine with a heavy work load. During virtual machine creation, you might want to add a hard disk that is preconfigured as a boot disk.

- Verify that you are familiar with configuration options and caveats for adding virtual hard disks. See [Virtual Disk Configuration](#).
- Before you add disks larger than 2TB in size to a virtual machine, see *vSphere Virtual Machine Administration*.
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the destination folder or datastore.

Power off the virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Optional: To delete an existing hard disk, move your pointer over the disk and click the **Remove** icon (X).
The disk is removed from the virtual machine. If other virtual machines share the disk, the disk files are not deleted.
4. On the **Virtual Hardware** tab, select **Add hard disk** and select **New standard hard disk** from the drop-down menu.
The hard disk appears in the Virtual Hardware devices list.

5. Expand **New Hard disk**.
6. Optional: Enter a value for the hard disk size and select the units from the drop-down menu.
7. Select the datastore location where you want to store the virtual machine files.
8. Select the format for the virtual machine disk.

Option	Description
Thick Provision Lazy Zeroed	Create a virtual disk in a default thick format. Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine.
Thick Provision Eager Zeroed	Create a thick disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out during creation. It might take much longer to create disks in this format than to create other types of disks.
Thin Provision	Use the thin provisioned format. At first, a thin provisioned disk uses only as much datastore space as the disk initially requires. If the thin disk needs more space later, it can grow to the maximum capacity allocated to it.

9. In the **Shares** drop-down menu, select a value for the shares to allocate to the virtual disk.
Shares is a value that represents the relative metric for controlling disk bandwidth. The values Low, Normal, High, and Custom are compared to the sum of all shares of all virtual machines on the host.
10. If you selected **Custom**, enter a number of shares in the text box.
11. In the **Limit IOPs** box, enter the upper limit of storage resources to allocate to the virtual machine, or select **Unlimited**.
This value is the upper limit of I/O operations per second allocated to the virtual disk.
12. Accept the default or select a different virtual device node.
In most cases, you can accept the default device node. For a hard disk, using a nondefault device node makes controlling the boot order or having different SCSI controller types easier. For example, you might want to boot from an LSI Logic controller and share a data disk with another virtual machine that is using a Buslogic controller with bus sharing turned on.
13. Optional: Select a disk mode.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent-Persistent	Disks in persistent mode behave like conventional physical computer disks. All data written to a disk in persistent mode are written permanently to the disk.
Independent-Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. The virtual disk returns to the same state every time you restart the virtual machine. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset.

14. Click **Save**.

Add an Existing Hard Disk to a Virtual Machine in the VMware Host Client

You can add an existing virtual hard disk to a virtual machine when you customize the virtual machine hardware during the virtual machine creation process or after the virtual machine is created. For example, you might want to add an existing hard disk that is preconfigured as a boot disk.

- Verify that you are familiar with controller and virtual device node behavior for different virtual hard disk configurations.
- Verify that you have the **Virtual machine > Change Configuration > Add existing disk** privilege on the destination folder or datastore.

Power off the virtual machine.

During virtual machine creation, a hard disk and a SCSI or SATA controller are added to the virtual machine by default, based on the guest operating system that you select. If this disk does not meet your needs, you can remove it and add an existing hard disk at the end of the creation process.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add hard disk** and select **Existing hard disk** from the drop-down menu.
4. Optional: To delete an existing hard disk, move your pointer over the disk and click the **Remove** icon (X).
The disk is removed from the virtual machine. If other virtual machines share the disk, the disk files are not deleted.
5. In the Datastore column, expand a datastore, select a virtual machine folder, and select the disk to add.
The disk file appears in the Contents column. The **File type** menu shows the compatibility file types for this disk.
6. Click **Select** and click **Save** to add the existing hard disk.

Add a Persistent Memory Disk in the Host Client

You can add a virtual hard disk to an existing virtual machine, or you can add a hard disk when you customize the virtual machine hardware during the virtual machine creation process. For example, you might need to provide additional disk space for an existing virtual machine with a heavy work load. During virtual machine creation, you might want to add a hard disk that is preconfigured as a boot disk.

- Verify that you are familiar with configuration options and caveats for adding virtual hard disks. See [Virtual Disk Configuration](#).
- Before you add disks larger than 2TB in size to a virtual machine, see *vSphere Virtual Machine Administration*.
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the destination folder or datastore.

Power off the virtual machine.

During virtual machine creation, a hard disk and a SCSI or SATA controller are added to the virtual machine by default, based on the guest operating system that you select. If this disk does not meet your needs, you can remove it and add an existing hard disk at the end of the creation process.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add hard disk** and select **New persistent memory disk** from the drop-down menu.
The hard disk appears in the Virtual Hardware devices list. By default, the disk is stored on the host-local PMem datastore and you cannot change the datastore.

4. Optional: Configure the settings for the new hard disk and click **Save** to close the wizard.

- a) Expand **New Hard disk**.
- b) Enter a value for the hard disk size and select the units from the drop-down menu.

Note: All persistent memory hard disks and NVDIMM modules that you add to the virtual machine share the same PMem resources. So, you must adjust the size of the newly added persistent memory devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.

- c) From the **Shares** drop-down menu, select a value for the shares to allocate to the virtual disk.
Shares is a value that represents the relative metric for controlling disk bandwidth. The values Low, Normal, High, and Custom are compared to the sum of all shares of all virtual machines on the host.
- d) From the **Controller location** drop-down menu, select the location of the controller that the new hard disk uses.
- e) Select a disk mode.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent-Persistent	Disks in persistent mode behave like conventional physical computer disks. All data written to a disk in persistent mode are written permanently to the disk.
Independent-Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. The virtual disk returns to the same state every time you restart the virtual machine. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset.

Use Disk Shares to Prioritize Virtual Machines in the VMware Host Client

You can change the disk resources for a virtual machine. If multiple virtual machines access the same VMFS datastore and the same logical unit number (LUN), use disk shares to prioritize the level of access that virtual machines have to resources. Disk shares distinguish high-priority from low-priority virtual machines.

Power off the virtual machine.

You can allocate the I/O bandwidth of the host to the virtual hard disks of a virtual machine. You cannot pool disk I/O across a cluster.

The shares value represents the relative metric for controlling disk bandwidth to all virtual machines.

Disk shares are relevant only within a given host. The shares assigned to virtual machines on one host have no effect on virtual machines on other hosts.

You can select an IOP limit, which sets an upper limit for storage resources that are allocated to a virtual machine. IOPs are the number of I/O operations per second.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand the hard disk to view the disk options.
4. In the **Shares** drop-down menu, select a value for the shares to allocate to the virtual machine.
5. If you selected **Custom**, enter a number of shares in the text box.
6. In the **Limit - IOPs** text box, enter the upper limit of storage resources to allocate to the virtual machine, or select **Unlimited**.
7. Click **Save**.

Virtual Machine Controller Configuration in the VMware Host Client

In the VMware Host Client, you can add various controllers to virtual machines, such as USB controllers, SCSI controllers, Paravirtual SCSI controllers, and SATA controllers. You can also change the SCSI Bus sharing configuration and the SCSI controller type.

Add a USB Controller to a Virtual Machine

To support USB passthrough from an ESXi host or from a client computer to a virtual machine, you can add a USB controller to the virtual machine.

- Verify that the ESXi host has USB controller hardware and modules that support USB 3.2, 2.0, and 3.1 devices.
- Verify that the client computers have USB controller hardware and modules that support USB 3.2, 2.0, and 3.1 devices present.
- To use the xHCI controller on a Linux guest operating system, verify that the Linux kernel version is 2.6.35 or later.
- Verify that the virtual machine is powered on.
- Required Privilege (ESXi host passthrough): **Virtual Machine > Configuration > Add or Remove Device**.

In the vSphere Client, you can add one xHCI controller and one EHCI+UHCI controller.

- From hardware version 11 to hardware version 16, the supported number of root hub ports per xHCI controller is eight (four local USB 3.1 SuperSpeed ports and four logical USB 2.0 ports).
- From hardware version 17 to hardware version 20, the supported number of root hub ports per xHCI controller is eight (four local USB 3.1 SuperSpeedPlus ports and four logical USB 2.0 ports).
- With hardware version 21, the supported number of root hub ports per xHCI controller is eight (four local USB 3.2 Gen 2x2 ports and four logical USB 2.0 ports).

The conditions for adding a controller vary, depending on the device version, the type of passthrough (host or client computer), and the guest operating system.

Table 561: USB Controller Support

Controller Type	Supported USB Device Version	Supported for Passthrough from ESXi host to a VM	Supported for Passthrough from Client Computer to a VM
EHCI+UHCI	2.0	Yes	Yes
xHCI	3.2, 3.1, 2.0	Yes USB 3.2, 3.1, and 2.0	Yes

Controller Type	Supported USB Device Version	Supported for Passthrough from ESXi host to a VM	Supported for Passthrough from Client Computer to a VM
			Windows 8 or later, Windows Server 2012 and later, or a Linux guest operating system with a 2.6.35 or later kernel.

For Mac OS X systems, the EHCI+UHCI controller is enabled by default and is required for access to a USB mouse and keyboard.

For virtual machines with Windows or Linux guest operating systems, you can add one or two controllers of different types. You cannot add two controllers of the same type.

For USB passthrough from an host to a virtual machine, the USB arbitrator can monitor a maximum of 15 USB controllers. If your system includes more than 15 controllers and you connect USB devices to them, the devices are not available to the virtual machine.

1. Right-click a virtual machine from the vSphere inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click **Add New Device** and from the drop-down menu select **USB Controller**.
The controller appears in the **Virtual Hardware** devices list.
3. To change the USB controller type, expand **New USB Controller**.
If compatibility errors appear, you must fix them before you can add the controller.
4. Click **OK**.

Add one or more USB devices to the virtual machine.

Add SCSI Controllers in the VMware Host Client

You can add SCSI controllers to an existing virtual machine by adding hard disks on unused SCSI Bus numbers.

Power off the virtual machine.

Adding a new hard disk on an unused SCSI Bus number creates a new SCSI controller.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add hard disk** and select **New hard disk** from the drop-down menu.
4. Expand the hard disk to view all options.
5. In the **Controller location** section, select an unused SCSI Bus number from the drop-down menu.
For example, bus and device numbers 0:0 - 0:15 are used by the initial SCSI controller. The second SCSI controller uses bus and device numbers 1:0 - 1:15.
6. Click **Save**.

The new hard disk and new SCSI controller are simultaneously created.

Change the SCSI Bus Sharing Configuration in the VMware Host Client

You can set the type of SCSI bus sharing for a virtual machine and indicate whether to share the SCSI bus. Depending on the type of sharing, virtual machines can access the same virtual disk simultaneously on the same server or on any server.

Power off the virtual machine.

You can change the SCSI controller configuration for a virtual machine only if the virtual machine is on an ESXi host.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand the SCSI Controller that you want to edit.
4. Select the type of sharing in the **SCSI Bus Sharing** list.

Option	Description
None	Virtual disks cannot be shared by other virtual machines.
Virtual	Virtual disks can be shared by virtual machines on the same server.
Physical	Virtual disks can be shared by virtual machines on any server.

5. Click **Save**.

Change the SCSI Controller Type in the VMware Host Client

You can attach virtual disks and RDMs to virtual machines by configuring virtual SCSI controller on the virtual machines.

Power off the virtual machine.

The choice of SCSI controller does not affect whether your virtual disk is an IDE or SCSI disk. The IDE adapter is always ATAPI. The default for your guest operating system is already selected. Older guest operating systems have BusLogic adapter as their default controller.

If you create an LSI Logic virtual machine and add a virtual disk that uses BusLogic adapters, the virtual machine boots from the BusLogic adapters disk. LSI Logic SAS is available only for virtual machines with hardware version 7 or later. Disks with snapshots might not experience performance gains when used on LSI Logic SAS, VMware Paravirtual, and LSI Logic Parallel adapters.

CAUTION: Changing the SCSI controller type might result in a virtual machine boot failure.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Click the **Virtual Hardware** tab and expand a SCSI controller.
4. Select a SCSI controller type from the drop-down menu.
5. Click **Save**.

About VMware Paravirtual SCSI Controllers

VMware Paravirtual SCSI controllers are high performance storage controllers that can result in greater throughput and lower CPU use. These controllers are best suited for high performance storage environments.

VMware Paravirtual SCSI controllers are available for virtual machines with ESXi 4.x and later compatibility. Disks on such controllers might not experience optimal performance gains if they have snapshots or if memory on the ESXi host is over committed. This behavior does not mitigate the overall performance gain of using VMware Paravirtual SCSI controllers as compared to other SCSI controller options.

For platform support for VMware Paravirtual SCSI controllers, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Add a Paravirtual SCSI Controller in the VMware Host Client

You can add a VMware Paravirtual SCSI high performance storage controller to provide greater throughput and lower CPU utilization.

- Verify that the virtual machine has a guest operating system with VMware Tools installed.
- Verify that the virtual machine has hardware version 7 or later.
- Familiarize yourself with VMware Paravirtual SCSI limitations. See *vSphere Virtual Machine Administration*.
- To access boot disk devices attached to a VMware Paravirtual SCSI controller, verify that the virtual machine has a Windows 2003 or Windows 2008 guest operating system.
- In some operating systems, before you change the controller type you must create a virtual machine with an LSI Logic controller and install VMware Tools.

Power off the virtual machine.

VMware Paravirtual SCSI controllers are best suited for environments, especially SAN environments, that run I/O-intensive applications.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, click **Add other device** and select **SCSI Controller** from the drop-down menu. The new SCSI Controllers appear in the Hardware list.
4. Click the **New SCSI Controller** and select **VMware Paravirtual** from the drop-down menu.
5. Click **Save**.

Add a SATA Controller to a Virtual Machine in the VMware Host Client

If a virtual machine has multiple hard disks or CD/DVD-ROM devices, you can add up to three additional SATA controllers to assign the devices to. When you assign the devices to different controllers, you can improve performance and avoid data traffic congestion. You can also add controllers if you need to exceed the thirty-device limit for a single controller.

- Verify that the virtual machine compatibility is ESXi 5.5 and later.
- Verify that you are familiar with storage controller behavior and limitations. See *vSphere Virtual Machine Administration*.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
- Power off the virtual machine.

You can boot virtual machines from SATA controllers and use them for large-capacity virtual hard disks.

Not all guest operating systems support AHCI SATA controllers. Typically, when you create virtual machines with ESXi 5.5 and later compatibility and Mac OS X guest operating systems, a SATA controller is added by default for the virtual hard disk and CD/DVD-ROM devices. Most guest operating systems, including Windows Vista and later, have a default SATA controller for CD/DVD-ROM devices. For verification, see the appropriate *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device**, and select **SATA Controller** from the drop-down menu. The SATA controller appears in the hardware list.

4. Click **Save**.

Add an NVMe Controller in the VMware Host Client

If a virtual machine has multiple hard disks, you can add up to four virtual NVMe controllers to which to assign the disks. Using a NVMe controller significantly reduces the software overhead for processing guest OS I/O, as compared to AHCI SATA or SCSI controllers.

- Verify that the virtual machine has a guest operating system that supports NVMe.
- Verify that the virtual machine compatibility is ESXi 6.5 or later.
- Verify that you are familiar with storage controllers behavior and limitations. For more information, see the *Virtual Machine Administration* guide.
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the virtual machine.

NVMe controllers perform best with virtual disks on an all-flash disk array, local NVMe SSD, and PMem storage.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Under the **Virtual Hardware** tab, click the **Add other device** icon and select **NVMe controller** from the drop-down menu.

A new NVMe controller is added to the virtual machine.

You can add a hard disk to the virtual machine and assign it to the NVMe controller.

Other Virtual Machine Device Configuration in the VMware Host Client

In addition to configuring virtual machine CPU and Memory, and adding hard disks and virtual network adapters, you can also add and configure virtual hardware, such as DVD/CD-ROM drives, floppy drives, and SCSI devices. You can also add a virtual Watchdog Timer (VWDT) device, Precision Clock device, and PCI devices.

Add a CD or DVD Drive to a Virtual Machine in the VMware Host Client

You can use a physical drive on a client or host, or you can use an ISO image to add a CD/DVD drive to a virtual machine.

Power off the virtual machine.

If you want to add a CD/DVD drive that is backed up by USB CD/DVD drive on the host, you must add the drive as a SCSI device. Hot adding or removing SCSI devices from an ESXi host is not supported.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **CD/DVD Drive** from the drop-down menu.
4. Expand **CD/DVD drive** and select an option.

Option	Description
Use physical drive	1. Select Client Device as the location. 2. From the Device Mode drop-down menu, select Emulate CD-ROM or Passthrough CD-ROM.
Use ISO Image	1. Select Datastore ISO File as the location. 2. Enter the path and filename for the image file, or click Browse to navigate to the file.

5. If you do not want the CD-ROM drive to connect when the virtual machine starts, deselect **Connect at power on**.
6. Select the virtual device node that the drive uses in the virtual machine.
7. Click **Save**.

Add a Floppy Drive to a Virtual Machine in the VMware Host Client

Use a physical floppy drive or a floppy image to add a floppy drive to a virtual machine.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.

ESXi does not support floppy drives that are backed up by a physical floppy drive on the host.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **Floppy Drive** from the drop-down menu.
The Floppy drive appears in the hardware list.
4. Expand **Floppy drive** and select the type of device you want to use.

Option	Description
Client Device	Select this option to connect the floppy device to a physical floppy device or a .flp floppy image on the system from which you access the VMware Host Client.
Use existing floppy image	1. Select this option to connect the virtual device to an existing image of a floppy drive on a datastore accessible to the host. 2. Click Browse and select the floppy image.

5. Optional: Select **Connect at power on** to configure the device to connect when the virtual machine powers on.
6. Click **Save**.

Add a USB Device to a Virtual Machine in the VMware Host Client

By using the VMware Host Client, you can add a USB device to a virtual machine.

- Verify that a USB controller is present. See [Add a USB Controller to a Virtual Machine](#).
- Add a physical USB device to the ESXi host where the virtual machine is located by plugging the USB device into the host.

Note: If the ESXi host does not have available USB devices, you cannot add a USB device to the virtual machine.


1. Click **Virtual Machines** in the inventory.
2. Right-click the virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **USB device** from the drop-down menu.
The USB device appears in the hardware list of available hardware devices for the virtual machine.

4. From the **USB device** drop-down, select which USB device to add to the virtual machine.
5. Click **Save**.

Add a Sound Controller to a Virtual Machine in the VMware Host Client

By using the VMware Host Client, you can add a sound controller to a virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **Sound controller** from the drop-down menu.



The sound controller appears in the list of available hardware devices for the virtual machine.

4. From the **Sound card** drop-down menu, select which sound controller to connect to the virtual machine.
5. Click **Save**.

Parallel and Serial Port Configuration in the VMware Host Client

Parallel and serial ports are interfaces for connecting peripherals to the virtual machine. The virtual serial port can connect to a physical serial port or to a file on the host computer.

You can also use the virtual serial port to establish a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. You can add parallel and serial ports and change the serial port configuration.

Add a Serial Port to a Virtual Machine in the VMware Host Client

A virtual machine can use up to four virtual serial ports. You can connect the virtual serial port to a physical serial port or to a file on the host computer. You can also use a host-side-named pipe to set up a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. In addition, you can use a port or a Virtual Serial Port Concentrator (vSPC) URI to connect a serial port over the network.

- Familiarize yourself with the different media types that the port can access, vSPC connections, and any conditions that might apply. See *vSphere Virtual Machine Administration*.
- To connect a serial port over a network, add a Firewall rule set. See *vSphere Virtual Machine Administration*.
- Required privilege: **Virtual Machine > Configuration > Add or Remove Device**

Power off the virtual machine.

- Click **Virtual Machines** in the inventory.
- Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
- On the **Virtual Hardware** tab, select **Add other device** and select **Serial Port**.
The Serial Port appears in the hardware list.
- In the hardware list, expand the serial port and select the type of media port to access.

Option	Description
Use output file	Browse to the location of the file on the host to store the output of the virtual serial port.
Use physical serial port	Select the port from the drop-down menu.
Use named pipe	- Type a name for the pipe in the Pipe name field - Select the Near End and Far End of the pipe from the drop-down menus.
Use network	- From the Direction drop-down menu, select Server or Client. - Type the port URI. The URI is the remote end of the serial port to which the serial port of the virtual machine should connect. - If vSPC is used as an intermediate step to access all virtual machines through a single IP address, select Use Virtual Serial Port Concentrator and enter the vSPC URI location.

- Optional: Deselect **Connect at power on** if you do not want the parallel port device to connect when the virtual machine powers on.
- Click **Save**.

Establishing Serial Port Network Connections to a Client or Server Without Authentication Parameters

If you do not use vSPC and you configure your virtual machine with a serial port connected as a server with a `telnet://:12345` URI, you can connect to your virtual machine's serial port from your Linux or Windows operating system.

```
telnetyourESXiServerIPAddress12345
```

Similarly, if you run the Telnet Server on your Linux system on port 23 (`telnet://yourLinuxBox:23`), you configure the virtual machine as a client URI.

```
telnet://yourLinuxBox:23
```

The virtual machine initiates the connection to your Linux system on port 23.

Add a Parallel Port to a Virtual Machine in the VMware Host Client

To connect peripheral devices to virtual machines, such as printers and scanners, you can use a parallel port. You send the output of such devices to a file on the host computer.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.

Note: To add a parallel port to a virtual machine that runs on an ESXi 4.1 or earlier host, you can also select to send output to a physical parallel port on the host. This option is not available with ESXi 5.0 and later host versions.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **Parallel Port**.
The parallel port appears in the hardware list.
4. Expand the parallel port and in the Connection field, browse to a folder to create the file in.
The file path appears in the **Connection** text box.
5. Optional: Select **Connect at power on** to configure the device to connect when the virtual machine powers on.
6. Click **Save**.

How do I Add a Virtual Watchdog Timer Device to a Virtual Machine

To ensure self-reliance related to the system performance within a virtual machine, you can add a virtual Watchdog Timer (VWDT) device.

If the guest operating system stops responding and cannot recover on its own due to software glitches or errors, the VWDT waits for a predefined period of time and then restarts the system.

You can enable the VWDT to start either by the guest operating system, or by the BIOS or EFI firmware. If you chose the VWDT to start by the BIOS or EFI firmware, it starts before the guest operating system boots.

The VWDT has an important role in guest-based clustering solutions where each virtual machine in the cluster can recover on its own if it fails.

Add a Virtual Watchdog Timer Device to a Virtual Machine in the VMware Host Client

You can add a virtual Watchdog Timer device to a virtual machine to prevent the virtual machine from a guest operating system failure for an extended period of time.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
- Verify that the guest operating system of the virtual machine supports the VWDT device.
- Verify that the virtual hardware version is 17.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and click **Watchdog Timer**.
The Watchdog Timer device appears in the hardware list.
4. Optional: Select **Start with BIOS/EFI boot** to start the watchdog timer by the BIOS or EFI firmware.
When you select this option, the VWDT device starts before the guest operating system. If the Guest OS boot takes too long or does not support the watchdog timer, the device might constantly restart the virtual machine.

5. Click **Save**.

Add a Precision Clock Device to a Virtual Machine in the VMware Host Client

A precision clock is a virtual device that runs on a virtual machine and accesses the system time of a host. By adding a precision clock to a virtual machine, you ensure time synchronization and high precision timestamping.

- Power off the virtual machine.
- Verify that the virtual hardware version is 17.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Modify device settings** privilege on the virtual machine.

1. In the VMware Host Client inventory, click **Virtual Machines**.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, click **Add other device** and select **Precision Clock**.
The precision clock device appears in the hardware list.
4. Optional: Select the time synchronization protocol.
5. Click **Save**.

Add a PCI Device to a Virtual Machine in the VMware Host Client

DirectPath I/O allows the guest operating system of a virtual machine to access the physical PCI and PCIe devices connected to a host directly. By using this technology, you can connect each virtual machine to up to sixteen physical PCI devices.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
- Verify that the PCI devices are connected to the host and marked as available for passthrough.
- If you want to add a dynamic PCI device to a virtual machine, verify that the virtual hardware version is 17.

You can use Dynamic DirectPath I/O to assign multiple PCI passthrough devices to a virtual machine. Starting with vSphere 7.0, you can identify the PCI passthrough devices by their vendor and model name.

Note: Some virtual machine operations become unavailable when you add a PCI or PCIe passthrough device to the virtual machine.

For information about the hardware label configuration, see [Change the Hardware Label in the VMware Host Client](#).

1. In the VMware Host Client inventory, click **Virtual Machines**.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, click **Add other device** and select a device.

Option	Action
PCI device	1. Click PCI device. A new device appears in the hardware list. 2. From the drop-down menu, select a PCI device to connect to the virtual machine.
Dynamic PCI device	1. Click Dynamic PCI device. A new device appears in the hardware list.

Option	Action
	2. Expand New PCI device, and from the drop-down menu, select the PCI passthrough devices to connect to the virtual machine. You can identify PCI passthrough devices by vendor, model name, and hardware label. Hardware labels, if present, are displayed in brackets. Note: When you add a PCI device to a virtual machine, the full memory size of the virtual machine is automatically reserved.

4. Click **Save**.

Securing Virtual Machines with Intel Software Guard Extensions

By using the vSphere Client, you can configure Virtual Intel® Software Guard Extensions (vSGX) for virtual machines and provide additional security to your workloads.

Some modern Intel CPUs implement a security extension called Intel® Software Guard Extensions (Intel SGX). Intel SGX is a processor-specific technology that defines private regions of memory, called enclaves. Intel SGX protects the enclave contents from disclosure or modification in such a way that code running outside the enclave cannot access them.

vSGX enables virtual machines to use Intel SGX technology if available on the hardware. To use vSGX, the ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the BIOS of the ESXi host. You can use the vSphere Client to enable SGX for a virtual machine.

Activate vSGX on a Virtual Machine in the VMware Host Client

To protect the enclave contents from disclosure and modifications, you can activate vSGX on a virtual machine in the VMware Host Client.

- Power off the virtual machine.
- Verify that the virtual machine uses EFI firmware.
- Verify that the ESXi host is version 7.0 or later.
- Verify that the guest operating system in the virtual machine is Linux, Windows 10 (64-bit) or later, or Windows Server 2016 (64-bit) or later.
- Verify that you have the **Virtual machine > Configuration > Modify device settings** privilege on the virtual machine.
- Verify that the ESXi host is installed on an SGX-capable CPU, and SGX is enabled in the BIOS of the ESXi host. For information about the supported CPUs, see <https://kb.vmware.com/s/article/71367>.

Secure Virtual Machines with vSGX

vSphere enables you to configure vSGX for virtual machines.

Some modern Intel CPUs implement a security extension called Intel® Software Guard Extension (Intel® SGX). Intel SGX allows user-level code to define private regions of memory, called enclaves. Intel SGX protects the enclave contents from disclosure or modification in such a way that code running outside the enclave cannot access them.

vSGX enables virtual machines to use Intel SGX technology if available on the hardware. To use vSGX, the ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the BIOS of the ESXi host. You can use the vSphere Client to enable SGX for a virtual machine. For more information, see the *vSphere Security* documentation.

Some operations and features are not compatible with SGX.

- Migration with Storage vMotion
- Suspending or resuming the virtual machine
- Taking a snapshot of the virtual machine
- Fault Tolerance
- Enabling Guest Integrity (GI, platform foundation for VMware AppDefense 1.0)

1. In the VMware Host Client inventory, click **Virtual Machines**.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **Security devices**.
4. Select the **Enable** check box.
5. Under **Enclave page cache size**, enter a new value in the text box and select the size in MB or GB from the drop-down menu.

Note: The enclave page cache size must be a multiple of 2.

6. From the **Launch control configuration** drop-down menu, select the appropriate mode.

Option	Action
Locked	Activates the launch enclave configuration. Under Launch enclave public key hash , enter a valid SHA256 hash. The SHA256 hash key must contain 64 characters.
Unlocked	Activates the launch enclave configuration of the guest operating system.

7. Click **Save**.

Deactivate vSGX on a Virtual Machine in the VMware Host Client

To deactivate vSGX on a virtual machine, you can use the VMware Host Client.

1. In the VMware Host Client inventory, click **Virtual Machines**.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **Security devices**.
4. Deselect the **Enable** check box and click **Save**.

vSGX is deactivated on the virtual machine.

Manage Virtual Machines in the VMware Host Client

After creating a virtual machine in the VMware Host Client, you can perform different management tasks on the virtual machine.

You can delete the virtual machine from the host, remove the virtual machine from a datastore, register it back on a datastore, and so on. You can also return the virtual machine to the host.

Access a Virtual Machine in the VMware Host Client

You can access the virtual machines on the host that you are logged in to, to configure the virtual machine hardware and options, to perform administrative tasks, and to perform basic troubleshooting tasks.

To display a virtual machine in the VMware Host Client inventory, power on the virtual machine.

To access the virtual machines that are available on the host that you are logged in to, click **Virtual Machines** in the VMware Host Client inventory.

The list of available virtual machines is displayed under **Virtual Machines**.

You can now edit the virtual machine settings and perform different administrative and troubleshooting tasks on the virtual machines in the list.

Power States of a Virtual Machine in the VMware Host Client

The basic power operations for a virtual machine include powering on, powering off, suspending, and reset.

- Verify that you have the **VirtualMachine > Interaction > PowerOn** privilege.
- Verify that you have the **VirtualMachine > Interaction > PowerOff** privilege.
- Verify that you have the **VirtualMachine > Interaction > Suspend** privilege.
- Verify that you have the **VirtualMachine > Interaction > Reset** privilege.

For information about how to change the virtual machine power states, see [Configure the Virtual Machine Power States in the VMware Host Client](#).


1. In the VMware Host Client inventory, click **Virtual Machines**.
2. Right-click a virtual machine and select a power operation.

Option	Description
Power On ()	Powers on a virtual machine when the virtual machine is stopped.
Power off ()	Powers off a virtual machine and shuts down the guest operating system. Powering off a virtual machine might cause loss of data.
Suspend ()	Suspends a running virtual machine and leaves it connected to the network. When you resume a suspended virtual machine, the virtual machine continues operating at the same point the virtual machine was at when it was suspended.
Reset ()	Shuts down and restarts the guest operating system without powering off the virtual machine.

Use Virtual Machine Column Configuration in the VMware Host Client

The virtual machines panel in the VMware Host Client allows you to configure the information that you want to display. You can show or hide different columns, such as status, used space, host name, host CPU, and so on.

1. Click **Virtual Machines** in the inventory.
2. In the list of virtual machines, click the down arrow icon next to any column title and select **Select**



columns.

The list with all available columns appears.

3. Select the information that you want to display in the virtual machine panel.

Remove Virtual Machines from a Host in the VMware Host Client

You can unregister a virtual machine if you want to keep it on the datastore, but you no longer want the VMware Host Client inventory to display the virtual machine.

Power off the virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click the virtual machine from the list and select **Unregister**.
3. To confirm that you want to remove the virtual machine from the inventory, click **Yes**.

The host removes the virtual machine from the inventory and no longer tracks its condition.

Remove Virtual Machines from a Datastore in the VMware Host Client

To free space on the datastore, you can remove the virtual machines that you no longer need. Removing a virtual machine from the VMware Host Client inventory deletes all virtual machine files from the datastore, including the configuration file and virtual disk files. You can delete multiple virtual machines

- Power off the virtual machine.
 - Verify that the virtual machine does not share the disk with another virtual machine. If two virtual machines share one disk, the disk files are not deleted.
1. Click **Virtual Machines** in the inventory.
 2. Select one or multiple check boxes next to the virtual machines that you want to remove and select **Actions > Delete**. The **Delete VMs** dialog box opens.

3. Click **Delete**.

Register a Virtual Machine in the VMware Host Client

If you remove a virtual machine or template from a host but do not remove it from the host datastore, you can return it to the host's inventory.

1. Click **Storage** in the VMware Host Client inventory.
2. Right-click a datastore from the list and click **Register a VM**.
3. Select the virtual machine you want to register from the list and click **Register**.

Manage Virtual Machines With Snapshots

Snapshots preserve the state and data of a virtual machine at the time you take the snapshot. When you take a snapshot of a virtual machine, an image of the virtual machine in a given state is copied and stored.

Snapshots are useful when you want to revert repeatedly to a virtual machine state, but you do not want to create multiple virtual machines.

You can take multiple snapshots of a virtual machine to create restoration positions in a linear process. With multiple snapshots, you can save many positions to accommodate many kinds of work processes. Snapshots operate on individual virtual machines. Taking snapshots of multiple virtual machines, for example, taking a snapshot of a VM for each member of a team, requires that you take a separate snapshot of each team member's virtual machine.

Snapshots are useful as a short term solution for testing software with unknown or potentially harmful effects. For example, you can use a snapshot as a restoration point during a linear or iterative process, such as installing update packages, or during a branching process, such as installing different versions of a program. Using snapshots ensures that each installation begins from an identical baseline.

With snapshots, you can preserve a baseline before you change a virtual machine.

Several operations for creating and managing virtual machine snapshots and snapshot trees are available in the vSphere Client. These operations enable you to create snapshots, revert any snapshot in the snapshot hierarchy, delete snapshots, and more. You can create snapshot trees where you save the virtual machine state at any specific time so that you can revert that virtual machine state later. Each branch in a snapshot tree can have up to 32 snapshots.

A snapshot preserves the following information:

- Virtual machine settings. The virtual machine directory, which includes the disks added or changed after you take the snapshot.
- Power state. The virtual machine can be powered on, powered off, or suspended.
- Disk state. State of all the virtual machine's virtual disks.
- (Optional) Memory state. The contents of the virtual machine's memory.

The Snapshot Hierarchy

The vSphere Client presents the snapshot hierarchy as a tree with one or more branches. Snapshots in the hierarchy have parent to child relationships. In linear processes, each snapshot has one parent snapshot and one child snapshot, except for the last snapshot, which has no child snapshot. Each parent snapshot can have more than one child. You can revert to the current parent snapshot or to any parent or child snapshot in the snapshot tree and create more snapshots from that snapshot. Each time you revert a snapshot and take another snapshot, a branch (child snapshot) is created.

Parent Snapshots

The first virtual machine snapshot that you create is the base parent snapshot. The parent snapshot is the most recently saved version of the current state of the virtual machine. Taking a snapshot creates a delta disk file for each disk attached to the

virtual machine and optionally, a memory file. The delta disk files and memory file are stored with the base `.vmdk` file. The parent snapshot is always the snapshot that appears immediately above the **You are here** icon in the Snapshot Manager. If you revert a snapshot, that snapshot becomes the parent of the **You are here** current state.

Note: The parent snapshot is not always the snapshot that you took most recently.

Child Snapshots

A snapshot of a virtual machine taken after the parent snapshot. Each child snapshot contains delta files for each attached virtual disk, and optionally a memory file that points from the present state of the virtual disk (You are here). Each child snapshot's delta files merge with each previous child snapshot until reaching the parent disks. A child disk can later be a parent disk for future child disks.

The relationship of parent and child snapshots can change if you have multiple branches in the snapshot tree. A parent snapshot can have more than one child. Many snapshots have no children.

CAUTION: Do not manually manipulate individual child disks or any of the snapshot configuration files because doing so can compromise the snapshot tree and result in data loss. This restriction includes disk resizing and making modifications to the base parent disk by using the `vmkfstools` command.

Snapshot Behavior

Taking a snapshot preserves the disk state at a specific time by creating a series of delta disks for each attached virtual disk or virtual RDM and optionally preserves the memory and power state by creating a memory file. Taking a snapshot creates a snapshot object in the Snapshot Manager that represents the virtual machine state and settings.

Each snapshot creates an additional delta `.vmdk` disk file. When you take a snapshot, the snapshot mechanism prevents the guest operating system from writing to the base `.vmdk` file and instead directs all writes to the delta disk file. The delta disk represents the difference between the current state of the virtual disk and the state that existed at the time that you took the previous snapshot. If more than one snapshot exists, delta disks can represent the difference between each snapshot. Delta disk files can expand quickly and become as large as the entire virtual disk if the guest operating system writes to every block of the virtual disk.

Snapshot Files and Limitations

When you take a snapshot, you capture the state of the virtual machine settings and the virtual disk. If you are taking a memory snapshot, you also capture the memory state of the virtual machine. These states are saved to files that reside with the base files of the virtual machines.

Snapshot Files

A snapshot consists of files that are stored on a supported storage device. A Take Snapshot operation creates `.vmdk`, `-delta.vmdk` or `-sesparse.vmdk`, `.vmsd`, and `.vmsn` files. By default, the first and all delta disks are stored with the base `.vmdk` file. The `.vmsd` and `.vmsn` files are stored in the virtual machine directory.

SEsparse is a default format for all delta disks on the VMFS6 datastores.

Delta disk files

A `.vmdk` file to which the guest operating system can write. The delta disk represents the difference between the current state of the virtual disk and the state that existed at the time that the previous snapshot was taken. When you take a snapshot, the

state of the virtual disk is preserved, the guest operating system stops writing to it, and a delta or child disk is created.

A delta disk has two files. One is a small descriptor file that contains information about the virtual disk, such as geometry and child-parent relationship information. The other one is a corresponding file that contains the raw data.

The files that make up the delta disk are called child disks or redo logs.

A `-flat.vmdk` file that is one of two files that comprises the base disk. The flat disk contains the raw data for the base disk. This file does not appear as a separate file in the Datastore Browser.

A `.vmsd` file that contains the virtual machine's snapshot information and is the primary source of information for the Snapshot Manager. This file contains line entries, which define the relationships between snapshots and between child disks for each snapshot.

A `.vmsn` file that includes the active state of the virtual machine. Capturing the memory state of the virtual machine lets you revert to a powered on virtual machine state. With nonmemory snapshots, you can only revert to a powered off virtual machine state. Memory snapshots take longer to create than nonmemory snapshots. The time the ESXi host takes to write the memory onto the disk depends on the amount of memory the virtual machine is configured to use.

Flat file

Database file

Memory file

A **Take Snapshot** operation creates `.vmdk`, `-delta.vmdk` or `-sesparse.vmdk`, and `.vmsn` files and updates `.vmsd` file.

File	Description
The base disk names in a VM are: <code>vmname_number.vmdk</code> and <code>vmname_number-flat.vmdk</code> The newly created disk names in a VM are: <code>vmname_number-number.vmdk</code> , <code>vmname_number-number-sesparse.vmdk</code> or <code>vmname_number-number-delta.vmdk</code>	Snapshot file that represents the difference between the current state of the virtual disk and the state that existed at the time the previous snapshot was taken. The base disk filename uses the following syntax, <code>vm_000001.vmdk</code> where <code>vm</code> is the name of the virtual machine and the six-digit number, <code>000001</code>, is based on the files that already exist in the directory. The number does not consider the number of disks that are attached to the virtual machine. The newly created disk filenames are, for example: <code>test-512-000001-sesparse.vmdk</code> <code>test-512-000001.vmdk</code> <code>test-512_1-000001-sesparse.vmdk</code> <code>test-512_1-000001.vmdk</code>
<code>vmname.vmsd</code>	Database of the virtual machine's snapshot information and the primary source of information for the Snapshot Manager.
<code>vmname-Snapshotnumber.vmsn</code>	Memory state of the virtual machine at the time you take the snapshot. The filename uses the following syntax, <code>vm1-Snapshot1.vmsn</code> , where <code>vm1</code> is the virtual machine name, and <code>Snapshot1</code> is the first snapshot.

File	Description
	Note: A <code>.vmsn</code> file is created each time you take a snapshot, regardless of the memory selection. A <code>.vmsn</code> file without memory is much smaller than one with memory.

Snapshot Limitations

Snapshots can affect the virtual machine performance and do not support some disk types or virtual machines configured with bus sharing. Snapshots are useful as short-term solutions for capturing point-in-time virtual machine states and are not appropriate for long-term virtual machine backups.

- VMware does not support snapshots of raw disks, RDM physical mode disks, or guest operating systems that use an iSCSI initiator in the guest.
- Virtual machines with independent disks must be powered off before you take a snapshot.
- You can take a memory snapshot of a virtual machine with an independent disk only to analyze the guest operating system behavior of a virtual machine. You cannot use such snapshots for virtual machine backups because restoring this type of snapshots is unsupported.
- Quiesced snapshots require VMware Tools installation and guest operating system support.
- Snapshots are not supported with PCI vSphere DirectPath I/O devices.
- VMware does not support snapshots of virtual machines configured for bus sharing. If you require bus sharing, consider running backup software in your guest operating system as an alternative solution. If your virtual machine currently has snapshots that prevent you from configuring bus sharing, delete (consolidate) the snapshots.
- Snapshots provide a point-in-time image of the disk that backup solutions can use, but Snapshots are not intended for a robust method of backup and recovery. If the files containing a virtual machine are lost, its snapshot files are also lost. Also, large numbers of snapshots are difficult to manage, consume large amounts of disk space, and are not protected if there is hardware failure.
- Snapshots can negatively affect the performance of a virtual machine. Performance degradation is based on how long the snapshot or snapshot tree is in place, the depth of the tree, and how much the virtual machine and its guest operating system have changed from the time you took the snapshot. Also, you might see a delay in the amount of time it takes the virtual machine to power on. Do not run production virtual machines from snapshots on a permanent basis.

Snapshot Limitations

Snapshots can affect the virtual machine performance and do not support some disk types or virtual machines configured with bus sharing. Snapshots are useful as short-term solutions for capturing point-in-time virtual machine states and are not appropriate for long-term virtual machine backups.

- VMware does not support snapshots of raw disks, RDM physical mode disks, or guest operating systems that use an iSCSI initiator in the guest.
- Virtual machines with independent disks must be powered off before you take a snapshot.
- Starting with ESXi 7.0 Update 3l, you can take a memory snapshot of a virtual machine with an independent disk only to analyze the guest operating system behavior of a virtual machine. You cannot use such snapshots for virtual machine backups because restoring this type of snapshots is unsupported.
- Quiesced snapshots require VMware Tools installation and guest operating system support.
- Snapshots are not supported with PCI vSphere DirectPath I/O devices.
- VMware does not support snapshots of virtual machines configured for bus sharing. If you require bus sharing, consider running backup software in your guest operating system as an alternative solution. If your virtual machine currently has snapshots that prevent you from configuring bus sharing, delete (consolidate) the snapshots.
- Snapshots provide a point-in-time image of the disk that backup solutions can use, but Snapshots are not meant to be a robust method of backup and recovery. If the files containing a virtual machine are lost, its snapshot files are also

lost. Also, large numbers of snapshots are difficult to manage, consume large amounts of disk space, and are not protected if there is hardware failure.

- Snapshots can negatively affect the performance of a virtual machine. Performance degradation is based on how long the snapshot or snapshot tree is in place, the depth of the tree, and how much the virtual machine and its guest operating system have changed from the time you took the snapshot. Also, you might see a delay in the amount of time it takes the virtual machine to power on. Do not run production virtual machines from snapshots on a permanent basis.
- If a virtual machine has virtual hard disks larger than 2 TB, snapshot operations can take much longer to finish.

Take Snapshots of a Virtual Machine

You can take one or more snapshots of a virtual machine to capture the settings state, disk state, and memory state at specific times. When you take a snapshot, you can also quiesce the virtual machine files and exclude the virtual machine disks from snapshots.

When you take a snapshot, other activity that is occurring in the virtual machine might affect the snapshot process when you revert to that snapshot. The best time to take a snapshot from a storage perspective, is when you are not incurring a large I/O load. The best time to take a snapshot from a service perspective is when no applications in the virtual machine are communicating with other computers. The potential for problems is greatest if the virtual machine is communicating with another computer, especially in a production environment. For example, if you take a snapshot while the virtual machine is downloading a file from a server on the network, the virtual machine continues downloading the file and communicating its progress to the server. If you revert to the snapshot, communications between the virtual machine and the server are confused and the file transfer fails. Depending on the task that you are performing, you can create a memory snapshot or you can quiesce the file system in the virtual machine.

Memory Snapshots

The default selection for taking snapshots. When you capture the virtual machine's memory state, the snapshot retains the live state of the virtual machine. Memory snapshots create a snapshot at a precise time, for example, to upgrade software that is still working. If you take a memory snapshot and the upgrade does not complete as expected, or the software does not meet your expectations, you can revert the virtual machine to its previous state.

When you capture the memory state, the virtual machine's files do not require quiescing. If you do not capture the memory state, the snapshot does not save the live state of the virtual machine and the disks are crash consistent unless you quiesce them.

Quiesced Snapshots

When you quiesce a virtual machine, VMware Tools quiesces the file system of the virtual machine. A quiesce operation ensures that a snapshot disk represents a consistent state of the guest file systems. Quiesced snapshots are appropriate for automated or periodic backups. For example, if you are unaware of the virtual machine's activity, but want several recent backups to revert to, you can quiesce the files.

If the virtual machine is powered off or VMware Tools is not available, the `Quiesce` parameter is not available. You cannot quiesce virtual machines that have large capacity disks.

Important: Do not use snapshots as your only backup solution or as a long-term backup solution.

Change Disk Mode to Exclude Virtual Disks from Snapshots

You can set a virtual disk to independent mode to exclude the disk from any snapshots taken of its virtual machine.

Prerequisites

Power off the virtual machine and delete any existing snapshots before you change the disk mode. Deleting a snapshot involves committing the existing data on the snapshot disk to the parent disk.

Required privileges:

- **Virtual machine > Snapshot management > Remove Snapshot**
- **Virtual machine > Configuration > Modify device settings**

Note: An independent disk does not participate in virtual machine snapshots. That is, the disk state is independent of the snapshot state and creating, consolidating, or reverting to snapshots does not have effect on the disk.

You can take a memory snapshot of a virtual machine with an independent disk, but the snapshot cannot be restored. The main purpose of these snapshots is to capture the state of a virtual machine for further analysis through the Vmss2core or similar tools.

Procedure

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Hard disk**, and select an independent disk mode option.

Option	Description
Dependent	Dependent mode is the default disk mode. When you take a snapshot of a virtual machine, dependent disks are included in the snapshot. When you revert to the previous snapshot, all data are reverted to the point of taking a snapshot.
Independent - Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk even if you revert a snapshot. When you power off or reset a virtual machine, the disk and all its snapshots are preserved.
Independent - Nonpersistent	Disks in nonpersistent mode behave like read-only disks. Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset the virtual machine, or when you delete a snapshot.

3. Click **OK**.

Take a Snapshot of a Virtual Machine

You can take a snapshot when a virtual machine is powered on, powered off, or suspended. If you are suspending a virtual machine, wait until the suspend operation finishes before you take a snapshot.

When you create a memory snapshot, the snapshot captures the state of the virtual machine memory and the virtual machine power settings. When you capture the virtual machine memory state, the snapshot operation takes longer to complete. You might also see a momentary lapse in response over the network.

When you quiesce a virtual machine, VMware Tools quiesces the file system in the virtual machine. The quiesce operation pauses or alters the state of running processes on the virtual machine, especially processes that might modify information stored on the disk during a revert operation.

Application-consistent quiescing is not supported for virtual machines with IDE or SATA disks.

Note: If you take a snapshot of a Dynamic Disk (a Microsoft-specific disk type), the snapshot technology preserves the quiesce state of the file system, but does not preserve the quiesce state of the application.

Prerequisites

- If you are taking a memory snapshot of a virtual machine that has multiple disks in different disk modes, verify that the virtual machine is powered off. For example, if you have a special purpose configuration that requires you to use an independent disk, you must power off the virtual machine before taking a snapshot.
- To capture the memory state of the virtual machine, verify that the virtual machine is powered on.
- To quiesce the virtual machine files, verify that the virtual machine is powered on and that VMware Tools is installed.
- Verify that you have the **Virtual machine > Snapshot management > Create snapshot** privilege on the virtual machine.

Procedure

1. In the vSphere Client, navigate to a virtual machine and click the **Snapshots** tab.
2. Click **Take Snapshot**.
Result: The **Take snapshot** dialog box opens.
3. Enter a name for the snapshot.
4. (Optional) Enter a description for the snapshot.
5. (Optional) To capture the memory of the virtual machine, select the **Snapshot the virtual machine's memory** check box.
6. (Optional) To pause running processes on the guest operating system so that file system contents are in a known consistent state when you take a snapshot, select the **Quiesce guest file system (requires VMware Tools)** check box.

Note: You can quiesce the virtual machine files only when the virtual machine is powered on and the **Snapshot the virtual machine's memory** check box is deselected.

For more information about the virtual machine behavior when you quiesce a virtual machine, see the KB article <https://kb.vmware.com/s/article/5962168>.

7. Click **Create**.

Take a Snapshot of a Virtual Machine in the VMware Host Client

You can take one or more snapshots of a virtual machine to capture the settings state, disk state, and memory state at the time you take the snapshot. When you take a snapshot, you can also quiesce the virtual machine files and exclude the virtual machine disks from snapshots. You can take a snapshot when a virtual machine is powered on, powered off, or suspended. To take a snapshot of a suspended virtual machine, wait until the suspend operation finishes before you take a snapshot.

- If you are taking a memory snapshot of a virtual machine that has multiple disks in different disk modes, verify that the virtual machine is powered off. For example, if you have a special purpose configuration that requires you to use an independent disk, you must power off the virtual machine before taking a snapshot.
- To capture the memory state of the virtual machine, verify that the virtual machine is powered on.
- To quiesce the virtual machine files, verify that the virtual machine is powered on and that VMware Tools is installed.
- Verify that you have the **Virtual machine > Snapshot management > Create snapshot** privilege on the virtual machine.

When you take a snapshot, other activity that is occurring in the virtual machine might affect the snapshot process when you revert to that snapshot. The best time to take a snapshot from a storage perspective, is when you are not incurring a large I/O load. The best time to take a snapshot from a service perspective is when no applications in the virtual machine are communicating with other computers. The potential for problems is greatest if the virtual machine is communicating with another computer, especially in a production environment. For example, if you take a snapshot while the virtual machine is downloading a file from a server on the network, the virtual machine continues downloading the file and communicating its progress to the server. If you revert to the snapshot, communications between the virtual machine and the server are confused and the file transfer fails. Depending on the task that you are performing, you can create a memory snapshot or you can quiesce the file system in the virtual machine.

Memory Snapshots

The default selection for taking snapshots. When you capture the virtual machine memory state, the snapshot retains the live state of the virtual machine. Memory snapshots create a snapshot at a precise time, for example, to upgrade software that is still working. If you take a memory snapshot and the upgrade does not complete as expected, or the software does not meet your expectations, you can revert the virtual machine to its previous state.

When you capture the memory state, the virtual machine files do not require quiescing. If you do not capture the memory state, the snapshot does not save the live state of the virtual machine and the disks are crash consistent unless you quiesce them.

Snapshots that capture the memory state of a virtual machine take longer to complete. You might also see a momentary lapse in response over the network.

Quiesced Snapshots

When you quiesce a virtual machine, VMware Tools quiesces the file system of the virtual machine. A quiesce operation ensures that a snapshot disk represents a consistent state of the guest file systems. The quiesce operation pauses or alters the state of running processes on the virtual machine, especially processes that might modify information stores on the disk during a restore operation. Quiesced snapshots are appropriate for automated or periodic backups. For example, if you are unaware of the virtual machine activity, but want several recent backups to revert to, you can quiesce the files.

If the virtual machine is powered off or VMware Tools is not available, the `Quiesce` parameter is not available. You cannot quiesce virtual machines that have large capacity disks.

Application-consistent quiescing is not supported for virtual machines with IDE or SATA disks.

Important: Do not use snapshots as your only backup solution or as a long-term backup solution.

Note: If you take a snapshot of a Dynamic Disk (Microsoft-specific disk type), the snapshot technology preserves the quiesced state of the file system, but does not preserve the quiesced state of the application.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine from the list and select **Snapshots > Take snapshot**.
3. Enter a name for the snapshot.
4. Optional: Type a description for the snapshot.
5. Optional: Select the **Snapshot the virtual machine's memory** check box to capture the memory of the virtual machine.
6. Optional: Deselect **Snapshot the virtual machine's memory** and select **Quiesce guest file system (needs VMware Tools installed)** check box to pause running processes on the guest operating system so that file system contents are in a known consistent state when you take the snapshot.

Quiesce the virtual machine files only when the virtual machine is powered on and you do not want to capture the virtual machine's memory.

7. Click **Take snapshot**.

Restoring Snapshots

To return a virtual machine to its original state, or to return to another snapshot in the snapshot hierarchy, you can restore a snapshot.

When you restore a snapshot, you return the virtual machine's memory, settings, and the state of the virtual machine disks to the state they were in at the time you took the snapshot. If you want the virtual machine to be suspended, powered on, or powered off when you start it, make sure that it is in the correct state when you take the snapshot.

You can restore snapshots in the following ways:

Revert to Latest Snapshot

Restores the parent snapshot, one level up in the hierarchy from the **You are Here** position. **Revert to Latest Snapshot** activates the parent snapshot of the current state of the virtual machine.

Revert To

Lets you restore any snapshot in the snapshot tree and makes that snapshot the parent snapshot of the current state of the virtual machine. Subsequent snapshots from this point create a new branch of the snapshot tree.

Restoring snapshots has the following effects:

- The current disk and memory states are discarded, and the virtual machine reverts to the disk and memory states of the parent snapshot.
- Existing snapshots are not removed. You can restore those snapshots at any time.
- If the snapshot includes the memory state, the virtual machine will be in the same power state as when you created the snapshot.

Table 562: Virtual Machine Power State After Restoring a Snapshot

Virtual Machine State When Parent Snapshot Is Taken	Virtual Machine State After Restoration
Powered on (includes memory)	Reverts to the parent snapshot, and the virtual machine is powered on and running.
Powered on (does not include memory)	Reverts to the parent snapshot and the virtual machine is powered off.
Powered off (does not include memory)	Reverts to the parent snapshot and the virtual machine is powered off.

Virtual machines running certain kinds of workloads can take several minutes to resume responsiveness after reverting from a snapshot.

Note: vApp metadata for virtual machines in vApps does not follow the snapshot semantics for virtual machine configuration. vApp properties that are deleted, modified, or defined after a snapshot is taken remain intact (deleted, modified, or defined) after the virtual machine reverts to that snapshot or any previous snapshots.

Revert to the Latest Snapshot in the VMware Host Client

To return a virtual machine to its original state, or to return to another snapshot in the snapshot hierarchy, you can restore a snapshot.

Verify that you have the **Virtual machine > Snapshot management > Revert to snapshot** privilege on the virtual machine.

When you restore a snapshot, you return the virtual machine memory, settings, and the state of the virtual machine disks to the state they were in at the time you took the snapshot. If you want the virtual machine to be suspended, powered on, or powered off when you start it, make sure that it is in the correct state when you take the snapshot.

You can restore snapshots in the following ways:

Revert to Latest Snapshot

Restores the parent snapshot, one level up in the hierarchy from the **You are Here** position. **Revert to Latest Snapshot** activates the parent snapshot of the current state of the virtual machine.

Revert To

Lets you restore any snapshot in the snapshot tree and makes that snapshot the parent snapshot of the current state of the virtual machine. Subsequent snapshots from this point create a new branch of the snapshot tree.

Restoring snapshots has the following effects:

- The current disk and memory states are discarded, and the virtual machine reverts to the disk and memory states of the parent snapshot.
- Existing snapshots are not removed. You can restore those snapshots at any time.
- If the snapshot includes the memory state, the virtual machine will be in the same power state as when you created the snapshot.

Table 563: Virtual Machine Power State After Restoring a Snapshot

Virtual Machine State When Parent Snapshot Is Taken	Virtual Machine State After Restoration
Powered on (includes memory)	Reverts to the parent snapshot, and the virtual machine is powered on and running.
Powered on (does not include memory)	Reverts to the parent snapshot and the virtual machine is powered off.
Powered off (does not include memory)	Reverts to the parent snapshot and the virtual machine is powered off.

Virtual machines running certain kinds of workloads can take several minutes to resume responsiveness after reverting from a snapshot.

Note: vApp metadata for virtual machines in vApps does not follow the snapshot semantics for virtual machine configuration. vApp properties that are deleted, modified, or defined after a snapshot is taken remain intact (deleted, modified, or defined) after the virtual machine reverts to that snapshot or any previous snapshots.

1. Click **Virtual Machines** in the inventory.
2. Right-click the virtual machine in the list and select **Snapshots > Restore snapshot**.

Note: The current state of the virtual machine will be lost unless you save it in a snapshot.

3. Click **Restore** to revert the virtual machine to the most recent snapshot.

Deleting Snapshots

Deleting a snapshot removes the snapshot from the Snapshot Manager. The snapshot files are consolidated and written to the parent snapshot disk and merge with the virtual machine base disk.

Deleting a snapshot does not change the virtual machine or other snapshots. Deleting a snapshot consolidates the changes between snapshots and previous disk states and writes all the data from the delta disk that contains the

information about the deleted snapshot to the parent disk. When you delete the base parent snapshot, all changes merge with the base virtual machine disk.

To delete a snapshot, a large amount of information needs to be read and written to a disk. This process can reduce virtual machine performance until consolidation is complete. Consolidating snapshots removes redundant disks, which improves virtual machine performance and saves storage space. The time it takes to delete snapshots and consolidate the snapshot files depends on the amount of data that the guest operating system writes to the virtual disks after you take the last snapshot. The required time is proportional to the amount of data the virtual machine is writing during consolidation if the virtual machine is powered on.

Failure of disk consolidation can reduce the performance of virtual machines. You can check whether any virtual machines require separate consolidation operations by viewing a list. For information about locating and viewing the consolidation state of multiple virtual machines and running a separate consolidation operation, see *vSphere Virtual Machine Administration*.

Delete

Use the **Delete** option to remove a single parent or child snapshot from the snapshot tree. **Delete** writes disk changes that occur between the state of the snapshot and the previous disk state to the parent snapshot.

Note: Deleting a single snapshot preserves the current state of the virtual machine and does not affect any other snapshot.

You can also use the **Delete** option to remove a corrupt snapshot and its files from an abandoned branch of the snapshot tree without merging them with the parent snapshot.

Delete All

Use the **Delete All** option to delete all snapshots from the Snapshot Manager. **Delete all** consolidates and writes the changes that occur between snapshots and the previous delta disk states to the base parent disk and merges them with the base virtual machine disk.

To prevent snapshot files from merging with the parent snapshot if, for example, an update or installation fails, first use the **Restore** command to restore to a previous snapshot. This action invalidates the snapshot delta disks and deletes the memory file. You can then use the **Delete** option to remove the snapshot and any associated files.

Delete a Snapshot in the VMware Host Client

You can use the Snapshot Manager to delete a single snapshot or all snapshots in a tree. Deleting a snapshot removes the snapshot from the Snapshot Manager. The snapshot files are consolidated and written to the parent snapshot disk and merge with the virtual machine base disk.

Deleting a snapshot does not change the virtual machine or other snapshots. Deleting a snapshot consolidates the changes between snapshots and previous disk states and writes all the data from the delta disk that contains the information about the deleted snapshot to the parent disk. When you delete the base parent snapshot, all changes merge with the base virtual machine disk.

To delete a snapshot, a large amount of information needs to be read and written to a disk. This process can reduce virtual machine performance until consolidation is complete. Consolidating snapshots removes redundant disks, which improves virtual machine performance and saves storage space. The time it takes to delete snapshots and consolidate the snapshot files depends on the amount of data that the guest operating system writes to the virtual disks after you take the last snapshot. The required time is proportional to the amount of data the virtual machine is writing during consolidation if the virtual machine is powered on.

Failure of disk consolidation can reduce the performance of virtual machines. You can check whether any virtual machines require separate consolidation operations by viewing a list. For information about locating and viewing the consolidation state of multiple virtual machines and running a separate consolidation operation, see *vSphere Virtual Machine Administration*.

Delete

Use the **Delete** option to remove a single parent or child snapshot from the snapshot tree. **Delete** writes disk changes that occur between the state of the snapshot and the previous disk state to the parent snapshot.

Note: Deleting a single snapshot preserves the current state of the virtual machine and does not affect any other snapshot.

You can also use the **Delete** option to remove a corrupt snapshot and its files from an abandoned branch of the snapshot tree without merging them with the parent snapshot.

Delete All

Use the **Delete All** option to delete all snapshots from the Snapshot Manager. **Delete all** consolidates and writes the changes that occur between snapshots and the previous delta disk states to the base parent disk and merges them with the base virtual machine disk.

To prevent snapshot files from merging with the parent snapshot if, for example, an update or installation fails, first use the **Restore** command to restore to a previous snapshot. This action invalidates the snapshot delta disks and deletes the memory file. You can then use the **Delete** option to remove the snapshot and any associated files.

Be careful not to accidentally delete a snapshot that you need. You cannot restore a deleted snapshot. For example, you might want to install several browsers, a, b, and c, and capture the virtual machine state after you install each browser. The first, or base snapshot, captures the virtual machine with browser a and the second snapshot captures browser b. If you restore the base snapshot that includes browser a and take a third snapshot to capture browser c and delete the snapshot that contains browser b, you cannot return to the virtual machine state that includes browser b.

1. Click **Virtual Machines** in the inventory.
2. Right-click the virtual machine in the list and select **Snapshots > Manage Snapshots**.
3. Click the snapshot that you want to delete and click **Delete snapshot**.
4. Optional: In the **Delete Snapshot** dialog box, select the **Remove all children snapshots** check-box to delete the selected snapshot together with all its children snapshots.
5. Click **Remove** to confirm the deletion.
6. Click **Close** to navigate out of the Snapshot Manager.

Why to Use Snapshot Manager in the VMware Host Client

You can review all snapshots for your virtual machines and use the Snapshot Manager to manage the snapshots.

After you take a snapshot, you can right-click a virtual machine and click **Revert to snapshot** to restore the virtual machine to the state of the snapshot at any time.

If you have a series of snapshots, you can use the Snapshot Manager to restore any parent or child snapshot. Subsequent child snapshots that you take from the restored snapshot create a branch in the snapshot tree. Use the Snapshot Manager to delete a snapshot from the tree.

Table 564: Snapshot Manager

Option	Description
Snapshot tree	Displays all snapshots for the virtual machine.
You are here icon	The You are here icon represents the current and active state of the virtual machine. The Restore , Delete , and Edit actions are deactivated for the You are here state.

Option	Description
Take, Restore, Delete, Edit	Snapshot options.
Details	Shows the snapshot name and description, the date you created the snapshot. The Console shows the power state of the virtual machine when a snapshot was taken. The Name, Description, and Created text boxes are blank if you do not select a snapshot.

Monitor a Virtual Machine in the VMware Host Client

You can monitor various performance aspects and keep track of the actions that take place on virtual machines that you create in the VMware Host Client.

View Virtual Machine Performance Charts in the VMware Host Client

You can view line charts with information about the resource usage of virtual machines that you create in the VMware Host Client.

1. Click **Virtual Machines** in the inventory.
2. Click a virtual machine from the list.
3. Expand the virtual machine in the inventory and click **Monitor**.
4. Click **Performance**.
5. To view the virtual machine resource usage for the last hour, select an option from the drop-down menu.
 - To view the percentage of CPU that the virtual machine used during the last hour, select **CPU usage**.
 - To view the memory that the host consumed during the last hour, select **Memory usage**.

View Virtual Machine Events in the VMware Host Client

Events are records of the actions that a user performs on a virtual machine. When you create a virtual machine in the VMware Host Client, you can view the events associated with the virtual machine.

Required privilege: **Read only**.

1. Click **Virtual Machines** in the inventory.
2. Click a virtual machine from the list.
3. Expand the virtual machine in the inventory and click **Monitor**.
4. Click **Events**.

A list of all virtual machine events is displayed.

5. Optional: Click an event from the list to view event details.
6. Optional: Use the filter controls above the list to filter the list.
7. Optional: Click a column heading to sort the list.

View Virtual Machine Tasks in the VMware Host Client

When you create a virtual machine in the VMware Host Client, you can view all virtual machine tasks and information about the task target, initiator, queue time, start time, result, and time of completion.

1. Click **Virtual Machines** in the inventory.
2. Click a virtual machine from the list.
3. Expand the virtual machine in the inventory and click **Monitor**.
4. Click **Tasks**.
5. Optional: Click on a task from the list to view task details.
6. Optional: Use the filter controls above the list to filter the list.
7. Optional: Click a column heading to sort the list.

View Virtual Machine Log Browser in the VMware Host Client

Generate and monitor logs for the host that you are managing by using the VMware Host Client. Use the logs to diagnose and troubleshoot various issues with your host environment.

1. Click **Virtual Machines** in the inventory.
2. Click a virtual machine from the list.
3. Expand the virtual machine in the inventory and click **Monitor**.
4. Click **Logs**.
5. Optional: Click **Generate support bundle** to consolidate all the logs for troubleshooting.
6. Right-click a log from the list and select **Open in new window** to view the log.

View Virtual Machine Notifications in the VMware Host Client

You can view virtual machine notifications and information about related tasks, which you can perform, for virtual machines that you create in the VMware Host Client.

1. Click **Virtual Machines** in the inventory.
2. Click a virtual machine from the list.
3. Expand the virtual machine in the inventory and click **Monitor**.
4. Click **Notifications**.
A list with all virtual machine notifications is displayed.

5. Optional: Click a notification to view details.
6. Optional: Click a notification and click **Actions** to view suggested tasks.

Configure Virtual Machine Hardware in the VMware Host Client

You can configure most virtual machine hardware settings when you create a virtual machine or after you create the virtual machine and install the guest operating system.

Virtual CPU Configuration and Limitations

You can set most of the CPU parameters when you create virtual machines or after you install the guest operating system. Some actions require that you power off the virtual machine before you change the settings.

VMware uses the following terminology. Understanding these terms can help you plan your strategy for CPU resource allocation.

CPU	The CPU, or processor, is the component of a computer system that performs the tasks required for computer applications to run. The CPU is the primary element that performs the computer functions. CPUs contain cores.
CPU Socket	A CPU socket is a physical connector on a computer motherboard that connects to a single physical CPU. Some motherboards have multiple sockets and can connect multiple multicore processors (CPUs).
Core	A core contains a unit containing an L1 cache and functional units needed to run applications. Cores can independently run applications or threads. One or more cores can exist on a single CPU.
Resource sharing	Shares specify the relative priority or importance of a virtual machine or resource pool. If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when the two virtual machines are competing for resources.
Resource allocation	You can change CPU resource allocation settings, such as shares, reservation, and limit, when available resource capacity does not meet demands. For example, if at year end, the workload on accounting increases, you can increase the accounting resource pool reserve.
vSphere Virtual Symmetric Multiprocessing (Virtual SMP)	Virtual SMP or vSphere Virtual Symmetric Multiprocessing is a feature that enables a single virtual machine to have multiple processors.

Virtual CPU Limitations

The maximum number of virtual CPUs that you can assign to a virtual machine is 768. The number of virtual CPUs depends on the number of logical CPUs on the host, and the type of guest operating system that is installed on the virtual machine.

Be aware of the following limitations:

- A virtual machine cannot have more virtual CPUs than the number of logical cores on the host. The number of logical cores is equal to the number of physical cores if hyperthreading is deactivated or two times that number if hyperthreading is activated.
- If a running virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine.

By contrast, if a running virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.

- The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.
- Not every guest operating system supports Virtual SMP, and guest operating systems that support this functionality might support fewer processors than are available on the host. For information about Virtual SMP support, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Configuring Multicore Virtual CPUs

VMware multicore virtual CPU support lets you control the number of cores per virtual socket in a virtual machine. This capability lets operating systems with socket restrictions use more of the host CPU cores, which increases overall performance.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

Using multicore virtual CPUs can be useful when you run operating systems or applications that can take advantage of only a limited number of CPU sockets.

You can configure a virtual machine with ESXi 7.0 Update 1 and later compatibility to have up to 768 virtual CPUs. A virtual machine cannot have more virtual CPUs than the actual number of logical CPUs on the host. The number of logical CPUs means the number of physical processor cores or two times that number if hyperthreading is enabled. For example, if a host has 128 logical CPUs, you can configure the virtual machine for 128 virtual CPUs.

You configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single-core CPU, dual-core CPU, tri-core CPU, and so on. Your selection determines the number of sockets that the virtual machine has.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

For more information about multicore CPUs, see the *vSphere Resource Management* documentation.

Virtual CPU Limitations

The maximum number of virtual CPUs that you can assign to a virtual machine is 768. The number of virtual CPUs depends on the number of logical CPUs on the host, and the type of guest operating system that is installed on the virtual machine.

Be aware of the following limitations:

- A virtual machine cannot have more virtual CPUs than the number of logical cores on the host. The number of logical cores is equal to the number of physical cores if hyperthreading is deactivated or two times that number if hyperthreading is activated.
- If a running virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a running virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.
- The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.
- Not every guest operating system supports Virtual SMP, and guest operating systems that support this functionality might support fewer processors than are available on the host. For information about Virtual SMP support, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Configuring Multicore Virtual CPUs

VMware multicore virtual CPU support lets you control the number of cores per virtual socket in a virtual machine. This capability lets operating systems with socket restrictions use more of the host CPU cores, which increases overall performance.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

Using multicore virtual CPUs can be useful when you run operating systems or applications that can take advantage of only a limited number of CPU sockets.

You can configure a virtual machine with ESXi 7.0 Update 1 and later compatibility to have up to 768 virtual CPUs. A virtual machine cannot have more virtual CPUs than the actual number of logical CPUs on the host. The number of logical CPUs means the number of physical processor cores or two times that number if hyperthreading is enabled. For example, if a host has 128 logical CPUs, you can configure the virtual machine for 128 virtual CPUs.

You configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single-core CPU, dual-core CPU, tri-core CPU, and so on. Your selection determines the number of sockets that the virtual machine has.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

For more information about multicore CPUs, see the *vSphere Resource Management* documentation.

Change the Number of Virtual CPUs

A virtual machine with ESXi 7.0 Update 1 and later compatibility can have up to 768 virtual CPUs. You can change the number of virtual CPUs while your virtual machine is powered off. If virtual CPU hot add is enabled, you can increase the number of virtual CPUs while the virtual machine is running.

- If CPU hot add is not enabled, power off the virtual machine before adding virtual CPUs.
- To hot add multicore CPUs, verify that the virtual machine is compatible with ESXi 5.0 and later.
- Verify that you have the **Virtual Machine > Configuration > Change CPU Count** privilege.

Virtual CPU hot add is supported for virtual machines with multicore CPU support and ESXi 5.0 and later compatibility. When the virtual machine is powered on and CPU hot add is enabled, you can hot add virtual CPUs to the running virtual machine. You can add only multiples of the number of cores per socket.

If a virtual machine has 128 virtual CPUs or less, you cannot use hot adding to further increase the number of virtual CPUs. To change the number of virtual CPUs beyond that limit, you must first power off the virtual machine. By contrast, if a virtual machine already has more than 128 virtual CPUs, you can use hot adding to further increase the number of virtual CPUs to up to 768.

The maximum number of virtual CPU sockets that a virtual machine can have is 128. If you want to configure a virtual machine with more than 128 virtual CPUs, you must use multicore virtual CPUs.

Important: When you configure your virtual machine for multicore virtual CPU settings, you must ensure that your configuration complies with the requirements of the guest operating system EULA.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.

2. On the **Virtual Hardware** tab, expand **CPU**.
3. From the **CPU** drop-down menu, select the number of cores.
4. From the **Cores Per Socket** drop-down menu, select the number of cores per socket and click **OK**.

Allocate CPU Resources in the VMware Host Client

To manage workload demands, you can change the amount of CPU resources allocated to a virtual machine by using the shares, reservations, and limits settings.

Power off the virtual machine.

A virtual machine has the following user-defined settings that affect its CPU resource allocation.

Limit	Places a limit on the consumption of CPU time for a virtual machine. This value is expressed in MHz or GHz.
Reservation	Specifies the guaranteed minimum allocation for a virtual machine. The reservation is expressed in MHz or GHz.
Shares	Each virtual machine is granted CPU shares. The more shares a virtual machine has, the more often it receives a time slice of a CPU when there is no CPU idle time. Shares represent a relative metric for allocating CPU capacity.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **CPU**, and allocate CPU capacity for the virtual machine.

Option	Description
Reservation	Guaranteed CPU allocation for this virtual machine.
Limit	Upper limit for this virtual machine's CPU allocation. Select Unlimited to specify no upper limit.
Shares	CPU shares for this virtual machine in relation to the parent's total. Sibling virtual machines share resources according to their relative share values bounded by the reservation and limit. Select Low , Normal , or High , which specify share values respectively in a 1:2:4 ratio. Select Custom to give each virtual machine a specific number of shares, which express a proportional weight.

4. Click **Save**.

Virtual Memory Configuration

You can add, change, or configure virtual machine memory resources or options to improve virtual machine performance. You can set most of the memory parameters during virtual machine creation or after the guest operating system is installed.

Some actions require that you power off the virtual machine before changing the settings.

The memory resource settings for a virtual machine determine how much of the host's memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size. ESXi hosts limit the memory resource use to the maximum amount useful for the virtual machine, so that you can accept the default of Unlimited memory resources.

Change the Memory Configuration

You can reconfigure the amount of memory allocated to a virtual machine to improve performance.

Verify that you have the **Virtual machine > Change Configuration > Change Memory** privilege on the virtual machine.

Minimum memory size is 4 MB for virtual machines that use BIOS firmware. Virtual machines that use EFI firmware require at least 96 MB of RAM or they cannot power on.

Maximum memory size for virtual machines that use BIOS firmware is 24560 GB. You must use EFI firmware for virtual machines with memory size greater than 6128 GB.

Maximum memory size for a virtual machine depends on the physical memory of the host and the virtual machine compatibility settings.

If the virtual machine memory is greater than the host memory size, swapping occurs, which can have a severe effect on virtual machine performance. The maximum for best performance represents the threshold above which the physical memory of the host is insufficient to run the virtual machine at full speed. This value fluctuates as conditions on the host change, for example, as virtual machines are powered on or off.

The memory size must be a multiple of 4 MB.

Table 565: Maximum Virtual Machine Memory

Introduced in Host Version	Virtual Machine Compatibility	Maximum Memory Size
9.0	9.0 and later	24560 GB
8.0 Update 3	8.0 Update 3 and later	24560 GB
8.0 Update 2	8.0 Update 2 and later	24560 GB
8.0 Update 1	8.0 Update 1 and later	24560 GB
8.0	8.0 and later	24560 GB
7.0 Update 3	7.0 Update 3 and later	24560 GB
7.0 Update 2	7.0 Update 2 and later	24560 GB
7.0 Update 1	7.0 Update 1 and later	24560 GB
7.0	7.0 and later	6128 GB
6.7 Update 2	6.7 Update 2 and later	6128 GB
6.7	6.7 and later	6128 GB
6.5	6.5 and later	6128 GB

The ESXi host version indicates when support began for the increased memory size. For example, the memory size of a virtual machine with 6.0 and later compatibility running on 6.5 is restricted to 4080 GB.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Memory** and change the memory configuration.
 - a) In the **Memory** text box, enter the amount of RAM to assign to the virtual machine.
 - b) Select whether the memory is specified in MB, GB or TB.
3. Click **OK**.

Allocate Memory Resources to a Virtual Machine in the VMware Host Client

You can change the amount of memory resources allocated to a virtual machine by using the shares, reservations, and limits settings. The host determines the appropriate amount of physical RAM to allocate to virtual machines based on these settings. You can assign a high or low shares value to a virtual machine, depending on its load and status.

Power off the virtual machine.

The following user-defined settings affect the memory resource allocation of a virtual machine.

Limit	Places a limit on the consumption of memory for a virtual machine. This value is expressed in megabytes.
Reservation	Specifies the guaranteed minimum allocation for a virtual machine. The reservation is expressed in megabytes. If the reservation cannot be met, the virtual machine will not power on.
Shares	Each virtual machine is granted a number of memory shares. The more shares a virtual machine has, the greater share of host memory it receives. Shares represent a relative metric for allocating memory capacity. For more information about share values, see the <i>vSphere Resource Management</i> documentation.

You cannot assign a reservation to a virtual machine that is larger than the virtual machine configured memory. If you give a virtual machine a large reservation and reduce the virtual machine configured memory size, the reservation is reduced to match the new configured memory size.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **Memory**, and allocate the memory capacity for the virtual machine.

Option	Description
Reservation	Guaranteed memory allocation for this virtual machine.
Limit	Upper limit for this virtual machine's memory allocation. Select Unlimited to specify no upper limit.
Shares	The values Low , Normal , High , and Custom are compared to the sum of all shares of all virtual machines on the server.

4. Click **Save**.

Change Memory Hot Add Settings

Memory hot add lets you add memory resources for a virtual machine while that virtual machine is powered on.

- Power off the virtual machine.

- Verify that the virtual machine has a guest operating system that supports memory hot add functionality.
- Verify that the virtual machine compatibility is ESXi 4.x and later.
- Verify that VMware Tools is installed.

Enabling memory hot add produces some memory overhead on the ESXi host for the virtual machine.

Note: If the ESXi host is of version 7.0 Update 2 and earlier, hot-adding memory to a virtual machine with NVIDIA vGPU requires that the ESXi host has a free vGPU slot. Starting with vSphere 7.0 Update 3, the source host does not require to have a free vGPU slot.

1. Right-click a virtual machine in the inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, expand **Memory**, and select **Enable** to enable adding memory to the virtual machine while it is powered on.
3. Click **OK**.

You can now add memory to a virtual machine, even if the virtual machine is turned on.

Add an NVDIMM device to a VM in the VMware Host Client

Add a virtual NVDIMM device to a virtual machine to enable it to use non-volatile, or persistent, computer memory. Non-volatile memory (NVM), or persistent memory (PMem), combines the high data transfer rates of the volatile memory with the persistence and resiliency of traditional storage. The virtual NVDIMM device is a virtual NVM device that can retain stored data through reboots or power source failures.

- Verify that the guest OS of the virtual machine supports PMem.
- Verify that the virtual hardware version is 14 or later.
- Verify that you have the **Datastore > Allocate space** privilege.
- Verify that the host or the cluster on which the virtual machine resides has available PMem resources.

Virtual machines consume the PMem resource of the host through a virtual non-volatile dual in-line memory module (NVDIMM) or through a virtual persistent memory disk.

For more information about persistent memory, see [Manage Persistent Memory](#)

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Under the **Virtual Hardware** tab, click **Add other device** and select **NVDIMM** from the drop-down menu.
The NVDIMM device appears in the Virtual Hardware devices list. Each virtual machine can have a maximum of 64 NVDIMM devices.
4. Configure the newly added NVDIMM device.
 - a) In the Virtual Hardware devices list, expand **New NVDIMM**.
 - b) Enter the size of the new NVDIMM device.

Note: You can change the size of the NVDIMM device at a later time. The virtual machine must be powered off.

- c) Select the NVDIMM controller location or leave the default one.
5. Click **Save** to close the wizard.

Network Virtual Machine Configuration

ESXi networking features enable communication between virtual machines on the same host, between virtual machines on different hosts, and between other virtual and physical machines.

The networking features also allow management of ESXi hosts and enable communication between VMkernel services, such as NFS, iSCSI, or vSphere vMotion, and the physical network. When you configure networking for a virtual machine,

you select or change an adapter type, a network connection, and whether to connect the network when the virtual machine powers on.

Network Adapter Basics

When you configure a virtual machine, you can add network adapters (NICs) and specify the adapter type.

Network Adapter Types

The type of network adapters that are available depend on the following factors:

- The virtual machine compatibility, which depends on the host that created or most recently updated it.
- Whether the virtual machine compatibility has been updated to the latest version for the current host.
- The guest operating system.

Supported NICs currently differ between an on-premises environment and VMware Cloud on AWS. The following NIC types are supported in an on-premises deployment:

E1000E	Emulated version of the Intel 82574 Gigabit Ethernet NIC. E1000E is the default adapter for Windows 8 and Windows Server 2012.
E1000	Emulated version of the Intel 82545EM Gigabit Ethernet NIC, with drivers available in most newer guest operating systems, including Windows XP and later and Linux versions 2.4.19 and later.
Flexible	Identifies itself as a Vlan adapter when a virtual machine boots, but initializes itself and functions as either a Vlan or a VMXNET adapter, depending on which driver initializes it. With VMware Tools installed, the VMXNET driver changes the Vlan adapter to the higher performance VMXNET adapter.
Vlan	Emulated version of the AMD 79C970 PCnet32 LANCE NIC, an older 10 Mbps NIC with drivers available in 32-bit legacy guest operating systems. A virtual machine configured with this network adapter can use its network immediately.
VMXNET	Optimized for performance in a virtual machine and has no physical counterpart. Because operating system vendors do not provide built-in drivers for this card, you must install VMware Tools to have a driver for the VMXNET network adapter available.
VMXNET 2 (Enhanced)	Based on the VMXNET adapter but provides high-performance features commonly used on modern networks, such as jumbo frames and hardware offloads. VMXNET 2 (Enhanced) is available only for some guest operating systems on ESX/ ESXi 3.5 and later.
VMXNET 3	A paravirtualized NIC designed for performance. VMXNET 3 offers all the features available in VMXNET 2 and adds several new features, such as multiqueue support (also known as Receive Side Scaling in Windows), IPv6 offloads, and MSI/MSI-X interrupt delivery. VMXNET 3 is not related to VMXNET or VMXNET 2.
PVRDMA	A paravirtualized NIC that supports remote direct memory access (RDMA) between virtual machines through the OFED verbs API. All virtual machines must have a PVRDMA device and should be connected to a distributed switch. PVRDMA supports VMware vSphere vMotion and snapshot technology. It is available in virtual machines with hardware version 13 and guest operating system Linux kernel 4.6 and later.

SR-IOV passthrough

For information about assigning an PVRDMA network adapter to a virtual machine, see the *vSphere Networking* documentation.

Representation of a virtual function (VF) on a physical NIC with SR-IOV support. The virtual machine and the physical adapter exchange data without using the VMkernel as an intermediary. This adapter type is suitable for virtual machines where latency might cause failure or that require more CPU resources.

SR-IOV passthrough is available in ESXi 6.0 and later for guest operating systems Red Hat Enterprise Linux 6 and later, and Windows Server 2008 R2 with SP2. An operating system release might contain a default VF driver for certain NICs, while for others you must download and install it from a location provided by the vendor of the NIC or of the host.

For information about assigning an SR-IOV passthrough network adapter to a virtual machine, see the *vSphere Networking* documentation.

For network adapter compatibility considerations, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Legacy Network Adapters and ESXi Virtual Hardware Versions

The default network adapter types for all legacy virtual machines depend on the adapters available and compatible to the guest operating system and the version of virtual hardware on which the virtual machine was created.

If you do not upgrade a virtual machine to use a virtual hardware version, your adapter settings remain unchanged. If you upgrade your virtual machine to take advantage of newer virtual hardware, your default adapter settings will likely change to be compatible with the guest operating system and upgraded host hardware.

To verify the network adapters that are available to your supported guest operating system for a particular version of vSphere ESXi, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Network Adapters and Legacy Virtual Machines

Legacy virtual machines are virtual machines that are supported by the product in use, but are not current for that product.

The default network adapter types for all legacy virtual machines depend on the adapters available and compatible to the guest operating system and the version of virtual hardware on which the virtual machine was created.

If you do not upgrade a virtual machine to correspond with an upgrade to a newer version of an ESXi host, your adapter settings remain unchanged. If you upgrade your virtual machine to take advantage of newer virtual hardware, your default adapter settings will likely change to be compatible with the guest operating system and upgraded host hardware.

To verify the network adapters that are available to your supported guest operating system for a particular version of vSphere ESXi, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Change the Configuration of the Virtual Network Adapter in the VMware Host Client

You can configure the power-on connection setting, the MAC address, and the network connection of the virtual network adapter of a virtual machine.

Required privileges:

- **Virtual Machine > Configuration > Modify device settings** for editing the MAC address and network.

- **Virtual Machine > Interaction > Device connection** for changing **Connect** and **Connect at power on**.
- **Network > Assign network**

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Click the **Virtual Hardware** tab and select the appropriate Network Adapter (NIC) from the hardware list.
4. Optional: To connect the virtual NIC when the virtual machine is powered on, select **Connect at power on**.
5. Optional: Select the adapter type from the **Adapter Type** drop-down menu.
6. Select an option for MAC address configuration.

Option	Description
Automatic	vSphere assigns a MAC address automatically.
Manual	Enter the MAC address to use.

7. Click **Save**.

Add a Network Adapter to a Virtual Machine in the VMware Host Client

When you add a network adapter (NIC) to a virtual machine, you must select the adapter type, the network connection, and whether the device connects when the virtual machine is powered on.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Click the **Virtual Hardware** tab and click **Add network adapter**.
4. In the network connection panel, select either a network with a specific label or a legacy network.
5. Optional: To configure the virtual NIC to connect when the virtual machine is powered on, select **Connect at power on**.
6. Click **Save**.

Virtual Disk Configuration

You can add large-capacity virtual disks to virtual machines and add more space to existing disks, even when the virtual machine is running.

You can set most of the virtual disk parameters during virtual machine creation or after you install the guest operating system.

You can store virtual machine data in a new virtual disk, an existing virtual disk, or a mapped SAN LUN. A virtual disk appears as a single hard disk to the guest operating system. The virtual disk is composed of one or more files on the host file system. You can copy or move virtual disks on the same hosts or between hosts.

For virtual machines running on an ESXi host, you can store virtual machine data directly on a SAN LUN instead of using a virtual disk file. This option is useful if in your virtual machines you run applications that must detect the physical characteristics of the storage device. Mapping a SAN LUN also allows you to use existing SAN commands to manage storage for the disk.

When you map a LUN to a VMFS volume, vCenter Server or the ESXi host creates a raw device mapping (RDM) file that points to the raw LUN. Encapsulating disk information in a file allows vCenter Server or the ESXi host to lock the LUN so that only one virtual machine can write to it. This file has a `.vmdk` extension, but the file contains only disk information that describes the mapping to the LUN on the ESXi system. The actual data is stored on the LUN. You cannot deploy a virtual machine from a template and store its data on a LUN. You can store its data only in a virtual disk file.

The amount of free space in the datastore is always changing. Ensure that you leave sufficient space for virtual machine creation and other virtual machine operations, such as growth of sparse files, snapshots, and so on. To review space utilization for the datastore by file type, see the *vSphere Monitoring and Performance* documentation.

Thin provisioning lets you create sparse files with blocks that are allocated upon first access, which allows the datastore to be over-provisioned. The sparse files can continue growing and fill the datastore. If the datastore runs out of disk space while the virtual machine is running, it can cause the virtual machine to stop functioning.

About Virtual Disk Provisioning Policies

When you perform certain virtual machine management operations, you can specify a provisioning policy for the virtual disk file. The operations include creating a virtual disk, cloning a virtual machine to a template, or migrating a virtual machine.

NFS datastores with Hardware Acceleration and VMFS datastores support the following disk provisioning policies. On NFS datastores that do not support Hardware Acceleration, only thin format is available.

You can use Storage vMotion or cross-host Storage vMotion to transform virtual disks from one format to another.

Thick Provision Lazy Zeroed

Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand later on first write from the virtual machine. Virtual machines do not read stale data from the physical device.

Thick Provision Eager Zeroed

A type of thick virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the thick provision lazy zeroed format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take longer to create virtual disks in this format than to create other types of disks. Increasing the size of an Eager Zeroed Thick virtual disk causes a significant stun time for the virtual machine.

Thin Provision

Use this format to save storage space. For the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the virtual disk size. However, the thin disk starts small and at first, uses only as much datastore space as the disk needs for its initial operations. If the thin disk needs more space later, it can grow to its maximum capacity and occupy the entire datastore space provisioned to it. Thin provisioning is the fastest method to create a virtual disk because it creates a disk with just the header information. It does not allocate or zero out storage blocks. Storage blocks are allocated and zeroed out when they are first accessed.

Note: If a virtual disk supports clustering solutions such as Fault Tolerance, do not make the disk thin.

Change the Virtual Disk Configuration in the VMware Host Client

If you run out of disk space, you can increase the size of the disk. You can change the virtual device node and the persistence mode of virtual disk configuration of a virtual machine.

Power off the virtual machine.

Verify that you have the following privileges:

- **Virtual machine > Configuration > Modify device settings** on the virtual machine.
- **Virtual machine > Configuration > Extend virtual disk** on the virtual machine.
- **Datastore > Allocate space** on the datastore.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand the hard disk to view all disk options.
4. Optional: To change the size of the disk, enter a new value in the text box and select the units from the drop-down menu.
5. Optional: To change the way that disks are affected by snapshots, select a disk mode from the **Disk Mode** drop-down menu.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent-Persistent	Disks in persistent mode behave like conventional disks on your physical computer. All data written to a disk in persistent mode are written permanently to the disk.
Independent-Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. With nonpersistent mode, you can restart the virtual machine with a virtual disk in the same state every time. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset the virtual machine.

6. Click **Save**.

Add a New Standard Hard Disk to a Virtual Machine in the VMware Host Client

You can add a virtual hard disk to an existing virtual machine, or you can add a hard disk when you customize the virtual machine hardware during the virtual machine creation process. For example, you might need to provide additional disk space for an existing virtual machine with a heavy work load. During virtual machine creation, you might want to add a hard disk that is preconfigured as a boot disk.

- Verify that you are familiar with configuration options and caveats for adding virtual hard disks. See [Virtual Disk Configuration](#).
- Before you add disks larger than 2TB in size to a virtual machine, see *vSphere Virtual Machine Administration*.
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the destination folder or datastore.

Power off the virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Optional: To delete an existing hard disk, move your pointer over the disk and click the **Remove** icon (X).
The disk is removed from the virtual machine. If other virtual machines share the disk, the disk files are not deleted.
4. On the **Virtual Hardware** tab, select **Add hard disk** and select **New standard hard disk** from the drop-down menu.
The hard disk appears in the Virtual Hardware devices list.

5. Expand **New Hard disk**.
6. Optional: Enter a value for the hard disk size and select the units from the drop-down menu.
7. Select the datastore location where you want to store the virtual machine files.
8. Select the format for the virtual machine disk.

Option	Description
Thick Provision Lazy Zeroed	Create a virtual disk in a default thick format. Space required for the virtual disk is allocated during creation. Any data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine.
Thick Provision Eager Zeroed	Create a thick disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out during creation. It might take much longer to create disks in this format than to create other types of disks.
Thin Provision	Use the thin provisioned format. At first, a thin provisioned disk uses only as much datastore space as the disk initially requires. If the thin disk needs more space later, it can grow to the maximum capacity allocated to it.

9. In the **Shares** drop-down menu, select a value for the shares to allocate to the virtual disk.
Shares is a value that represents the relative metric for controlling disk bandwidth. The values Low, Normal, High, and Custom are compared to the sum of all shares of all virtual machines on the host.
10. If you selected **Custom**, enter a number of shares in the text box.
11. In the **Limit IOPs** box, enter the upper limit of storage resources to allocate to the virtual machine, or select **Unlimited**.
This value is the upper limit of I/O operations per second allocated to the virtual disk.
12. Accept the default or select a different virtual device node.
In most cases, you can accept the default device node. For a hard disk, using a nondefault device node makes controlling the boot order or having different SCSI controller types easier. For example, you might want to boot from an LSI Logic controller and share a data disk with another virtual machine that is using a Buslogic controller with bus sharing turned on.
13. Optional: Select a disk mode.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent-Persistent	Disks in persistent mode behave like conventional physical computer disks. All data written to a disk in persistent mode are written permanently to the disk.
Independent-Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. The virtual disk returns to the same state every time you restart the virtual machine. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset.

14. Click **Save**.

Add an Existing Hard Disk to a Virtual Machine in the VMware Host Client

You can add an existing virtual hard disk to a virtual machine when you customize the virtual machine hardware during the virtual machine creation process or after the virtual machine is created. For example, you might want to add an existing hard disk that is preconfigured as a boot disk.

- Verify that you are familiar with controller and virtual device node behavior for different virtual hard disk configurations.
- Verify that you have the **Virtual machine > Change Configuration > Add existing disk** privilege on the destination folder or datastore.

Power off the virtual machine.

During virtual machine creation, a hard disk and a SCSI or SATA controller are added to the virtual machine by default, based on the guest operating system that you select. If this disk does not meet your needs, you can remove it and add an existing hard disk at the end of the creation process.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add hard disk** and select **Existing hard disk** from the drop-down menu.
4. Optional: To delete an existing hard disk, move your pointer over the disk and click the **Remove** icon (X).
The disk is removed from the virtual machine. If other virtual machines share the disk, the disk files are not deleted.
5. In the Datastore column, expand a datastore, select a virtual machine folder, and select the disk to add.
The disk file appears in the Contents column. The **File type** menu shows the compatibility file types for this disk.
6. Click **Select** and click **Save** to add the existing hard disk.

Add a Persistent Memory Disk in the Host Client

You can add a virtual hard disk to an existing virtual machine, or you can add a hard disk when you customize the virtual machine hardware during the virtual machine creation process. For example, you might need to provide additional disk space for an existing virtual machine with a heavy work load. During virtual machine creation, you might want to add a hard disk that is preconfigured as a boot disk.

- Verify that you are familiar with configuration options and caveats for adding virtual hard disks. See [Virtual Disk Configuration](#).
- Before you add disks larger than 2TB in size to a virtual machine, see *vSphere Virtual Machine Administration*.
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the destination folder or datastore.

Power off the virtual machine.

During virtual machine creation, a hard disk and a SCSI or SATA controller are added to the virtual machine by default, based on the guest operating system that you select. If this disk does not meet your needs, you can remove it and add an existing hard disk at the end of the creation process.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add hard disk** and select **New persistent memory disk** from the drop-down menu.
The hard disk appears in the Virtual Hardware devices list. By default, the disk is stored on the host-local PMem datastore and you cannot change the datastore.

4. Optional: Configure the settings for the new hard disk and click **Save** to close the wizard.

- a) Expand **New Hard disk**.
- b) Enter a value for the hard disk size and select the units from the drop-down menu.

Note: All persistent memory hard disks and NVDIMM modules that you add to the virtual machine share the same PMem resources. So, you must adjust the size of the newly added persistent memory devices in accordance with the amount of the PMem available to the host. If any part of the configuration requires attention, the wizard alerts you.

- c) From the **Shares** drop-down menu, select a value for the shares to allocate to the virtual disk.
Shares is a value that represents the relative metric for controlling disk bandwidth. The values Low, Normal, High, and Custom are compared to the sum of all shares of all virtual machines on the host.
- d) From the **Controller location** drop-down menu, select the location of the controller that the new hard disk uses.
- e) Select a disk mode.

Option	Description
Dependent	Dependent disks are included in snapshots.
Independent-Persistent	Disks in persistent mode behave like conventional physical computer disks. All data written to a disk in persistent mode are written permanently to the disk.
Independent-Nonpersistent	Changes to disks in nonpersistent mode are discarded when you power off or reset the virtual machine. The virtual disk returns to the same state every time you restart the virtual machine. Changes to the disk are written to and read from a redo log file that is deleted when you power off or reset.

Use Disk Shares to Prioritize Virtual Machines in the VMware Host Client

You can change the disk resources for a virtual machine. If multiple virtual machines access the same VMFS datastore and the same logical unit number (LUN), use disk shares to prioritize the level of access that virtual machines have to resources. Disk shares distinguish high-priority from low-priority virtual machines.

Power off the virtual machine.

You can allocate the I/O bandwidth of the host to the virtual hard disks of a virtual machine. You cannot pool disk I/O across a cluster.

The shares value represents the relative metric for controlling disk bandwidth to all virtual machines.

Disk shares are relevant only within a given host. The shares assigned to virtual machines on one host have no effect on virtual machines on other hosts.

You can select an IOP limit, which sets an upper limit for storage resources that are allocated to a virtual machine. IOPs are the number of I/O operations per second.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand the hard disk to view the disk options.
4. In the **Shares** drop-down menu, select a value for the shares to allocate to the virtual machine.
5. If you selected **Custom**, enter a number of shares in the text box.
6. In the **Limit - IOPs** text box, enter the upper limit of storage resources to allocate to the virtual machine, or select **Unlimited**.
7. Click **Save**.

Virtual Machine Controller Configuration in the VMware Host Client

In the VMware Host Client, you can add various controllers to virtual machines, such as USB controllers, SCSI controllers, Paravirtual SCSI controllers, and SATA controllers. You can also change the SCSI Bus sharing configuration and the SCSI controller type.

Add a USB Controller to a Virtual Machine

To support USB passthrough from an ESXi host or from a client computer to a virtual machine, you can add a USB controller to the virtual machine.

- Verify that the ESXi host has USB controller hardware and modules that support USB 3.2, 2.0, and 3.1 devices.
- Verify that the client computers have USB controller hardware and modules that support USB 3.2, 2.0, and 3.1 devices present.
- To use the xHCI controller on a Linux guest operating system, verify that the Linux kernel version is 2.6.35 or later.
- Verify that the virtual machine is powered on.
- Required Privilege (ESXi host passthrough): **Virtual Machine > Configuration > Add or Remove Device**.

In the vSphere Client, you can add one xHCI controller and one EHCI+UHCI controller.

- From hardware version 11 to hardware version 16, the supported number of root hub ports per xHCI controller is eight (four local USB 3.1 SuperSpeed ports and four logical USB 2.0 ports).
- From hardware version 17 to hardware version 20, the supported number of root hub ports per xHCI controller is eight (four local USB 3.1 SuperSpeedPlus ports and four logical USB 2.0 ports).
- With hardware version 21, the supported number of root hub ports per xHCI controller is eight (four local USB 3.2 Gen 2x2 ports and four logical USB 2.0 ports).

The conditions for adding a controller vary, depending on the device version, the type of passthrough (host or client computer), and the guest operating system.

Table 566: USB Controller Support

Controller Type	Supported USB Device Version	Supported for Passthrough from ESXi host to a VM	Supported for Passthrough from Client Computer to a VM
EHCI+UHCI	2.0	Yes	Yes
xHCI	3.2, 3.1, 2.0	Yes USB 3.2, 3.1, and 2.0	Yes

Controller Type	Supported USB Device Version	Supported for Passthrough from ESXi host to a VM	Supported for Passthrough from Client Computer to a VM
			Windows 8 or later, Windows Server 2012 and later, or a Linux guest operating system with a 2.6.35 or later kernel.

For Mac OS X systems, the EHCI+UHCI controller is enabled by default and is required for access to a USB mouse and keyboard.

For virtual machines with Windows or Linux guest operating systems, you can add one or two controllers of different types. You cannot add two controllers of the same type.

For USB passthrough from an host to a virtual machine, the USB arbitrator can monitor a maximum of 15 USB controllers. If your system includes more than 15 controllers and you connect USB devices to them, the devices are not available to the virtual machine.

1. Right-click a virtual machine from the vSphere inventory and select **Edit Settings**.
2. On the **Virtual Hardware** tab, click **Add New Device** and from the drop-down menu select **USB Controller**.
The controller appears in the **Virtual Hardware** devices list.
3. To change the USB controller type, expand **New USB Controller**.
If compatibility errors appear, you must fix them before you can add the controller.
4. Click **OK**.

Add one or more USB devices to the virtual machine.

Add SCSI Controllers in the VMware Host Client

You can add SCSI controllers to an existing virtual machine by adding hard disks on unused SCSI Bus numbers.

Power off the virtual machine.

Adding a new hard disk on an unused SCSI Bus number creates a new SCSI controller.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add hard disk** and select **New hard disk** from the drop-down menu.
4. Expand the hard disk to view all options.
5. In the **Controller location** section, select an unused SCSI Bus number from the drop-down menu.
For example, bus and device numbers 0:0 - 0:15 are used by the initial SCSI controller. The second SCSI controller uses bus and device numbers 1:0 - 1:15.
6. Click **Save**.

The new hard disk and new SCSI controller are simultaneously created.

Change the SCSI Bus Sharing Configuration in the VMware Host Client

You can set the type of SCSI bus sharing for a virtual machine and indicate whether to share the SCSI bus. Depending on the type of sharing, virtual machines can access the same virtual disk simultaneously on the same server or on any server.

Power off the virtual machine.

You can change the SCSI controller configuration for a virtual machine only if the virtual machine is on an ESXi host.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand the SCSI Controller that you want to edit.
4. Select the type of sharing in the **SCSI Bus Sharing** list.

Option	Description
None	Virtual disks cannot be shared by other virtual machines.
Virtual	Virtual disks can be shared by virtual machines on the same server.
Physical	Virtual disks can be shared by virtual machines on any server.

5. Click **Save**.

Change the SCSI Controller Type in the VMware Host Client

You can attach virtual disks and RDMs to virtual machines by configuring virtual SCSI controller on the virtual machines.

Power off the virtual machine.

The choice of SCSI controller does not affect whether your virtual disk is an IDE or SCSI disk. The IDE adapter is always ATAPI. The default for your guest operating system is already selected. Older guest operating systems have BusLogic adapter as their default controller.

If you create an LSI Logic virtual machine and add a virtual disk that uses BusLogic adapters, the virtual machine boots from the BusLogic adapters disk. LSI Logic SAS is available only for virtual machines with hardware version 7 or later. Disks with snapshots might not experience performance gains when used on LSI Logic SAS, VMware Paravirtual, and LSI Logic Parallel adapters.

CAUTION: Changing the SCSI controller type might result in a virtual machine boot failure.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. Click the **Virtual Hardware** tab and expand a SCSI controller.
4. Select a SCSI controller type from the drop-down menu.
5. Click **Save**.

About VMware Paravirtual SCSI Controllers

VMware Paravirtual SCSI controllers are high performance storage controllers that can result in greater throughput and lower CPU use. These controllers are best suited for high performance storage environments.

VMware Paravirtual SCSI controllers are available for virtual machines with ESXi 4.x and later compatibility. Disks on such controllers might not experience optimal performance gains if they have snapshots or if memory on the ESXi host is over committed. This behavior does not mitigate the overall performance gain of using VMware Paravirtual SCSI controllers as compared to other SCSI controller options.

For platform support for VMware Paravirtual SCSI controllers, see the *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

Add a Paravirtual SCSI Controller in the VMware Host Client

You can add a VMware Paravirtual SCSI high performance storage controller to provide greater throughput and lower CPU utilization.

- Verify that the virtual machine has a guest operating system with VMware Tools installed.
- Verify that the virtual machine has hardware version 7 or later.
- Familiarize yourself with VMware Paravirtual SCSI limitations. See *vSphere Virtual Machine Administration*.
- To access boot disk devices attached to a VMware Paravirtual SCSI controller, verify that the virtual machine has a Windows 2003 or Windows 2008 guest operating system.
- In some operating systems, before you change the controller type you must create a virtual machine with an LSI Logic controller and install VMware Tools.

Power off the virtual machine.

VMware Paravirtual SCSI controllers are best suited for environments, especially SAN environments, that run I/O-intensive applications.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, click **Add other device** and select **SCSI Controller** from the drop-down menu.
The new SCSI Controllers appear in the Hardware list.
4. Click the **New SCSI Controller** and select **VMware Paravirtual** from the drop-down menu.
5. Click **Save**.

Add a SATA Controller to a Virtual Machine in the VMware Host Client

If a virtual machine has multiple hard disks or CD/DVD-ROM devices, you can add up to three additional SATA controllers to assign the devices to. When you assign the devices to different controllers, you can improve performance and avoid data traffic congestion. You can also add controllers if you need to exceed the thirty-device limit for a single controller.

- Verify that the virtual machine compatibility is ESXi 5.5 and later.
- Verify that you are familiar with storage controller behavior and limitations. See *vSphere Virtual Machine Administration*.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
- Power off the virtual machine.

You can boot virtual machines from SATA controllers and use them for large-capacity virtual hard disks.

Not all guest operating systems support AHCI SATA controllers. Typically, when you create virtual machines with ESXi 5.5 and later compatibility and Mac OS X guest operating systems, a SATA controller is added by default for the virtual hard disk and CD/DVD-ROM devices. Most guest operating systems, including Windows Vista and later, have a default SATA controller for CD/DVD-ROM devices. For verification, see the appropriate *VMware Compatibility Guide* at <http://www.vmware.com/resources/compatibility>.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device**, and select **SATA Controller** from the drop-down menu.
The SATA controller appears in the hardware list.

- Click **Save**.

Add an NVMe Controller in the VMware Host Client

If a virtual machine has multiple hard disks, you can add up to four virtual NVMe controllers to which to assign the disks. Using a NVMe controller significantly reduces the software overhead for processing guest OS I/O, as compared to AHCI SATA or SCSI controllers.

- Verify that the virtual machine has a guest operating system that supports NVMe.
- Verify that the virtual machine compatibility is ESXi 6.5 or later.
- Verify that you are familiar with storage controllers behavior and limitations. For more information, see the *Virtual Machine Administration* guide.
- Verify that you have the **Virtual machine > Change Configuration > Add new disk** privilege on the virtual machine.

NVMe controllers perform best with virtual disks on an all-flash disk array, local NVMe SSD, and PMem storage.

- Click **Virtual Machines** in the inventory.
- Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
- Under the **Virtual Hardware** tab, click the **Add other device** icon and select **NVMe controller** from the drop-down menu.

A new NVMe controller is added to the virtual machine.

You can add a hard disk to the virtual machine and assign it to the NVMe controller.

Other Virtual Machine Device Configuration in the VMware Host Client

In addition to configuring virtual machine CPU and Memory, and adding hard disks and virtual network adapters, you can also add and configure virtual hardware, such as DVD/CD-ROM drives, floppy drives, and SCSI devices. You can also add a virtual Watchdog Timer (VWDT) device, Precision Clock device, and PCI devices.

Add a CD or DVD Drive to a Virtual Machine in the VMware Host Client

You can use a physical drive on a client or host, or you can use an ISO image to add a CD/DVD drive to a virtual machine.

Power off the virtual machine.

If you want to add a CD/DVD drive that is backed up by USB CD/DVD drive on the host, you must add the drive as a SCSI device. Hot adding or removing SCSI devices from an ESXi host is not supported.

- Click **Virtual Machines** in the inventory.
- Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
- On the **Virtual Hardware** tab, select **Add other device** and select **CD/DVD Drive** from the drop-down menu.
- Expand **CD/DVD drive** and select an option.

Option	Description
Use physical drive	- Select Client Device as the location. - From the Device Mode drop-down menu, select Emulate CD-ROM or Passthrough CD-ROM.
Use ISO Image	- Select Datastore ISO File as the location. - Enter the path and filename for the image file, or click Browse to navigate to the file.

5. If you do not want the CD-ROM drive to connect when the virtual machine starts, deselect **Connect at power on**.
6. Select the virtual device node that the drive uses in the virtual machine.
7. Click **Save**.

Add a Floppy Drive to a Virtual Machine in the VMware Host Client

Use a physical floppy drive or a floppy image to add a floppy drive to a virtual machine.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.

ESXi does not support floppy drives that are backed up by a physical floppy drive on the host.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **Floppy Drive** from the drop-down menu.
The Floppy drive appears in the hardware list.
4. Expand **Floppy drive** and select the type of device you want to use.

Option	Description
Client Device	Select this option to connect the floppy device to a physical floppy device or a .flp floppy image on the system from which you access the VMware Host Client.
Use existing floppy image	1. Select this option to connect the virtual device to an existing image of a floppy drive on a datastore accessible to the host. 2. Click Browse and select the floppy image.

5. Optional: Select **Connect at power on** to configure the device to connect when the virtual machine powers on.
6. Click **Save**.

Add a USB Device to a Virtual Machine in the VMware Host Client

By using the VMware Host Client, you can add a USB device to a virtual machine.

- Verify that a USB controller is present. See [Add a USB Controller to a Virtual Machine](#).
- Add a physical USB device to the ESXi host where the virtual machine is located by plugging the USB device into the host.

Note: If the ESXi host does not have available USB devices, you cannot add a USB device to the virtual machine.


1. Click **Virtual Machines** in the inventory.
2. Right-click the virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **USB device** from the drop-down menu.
The USB device appears in the hardware list of available hardware devices for the virtual machine.

4. From the **USB device** drop-down, select which USB device to add to the virtual machine.
5. Click **Save**.

Add a Sound Controller to a Virtual Machine in the VMware Host Client

By using the VMware Host Client, you can add a sound controller to a virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **Sound controller** from the drop-down menu.



The sound controller appears in the list of available hardware devices for the virtual machine.

4. From the **Sound card** drop-down menu, select which sound controller to connect to the virtual machine.
5. Click **Save**.

Parallel and Serial Port Configuration in the VMware Host Client

Parallel and serial ports are interfaces for connecting peripherals to the virtual machine. The virtual serial port can connect to a physical serial port or to a file on the host computer.

You can also use the virtual serial port to establish a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. You can add parallel and serial ports and change the serial port configuration.

Add a Serial Port to a Virtual Machine in the VMware Host Client

A virtual machine can use up to four virtual serial ports. You can connect the virtual serial port to a physical serial port or to a file on the host computer. You can also use a host-side-named pipe to set up a direct connection between two virtual machines or a connection between a virtual machine and an application on the host computer. In addition, you can use a port or a Virtual Serial Port Concentrator (vSPC) URI to connect a serial port over the network.

- Familiarize yourself with the different media types that the port can access, vSPC connections, and any conditions that might apply. See *vSphere Virtual Machine Administration*.
- To connect a serial port over a network, add a Firewall rule set. See *vSphere Virtual Machine Administration*.
- Required privilege: **Virtual Machine > Configuration > Add or Remove Device**
Power off the virtual machine.

- Click **Virtual Machines** in the inventory.
- Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
- On the **Virtual Hardware** tab, select **Add other device** and select **Serial Port**.
The Serial Port appears in the hardware list.
- In the hardware list, expand the serial port and select the type of media port to access.

Option	Description
Use output file	Browse to the location of the file on the host to store the output of the virtual serial port.
Use physical serial port	Select the port from the drop-down menu.
Use named pipe	- Type a name for the pipe in the Pipe name field - Select the Near End and Far End of the pipe from the drop-down menus.
Use network	- From the Direction drop-down menu, select Server or Client. - Type the port URI. The URI is the remote end of the serial port to which the serial port of the virtual machine should connect. - If vSPC is used as an intermediate step to access all virtual machines through a single IP address, select Use Virtual Serial Port Concentrator and enter the vSPC URI location.

- Optional: Deselect **Connect at power on** if you do not want the parallel port device to connect when the virtual machine powers on.
- Click **Save**.

Establishing Serial Port Network Connections to a Client or Server Without Authentication Parameters

If you do not use vSPC and you configure your virtual machine with a serial port connected as a server with a `telnet://:12345` URI, you can connect to your virtual machine's serial port from your Linux or Windows operating system.

```
telnetyourESXiServerIPAddress12345
```

Similarly, if you run the Telnet Server on your Linux system on port 23 (`telnet://yourLinuxBox:23`), you configure the virtual machine as a client URI.

```
telnet://yourLinuxBox:23
```

The virtual machine initiates the connection to your Linux system on port 23.

Add a Parallel Port to a Virtual Machine in the VMware Host Client

To connect peripheral devices to virtual machines, such as printers and scanners, you can use a parallel port. You send the output of such devices to a file on the host computer.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.

Note: To add a parallel port to a virtual machine that runs on an ESXi 4.1 or earlier host, you can also select to send output to a physical parallel port on the host. This option is not available with ESXi 5.0 and later host versions.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and select **Parallel Port**.
The parallel port appears in the hardware list.
4. Expand the parallel port and in the Connection field, browse to a folder to create the file in.
The file path appears in the **Connection** text box.
5. Optional: Select **Connect at power on** to configure the device to connect when the virtual machine powers on.
6. Click **Save**.

How do I Add a Virtual Watchdog Timer Device to a Virtual Machine

To ensure self-reliance related to the system performance within a virtual machine, you can add a virtual Watchdog Timer (VWDT) device.

If the guest operating system stops responding and cannot recover on its own due to software glitches or errors, the VWDT waits for a predefined period of time and then restarts the system.

You can enable the VWDT to start either by the guest operating system, or by the BIOS or EFI firmware. If you chose the VWDT to start by the BIOS or EFI firmware, it starts before the guest operating system boots.

The VWDT has an important role in guest-based clustering solutions where each virtual machine in the cluster can recover on its own if it fails.

Add a Virtual Watchdog Timer Device to a Virtual Machine in the VMware Host Client

You can add a virtual Watchdog Timer device to a virtual machine to prevent the virtual machine from a guest operating system failure for an extended period of time.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
- Verify that the guest operating system of the virtual machine supports the VWDT device.
- Verify that the virtual hardware version is 17.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, select **Add other device** and click **Watchdog Timer**.
The Watchdog Timer device appears in the hardware list.

4. Optional: Select **Start with BIOS/EFI boot** to start the watchdog timer by the BIOS or EFI firmware.
When you select this option, the VMDT device starts before the guest operating system. If the Guest OS boot takes too long or does not support the watchdog timer, the device might constantly restart the virtual machine.
5. Click **Save**.

Add a Precision Clock Device to a Virtual Machine in the VMware Host Client

A precision clock is a virtual device that runs on a virtual machine and accesses the system time of a host. By adding a precision clock to a virtual machine, you ensure time synchronization and high precision timestamping.

- Power off the virtual machine.
 - Verify that the virtual hardware version is 17.
 - Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
 - Verify that you have the **Virtual machine > Configuration > Modify device settings** privilege on the virtual machine.
1. In the VMware Host Client inventory, click **Virtual Machines**.
 2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
 3. On the **Virtual Hardware** tab, click **Add other device** and select **Precision Clock**.
The precision clock device appears in the hardware list.
 4. Optional: Select the time synchronization protocol.
 5. Click **Save**.

Add a PCI Device to a Virtual Machine in the VMware Host Client

DirectPath I/O allows the guest operating system of a virtual machine to access the physical PCI and PCIe devices connected to a host directly. By using this technology, you can connect each virtual machine to up to sixteen physical PCI devices.

- Power off the virtual machine.
- Verify that you have the **Virtual machine > Configuration > Add or remove device** privilege on the virtual machine.
- Verify that the PCI devices are connected to the host and marked as available for passthrough.
- If you want to add a dynamic PCI device to a virtual machine, verify that the virtual hardware version is 17.

You can use Dynamic DirectPath I/O to assign multiple PCI passthrough devices to a virtual machine. Starting with vSphere 7.0, you can identify the PCI passthrough devices by their vendor and model name.

Note: Some virtual machine operations become unavailable when you add a PCI or PCIe passthrough device to the virtual machine.

For information about the hardware label configuration, see [Change the Hardware Label in the VMware Host Client](#).

1. In the VMware Host Client inventory, click **Virtual Machines**.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, click **Add other device** and select a device.

Option	Action
PCI device	1. Click PCI device. A new device appears in the hardware list. 2. From the drop-down menu, select a PCI device to connect to the virtual machine.
Dynamic PCI device	1. Click Dynamic PCI device. A new device appears in the hardware list.

Option	Action
	2. Expand New PCI device, and from the drop-down menu, select the PCI passthrough devices to connect to the virtual machine. You can identify PCI passthrough devices by vendor, model name, and hardware label. Hardware labels, if present, are displayed in brackets. Note: When you add a PCI device to a virtual machine, the full memory size of the virtual machine is automatically reserved.

4. Click **Save**.

Securing Virtual Machines in the VMware Host Client

The guest operating system that runs in the virtual machine is vulnerable to the same security risks as any physical system.

To boost security in your virtual environment, you can add a virtual Trusted Platform Module (vTPM) to your ESXi hosts. You can also enable virtualization-based security (VBS) for the virtual machines that run the latest Windows 10 and Windows Server 2016 operating systems. You can provide additional security to your workloads by using Virtual Intel® Software Guard Extensions (vSGX) for virtual machines.

Activate vSGX on a Virtual Machine in the VMware Host Client

To protect the enclave contents from disclosure and modifications, you can activate vSGX on a virtual machine in the VMware Host Client.

- Power off the virtual machine.
- Verify that the virtual machine uses EFI firmware.
- Verify that the ESXi host is version 7.0 or later.
- Verify that the guest operating system in the virtual machine is Linux, Windows 10 (64-bit) or later, or Windows Server 2016 (64-bit) or later.
- Verify that you have the **Virtual machine > Configuration > Modify device settings** privilege on the virtual machine.
- Verify that the ESXi host is installed on an SGX-capable CPU, and SGX is enabled in the BIOS of the ESXi host. For information about the supported CPUs, see <https://kb.vmware.com/s/article/71367>.

Secure Virtual Machines with vSGX

vSphere enables you to configure vSGX for virtual machines. Some modern Intel CPUs implement a security extension called Intel® Software Guard Extension (Intel® SGX). Intel SGX allows user-level code to define private regions of memory, called enclaves. Intel SGX protects the enclave contents from disclosure or modification in such a way that code running outside the enclave cannot access them.

vSGX enables virtual machines to use Intel SGX technology if available on the hardware. To use vSGX, the ESXi host must be installed on an SGX-capable CPU and SGX must be enabled in the BIOS of the ESXi host. You can use the vSphere Client to enable SGX for a virtual machine. For more information, see the *vSphere Security* documentation.

Some operations and features are not compatible with SGX.

- Migration with Storage vMotion

- Suspending or resuming the virtual machine
 - Taking a snapshot of the virtual machine
 - Fault Tolerance
 - Enabling Guest Integrity (GI, platform foundation for VMware AppDefense 1.0)
1. In the VMware Host Client inventory, click **Virtual Machines**.
 2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
 3. On the **Virtual Hardware** tab, expand **Security devices**.
 4. Select the **Enable** check box.
 5. Under **Enclave page cache size**, enter a new value in the text box and select the size in MB or GB from the drop-down menu.
- Note:** The enclave page cache size must be a multiple of 2.
6. From the **Launch control configuration** drop-down menu, select the appropriate mode.

Option	Action
Locked	Activates the launch enclave configuration. Under Launch enclave public key hash, enter a valid SHA256 hash. The SHA256 hash key must contain 64 characters.
Unlocked	Activates the launch enclave configuration of the guest operating system.

7. Click **Save**.

Deactivate vSGX on a Virtual Machine in the VMware Host Client

To deactivate vSGX on a virtual machine, you can use the VMware Host Client.

1. In the VMware Host Client inventory, click **Virtual Machines**.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, expand **Security devices**.
4. Deselect the **Enable** check box and click **Save**.

vSGX is deactivated on the virtual machine.

Remove a vTPM device from a VM in the VMware Host Client

The Trusted Platform Module (TPM) is a specialized chip that stores host-specific sensitive information, for example private keys and OS secrets. The TPM chip is also used to perform cryptographic tasks and attest the integrity of the platform. In the VMware Host Client, you can only remove the vTPM device from a virtual machine.

- The virtual machine must be of hardware version 14 or later.
- The guest OS must be Windows 10 or Windows Server 2016 and later.
- The virtual machine must be powered off.

The virtual TPM device is a software emulation of the TPM functionality. You can add a virtual TPM (vTPM) device to the virtual machines in your environment. The vTPM implementation does not require a physical TPM chip on the host. ESXi uses the vTPM device to exert the TPM functionality in your vSphere environment.

vTPM is available to virtual machines that have Windows 10 and Windows Server 2016 operating systems. The virtual machine must be of hardware version 14 or later.

You can add a virtual TPM device to a virtual machine only in the vCenter Server instance. For more information, see the *vSphere Security* documentation.

In the VMware Host Client, you can only remove the virtual TPM device from a virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **Virtual Hardware** tab, find the TPM device and click the **Remove** icon.
The virtual TPM device is removed from the virtual machine.
4. Click **Save** to close the wizard.

Activate or Deactivate Virtualization-based Security on an Existing VM in the VMware Host Client

Virtualization-based security (VBS) uses the Microsoft Hyper-V based virtualization technology to isolate core Windows OS services in a separate virtualized environment. Such isolation provides an additional level of protection, because it makes it impossible for the key services in your environment to be manipulated.

Configuring VBS is a process that involves first activating VBS in the virtual machine then activating VBS in the guest operating system.

Note: New virtual machines configured for Windows 10, Windows Server 2016, and Windows Server 2019 on hardware versions less than version 14 are created using Legacy BIOS by default. If you change the firmware type of a virtual machine from Legacy BIOS to UEFI, you must reinstall the guest operating system.

You can activate VBS on a virtual machine only if the TPM validation of the host is successful.

Using Intel CPUs for VBS requires vSphere 6.7 or later. The virtual machine must have been created using hardware version 14 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit) or later releases
- Windows Server 2016 (64 bit) or later releases

Using AMD CPUs for VBS requires vSphere 7.0 Update 2 or later. The virtual machine must have been created using hardware version 19 or later and one of the following supported guest operating systems:

- Windows 10 (64 bit), version 1809 or later releases
- Windows Server 2019 (64 bit) or later releases

Ensure that you install the latest patches for Windows 10, version 1809, and Windows Server 2019, before enabling VBS.

For more information about activating VBS on virtual machines on AMD platforms, see the VMware KB article at <https://kb.vmware.com/s/article/89880>.

You can change the level of security of a virtual machine by enabling or disabling Microsoft virtualization-based security (VBS) on existing virtual machines for supported Windows guest operating systems.

Activating VBS on a virtual machine automatically activates the virtual hardware that Windows requires for the VBS feature. By enabling VBS, a variant of Hyper-V starts in the virtual machine and Windows starts running inside the Hyper-V root partition.

VBS is available on the latest Windows OS versions, for example Windows 10 and Windows Server 2016. To use VBS on a virtual machine, the virtual machine compatibility must be ESXi 6.7 and later.

In the VMware Host Client, you can activate VBS during a virtual machine creation. Alternatively, you can activate or deactivate VBS for an existing virtual machine.

1. Click **Virtual Machines** in the inventory.
2. Right-click a virtual machine in the list and select **Edit settings** from the pop-up menu.
3. On the **VM Options** tab, activate or deactivate VBS for the virtual machine.
 - To activate VBS for the virtual machine, select the **Enable Virtualization Based Security** check box.
 - To deactivate VBS for the virtual machine, deselect the **Enable Virtualization Based Security** check box.

When you activate VBS, several options are automatically selected and become dimmed in the wizard.

4. Click **Save** to close the wizard.

Managing Storage in the VMware Host Client

When you connect to an ESXi host by using the VMware Host Client, you can perform different storage management tasks on the ESXi host, including configuring adapters, creating datastores, and viewing storage device information.

Datastores in the VMware Host Client

Datastores are logical containers, similar to file systems, that contain specific information of each storage device and provide a uniform model for storing virtual machine files.

You can use datastores to store ISO images, virtual machine templates, and floppy images.

Depending on the type of storage you use, datastores can be of the following types:

- Virtual Machine File System (VMFS)
- Network File System (NFS)

You can increase the datastore capacity after you create a datastore but only if it is a VMFS datastore.

Block storage devices, Fibre Channel and iSCSI, and NAS devices support the hardware acceleration.

The hardware acceleration functionality enables the ESXi host to integrate with compliant storage systems. The host can offload certain virtual machine and storage management operations to the storage systems. With the storage hardware assistance, your host performs these operations faster and consumes less CPU, memory, and storage fabric bandwidth.

For more information, see the VMware knowledge base article at <http://kb.vmware.com/kb/1021976>.

View Datastore Information in the VMware Host Client

Use the VMware Host Client to display the datastores available to the hosts and analyze their properties.

1. Click **Storage** in the VMware Host Client inventory and click **Datastores**.
2. To view the details for a specific datastore, select the datastore from the list.

Create a VMFS Datastore in the VMware Host Client

VMFS datastores serve as repositories for virtual machines. You can set up VMFS datastores on any SCSI-based storage devices that your host detects, including Fibre Channel, iSCSI, and local storage devices.

Install and configure any adapters that your storage requires. Rescan the adapters to discover newly added storage devices.

To create datastores in the VMware Host Client, you can use the **New datastore** wizard.

1. Click **Storage** in the inventory and click **Datastores**.

2. Click **New datastore**.

The **New datastore** wizard opens.

3. On the Select creation type page, select **Create new VMFS datastore** and click **Next**.

Option	Description
Create new VMFS datastore	Creates a new VMFS datastore on a local disk device.
Add an extent to existing VMFS datastore	Increases the size of an existing datastore by adding a new extent on another disk.
Expand an existing VMFS datastore extent	Increases the size of an existing datastore extent.
Mount NFS datastore	Creates a new datastore by mounting a remote NFS volume.

4. On the Select device page, select where to create the new VMFS partition.

- a) Enter a name for the new datastore.
- b) Select a device to add the datastore to.

The list contains only devices that have enough available space.

- c) Click **Next**.

5. On the Select partitioning options page, select how to partition the device and click **Next**.

Option	Description
Use Full Disk	It shows you all the free space that is available on the device.
Custom	Click the Free space bar and use the horizontal scroller to partition the device.

6. On the Ready to complete page, review the configuration details and click **Finish**.

Increasing VMFS Datastore Capacity

If your VMFS datastore requires more space, increase the datastore capacity. You can dynamically increase the capacity by growing a datastore extent or by adding an extent.

Use one of the following methods to increase the datastore capacity:

- Dynamically grow any expandable datastore extent, so that it fills the available adjacent capacity. The extent is considered expandable when the underlying storage device has free space immediately after the extent.
- Dynamically add the extent. The datastore can span over up to 32 extents with the size of each extent of more than 2 TB, yet appear as a single volume. The spanned VMFS datastore can use any or all its extents at any time. It does not need to fill up a particular extent before using the next one.

Increase an Existing VMFS Datastore in the VMware Host Client

When you need to add virtual machines to a datastore, or when the virtual machines running on a datastore require more space, you can dynamically increase the capacity of a VMFS datastore.

If a shared datastore has powered on virtual machines and becomes 100% full, you can increase the datastore's capacity only from the host that the powered on virtual machines are registered on.

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **New datastore**.
3. On the Select creation type page, click **Add an extent to existing VMFS datastore** and click **Next**.
4. On the Select datastore page, select the datastore to expand and click **Next**.
5. On the Select device page, select a device to create the new VMFS partition on and click **Next**.
6. On the Select partitioning options page, select how to partition the device and click **Next**.

Option	Description
Use Full Disk	It shows you all the free space that is available on the device.
Custom	Click the Free space bar and use the horizontal scroller to partition the device.

7. On the Ready to complete page, review the configuration details and click **Finish**.

Mounting a Network File System Datastore in the VMware Host Client

With the VMware Host Client, you can create a Network File System (NFS) datastore to store virtual disks and to use as a central repository for ISO images, virtual machines, and so on.

An NFS client built into ESXi uses the Network File System (NFS) protocol over TCP/IP to access a designated NFS volume that is located on a NAS server. vSphere supports versions 3 and 4.1 of the NFS protocol.

The ESXi host can mount an NFS volume and use it for its storage needs.

Typically, the NFS volume or directory is created by a storage administrator and is exported from the NFS server. You do not need to format the NFS volume with a local file system, such as VMFS. Instead, you mount the volume directly on the ESXi hosts and use it to store and boot virtual machines in the same way that you use the VMFS datastores.

In addition to storing virtual disks on NFS datastores, you can use NFS as a central repository for ISO images, virtual machine templates, and so on. If you use the datastore for the ISO images, you can connect the CD-ROM device of the virtual machine to an ISO file on the datastore. You then can install a guest operating system from the ISO file.

When you use NFS storage, follow the specific guidelines related to NFS server configuration, networking, NFS datastores, and so on.

Mount an NFS Datastore in the VMware Host Client

Use the **New datastore** wizard to mount a Network File System (NFS) datastore in the VMware Host Client.

Because NFS requires network connectivity to access data on remote servers, before configuring NFS, you must first configure VMkernel networking.

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **New datastore**.

The **New datastore** wizard opens.

3. On the Select creation type page, click **Mount NFS datastore** and click **Next**.
4. On the Provide NFS mount details page, provide the details for the NFS you mount.

- a) Enter a name for the NFS datastore.
- b) Enter the NFS server name.

For the server name, you can enter an IP address, a DNS name, or an NFS UUID.

Note: When you mount the same NFS volume on different hosts, make sure that the server and folder names are identical across the hosts. If the names do not match, the hosts detect the same NFS volume as two different datastores. This might result in a failure of features such as vMotion. An example of such discrepancy is if you enter `filer` as the server name on one host and `filer.domain.com` on the other.

- c) Specify the NFS share.
- d) Specify the NFS version.
- e) Click **Next**.

5. On the Ready to complete page, review the configuration settings for the NFS datastore and click **Finish**.

Unmount a Datastore in the VMware Host Client

When you unmount a datastore in the VMware Host Client, it remains intact, but you can no longer view it in the inventory of the host that you manage. The datastore continues to appear on other hosts that it remains mounted on.

Note: Make sure that the datastore is not used by vSphere HA heartbeating. vSphere HA heartbeating does not prevent you from unmounting the datastore. However, if the datastore is used for heartbeating, unmounting it might cause the host to fail and restart all active virtual machines.

Before unmounting a datastore, also make sure that the following prerequisites are met:

- No virtual machines reside on the datastore.
- Storage DRS does not manage the datastore.
- Storage I/O Control is deactivated for this datastore.

Do not perform any configuration operations that might result in I/O to the datastore while the unmounting is in progress.

1. Click **Storage** in the inventory and click **Datastores**.
2. Right-click the datastore to unmount from the list and click **Unmount**.
3. Confirm that you want to unmount the datastore.

Unmounting or Removing Datastore Fails

When you try to unmount or remove a datastore, the operation fails.

The operation to unmount or remove a datastore fails if the datastore has any opened files. For these user operations, the vSphere HA agent closes all the files that it has opened, for example, heartbeat files. If the agent is not reachable by vCenter Server or the agent cannot flush out pending I/Os to close the files, a The HA agent on host '{hostName}' failed to quiesce file activity on datastore '{dsName}' fault is triggered.

If the datastore to be unmounted or removed is used for heartbeating, vCenter Server excludes it from heartbeating and chooses a new one. However, the agent does not receive the updated heartbeat datastores if it is not reachable, that is, if the host is isolated or in a network partition. In such cases, heartbeat files are not closed and the user operation fails. The operation can also fail if the datastore is not accessible because of storage failures such as all paths down.

Note: When you remove a VMFS datastore, the datastore is removed from all the hosts in inventory. So if there are any hosts in a vSphere HA cluster that are unreachable or that cannot access the datastore, the operation fails.

Ensure that the datastore is accessible and the affected hosts are reachable.

Using Datastore File Browser in the VMware Host Client

Use the datastore file browser to manage the contents of your datastore. You can perform a number of tasks that include uploading files to the datastore, downloading datastore files to your system, moving and copying datastore folders or files, and creating new datastore directories.

Upload Files to a Datastore in the VMware Host Client

Use the datastore file browser to upload files to datastores on your host.

Required privilege: **Datastore > Browse Datastore**

Note: Virtual Volumes do not support uploading files directly to the virtual datastores. You must first create a folder on the virtual datastore, and then upload the files into the folder.

In addition to their traditional use as a storage for virtual machine files, datastores can serve to store data or files related to virtual machines. For example, you can upload ISO images of operating systems from a local computer to a datastore on the host. You then use these images to install guest operating systems on the new virtual machines.

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **Datastore browser**.
3. Select the datastore that you want to store the file on.
4. Optional: Click **Create directory** to create a new datastore directory to store the file.
5. Select the target folder and click **Upload**.
6. Locate the item that you want to upload from your local computer and click **Open**.
The file uploads to the datastore that you selected.
7. Optional: Refresh the datastore file browser to see the uploaded file on the list.
8. Click **Close** to exit the file browser.

Download Files from a Datastore to Your System in the VMware Host Client

Use the datastore file browser to download files from the datastores available on the host that you are managing to your local system.

Required privilege: **Datastore > Browse Datastore**

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **Datastore browser**.
3. Select the target datastore.
4. Click the folder that contains the file that you want to download.
The available files in the folder are displayed.

5. Click the file that you want to download.
6. Click **Download**.
The file is downloaded to your system.
7. Click **Close** to exit the file browser.

Delete Files from a Datastore in the VMware Host Client

You can permanently remove files from any datastore if you no longer need them.

Required privilege: **Datastore > Browse Datastore**

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **Datastore browser**.
3. Select the target datastore.
4. Select the folder that contains the file that you want to delete.
The available files in the folder are displayed.
5. Click the file that you want to remove from the datastore, click **Delete**, and click **Delete** again.
6. Click **Close** to exit the file browser.

Move Datastore Folders or Files in the VMware Host Client

Use the datastore file browser to move files or folders to a new location, either on the same datastore or on a different datastore.

Required privilege: **Datastore > Browse Datastore**

Note: Virtual disk files are moved and copied without format conversion. If you move a virtual disk to a datastore on a type of host that is different from the type of the source host, you might need to convert the virtual disks before you can use them.

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **Datastore browser**.
3. Select the target datastore.
4. Select the file or folder that you want to move to another location and click **Move**.
5. Select your target destination and click **Move**.
6. Click **Close** to exit the file browser.

Copy Datastore Folders or Files in the VMware Host Client

Use the datastore file browser to copy folders or files to a new location, either on the same datastore or on a different datastore.

Required privilege: **Datastore > Browse Datastore**

Note: Virtual disk files are moved and copied without format conversion. If you move a virtual disk to a datastore on a type of host that is different from the type of the source host, you might need to convert the virtual disks.

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **Datastore browser**.
3. Select the target datastore.
4. Select the file or folder that you want to move to another location and click **Copy**.
5. Select your target destination and click **Copy**.
6. Click **Close** to exit the file browser.

Create a New Datastore Directory in the VMware Host Client

You can create new datastore directories if you want to store files in a particular location.

Required privilege: **Datastore > Browse Datastore**

1. Click **Storage** in the inventory and click **Datastores**.
2. Click **Datastore browser**.
3. Click **Create directory**.
4. Select the target datastore.
5. Optional: Enter a name for the new directory.
6. Click **Create directory**.
7. Click **Close** to exit the file browser.

Rename a Datastore in the VMware Host Client

You can change the display name of a datastore in the VMware Host Client.

Note: If the host is managed by vCenter Server, you cannot rename the datastore from the VMware Host Client. You can only perform the task from the vCenter Server instance that manages the host.

1. Click **Storage** in the VMware Host Client inventory and click **Datastores**.
2. Right-click a datastore in the list and select **Rename** from the drop-down menu.
3. Enter a new name for the datastore and click **Save** to apply your changes.
4. Optional: Click **Refresh** to see the new name of the datastore in the list of available datastores.

Delete a VMFS Datastore in the VMware Host Client

You can delete any type of VMFS datastore, including copies that you have mounted without resignaturing. When you delete a datastore, the datastore and all the files associated to the datastore are removed from the host.

Remove all virtual machines from the datastore.

Note: The datastore delete operation permanently deletes all files associated with virtual machines on the datastore. Although you can delete the datastore without unmounting, it is preferable that you unmount the datastore first.

1. Click **Storage** in the inventory and click **Datastores**.
2. Right-click the datastore from the list and select **Delete** from the drop-down menu.
3. Click **Confirm** to delete the datastore.

Storage Thin Provisioning in the VMware Host Client

With ESXi, you can use two models of thin provisioning, array-level and virtual-disk level thin provisioning.

Thin provisioning is a method that optimizes storage utilization by allocating storage space in a flexible on-demand manner. Thin provisioning contrasts with the traditional model, called thick provisioning. With thick provisioning, a large amount of storage space is provided in advance in anticipation of future storage needs. However, the space might remain unused causing underutilization of storage capacity.

The VMware thin provisioning features help you eliminate storage underutilization problems at the datastore and storage array level.

Create Thin Provisioned Virtual Disks in the VMware Host Client

To save storage space, you can create thin provisioned virtual disks. The thin provisioned virtual disk starts small and grows as more disk space is required. You can create thin disks only on the datastores that support disk-level thin provisioning.

The following procedure assumes that you are creating a new virtual machine. For more information, see [Create a Virtual Machine in the VMware Host Client](#).

1. Right-click **Host** in the VMware Host Client inventory and select **Create/Register VM**.
The **New Virtual Machine** wizard opens.
2. Select a method for adding a new virtual machine on the host and click **Next**.
3. Enter a name for your virtual machine.
4. Select the virtual machine compatibility from the **Compatibility** drop-down menu.
5. Select a guest operating system version from the **Guest OS version** drop-down menu and click **Next**.
6. From the list of accessible datastores on the Select storage page of the **New Virtual Machine** wizard, select the destination datastore for the virtual machine configuration files and all of the virtual disks.
7. On the **Virtual Hardware** tab, expand **Hard disk**.
8. Under **Disk Provisioning**, select the **Thin provisioned** radio button and click **Next**.
9. On the Ready to complete page of the **New Virtual Machine** wizard, review the configuration settings for the virtual machine and click **Finish** to save your settings.

View Virtual Machine Storage Resources in the VMware Host Client

You can view how datastore storage space is allocated for your virtual machines in the VMware Host Client.

Resource Consumption shows how much datastore space is occupied by virtual machine files, including configuration files, log files, snapshots, virtual disks, and so on. When the virtual machine is running, the used storage space also includes swap files.

For virtual machines with thin disks, the actual storage usage value might be less than the size of the virtual disk.

1. Click the virtual machine in the VMware Host Client inventory.
2. Review the Resource Consumption information in the lower right area of the virtual machine summary page.

Determine the Disk Format of a Virtual Machine in VMware Host Client

You can determine whether your virtual disk is provisioned in thick or thin format.

1. Right-click the virtual machine in the VMware Host Client inventory and select **Edit settings**.
2. On the **Virtual Hardware** tab, expand **Hard disk**.

The **Type** text box shows the format of your virtual disk.

Managing Storage Adapters in the VMware Host Client

When you connect to a host or to vCenter Server by using the VMware Host Client, you can perform various tasks on your storage adapters, such as configuring various iSCSI components.

When you enable iSCSI on the host that you are managing in your VMware Host Client environment, you can configure and add new network port bindings, static and dynamic targets, you can manage CHAP authentication, and configure various advanced settings on your host storage.

View Storage Adapters in the VMware Host Client

View the storage adapters that your host uses and related information.

1. Click **Storage** in the VMware Host Client inventory and click **Adapters**.

All storage adapters available to the host are listed under **Adapters**.

2. To view details for a specific adapter, select the adapter from the list.

Configuring Software iSCSI Adapters in the VMware Host Client

With the software-based iSCSI implementation, you can use standard NICs to connect your host to a remote iSCSI target on the IP network. The software iSCSI adapter that is built into ESXi communicates with the physical NICs through the network stack.

Note: Before you can use the software iSCSI adapter, you must set up networking, activate the adapter, and configure parameters such as CHAP.


The iSCSI adapter configuration workflow includes the following procedures:

- Enabling iSCSI on your host. See [Enable iSCSI for an ESXi Host in the VMware Host Client](#).
- Adding a port binding. See [Add Port Binding in the VMware Host Client](#).
- Removing port binding. See [Remove Port Binding in the VMware Host Client](#).

Setting Up Network for iSCSI and iSER with ESXi

Certain types of iSCSI adapters depend on the VMkernel networking. These adapters include the software or dependent hardware iSCSI adapters, and the VMware iSCSI over RDMA (iSER) adapter. If your ESXi environment includes any of these adapters, you must configure connections for the traffic between the iSCSI or iSER component and the physical network adapters.

Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You use 1:1 mapping between each virtual and physical network adapter. You then associate the VMkernel adapter with an appropriate iSCSI or iSER adapter. This process is called port binding.



Follow these rules when configuring the port binding:

- You can connect the software iSCSI adapter with any physical NICs available on your host.
- The dependent iSCSI adapters must be connected only to their own physical NICs.
- You must connect the iSER adapter only to the RDMA-capable network adapter.

For specific considerations on when and how to use network connections with software iSCSI, see the VMware knowledge base article at <http://kb.vmware.com/kb/2038869>.

Enable iSCSI for an ESXi Host in the VMware Host Client

Enable iSCSI for your host in your VMware Host Client environment to configure storage adapters parameters, such as CHAP authentication, network port bindings, static and dynamic targets, and various advanced settings.

1. Click **Storage** in the inventory, click **Adapters**, and click **Configure iSCSI**.
2. Select the **Enabled** radio button.
3. Optional: Configure the parameters and components that you want to change.
4. Click **Save configuration**.

Best Practices for Configuring Networking with Software iSCSI

When you configure networking with software iSCSI, consider several best practices.

Software iSCSI Port Binding

You can bind the software iSCSI initiator on the ESXi host to a single or multiple VMkernel ports, so that iSCSI traffic flows only through the bound ports. Unbound ports are not used for iSCSI traffic.

When port binding is configured, the iSCSI initiator creates iSCSI sessions from all bound ports to all configured target portals.

See the following examples.

VMkernel Ports	Target Portals	iSCSI Sessions
2 bound VMkernel ports	2 target portals	4 sessions (2 x 2)
4 bound VMkernel ports	1 target portal	4 sessions (4 x 1)
2 bound VMkernel ports	4 target portals	8 sessions (2 x 4)

Note: Make sure that all target portals are reachable from all VMkernel ports when port binding is used. Otherwise, iSCSI sessions might fail to create. As a result, the rescan operation might take longer than expected.

No Port Binding

If you do not use port binding, the ESXi networking layer selects the best VMkernel port based on its routing table. The host uses the port to create an iSCSI session with the target portal. Without the port binding, only one session per each target portal is created.


See the following examples.

VMkernel Ports	Target Portals	iSCSI Sessions
2 unbound VMkernel ports	2 target portals	2 sessions
4 unbound VMkernel ports	1 target portal	1 session
2 unbound VMkernel ports	4 target portals	4 sessions

Software iSCSI Multipathing

Example 1. Multiple paths for an iSCSI target with a single network portal

If your target has only one network portal, you can create multiple paths to the target by adding multiple VMkernel ports on your ESXi host and binding them to the iSCSI initiator.




In this example, all initiator ports and the target portal are configured in the same subnet. The target is reachable through all bound ports. You have four VMkernel ports and one target portal, so total of four paths are created.

Without the port binding, only one path is created.

Example 2. Multiple paths with VMkernel ports in different subnets

You can create multiple paths by configuring multiple ports and target portals on different IP subnets. By keeping initiator and target ports in different subnets, you can force ESXi to create paths through specific ports. In this configuration, you do not use port binding because port binding requires that all initiator and target ports are on the same subnet.



ESXi selects vmk1 when connecting to Port 0 of Controller A and Controller B because all three ports are on the same subnet. Similarly, vmk2 is selected when connecting to Port 1 of Controller A and B. You can use NIC teaming in this configuration.

Total of four paths are created.


Paths	Description
Path 1	vmk1 and Port0 of Controller A
Path 2	vmk1 and Port0 of Controller B
Path 3	vmk2 and Port1 of Controller A
Path 4	vmk2 and Port1 of Controller B

Routing with Software iSCSI

You can use the `esxcli` command to add static routes for your iSCSI traffic. After you configure static routes, initiator and target ports in different subnets can communicate with each other.

Example 1. Using static routes with port binding

In this example, you keep all bound VMkernel ports in one subnet (N1) and configure all target portals in another subnet (N2). You can then add a static route for the target subnet (N2).




Use the following command:

```
# esxcli network ip route ipv4 add -gateway 192.168.1.253 -network 10.115.179.0/24
```

Example 2. Using static routes to create multiple paths

In this configuration, you use static routing when using different subnets. You cannot use the port binding with this configuration.



You configure vmk1 and vmk2 in separate subnets, 192.168.1.0 and 192.168.2.0. Your target portals are also in separate subnets, 10.115.155.0 and 10.115.179.0.

You can add the static route for 10.115.155.0 from vmk1. Make sure that the gateway is reachable from vmk1.

```
# esxcli network ip route ipv4 add -gateway 192.168.1.253 -network 10.115.155.0/24
```

You then add static route for 10.115.179.0 from vmk2. Make sure that the gateway is reachable from vmk2.

```
# esxcli network ip route ipv4 add -gateway 192.168.2.253 -network 10.115.179.0/24
```

When connecting with Port 0 of Controller A, vmk1 is used.

When connecting with Port 0 of Controller B, vmk2 is used.

Example 3. Routing with a separate gateway per vmkernel port

Starting with vSphere 6.5, you can configure a separate gateway per VMkernel port. If you use DHCP to obtain IP configuration for a VMkernel port, gateway information can also be obtained using DHCP.

To see gateway information per VMkernel port, use the following command:

```
# esxcli network ip interface ipv4 address list
```

Name	IPv4 Address	IPv4 Netmask	IPv4 Broadcast	Address Type	Gateway	DHCP	DNS
vmk0	10.115.155.122	255.255.252.0	10.115.155.255	DHCP	10.115.155.253	true	
vmk1	10.115.179.209	255.255.252.0	10.115.179.255	DHCP	10.115.179.253	true	
vmk2	10.115.179.146	255.255.252.0	10.115.179.255	DHCP	10.115.179.253	true	

With separate gateways per VMkernel port, you use port binding to reach targets in different subnets.

Add Port Binding in the VMware Host Client

Use the VMware Host Client to bind an iSCSI adapter with a VMkernel adapter on your host.

- Create a virtual VMkernel adapter for each physical network adapter on your host. If you use multiple VMkernel adapters, set up the correct network policy.
- Required privilege: **Host > Configuration > Storage Partition Configuration**

1. Click **Storage** in the inventory, click **Adapters**, and click **Configure iSCSI**.
2. In the **Network port bindings** section, click **Add port binding**.
3. Select a VMkernel adapter to bind with the iSCSI adapter.

Note: Make sure that the network policy for the VMkernel adapter is compliant with the binding requirements.

You can bind the software iSCSI adapter to one or more VMkernel adapters. For a dependent hardware iSCSI adapter, only one VMkernel adapter associated with the correct physical NIC is available.

4. Click **Select**.
5. Click **Save configuration**.

Remove Port Binding in the VMware Host Client

Edit the iSCSI configuration on your host to remove a port binding.

1. Click **Storage** in the inventory, click **Adapters**, and click **Configure iSCSI**.
2. In the **Network port bindings** section, select a VMkernel NIC from the list.
3. Click **Remove port binding**.
4. Click **Save configuration**.

Set up a Dynamic Target in the VMware Host Client

You must set up target discovery addresses, so that the iSCSI adapter can determine which storage resource on the network is available for access. The ESXi host supports dynamic and static discovery methods. With Dynamic Discovery, each time the initiator contacts a particular iSCSI storage system, the initiator sends the SendTargets request to the iSCSI system. The iSCSI system responds by supplying a list of available targets to the initiator.

Required privilege: **Host > Configuration > Storage Partition Configuration**

Also known as SendTargets discovery. Each time the initiator contacts a specified iSCSI server, the initiator sends the SendTargets request to the server. The server responds by supplying a list of available targets to the initiator. The names and IP addresses of these targets appear on the **Static Discovery** tab. If you remove a static target added by dynamic discovery, the target might be returned to the list the next time a rescan happens, the iSCSI adapter is reset, or the host is rebooted.

Note: With software and dependent hardware iSCSI, ESXi filters target addresses based on the IP family of the iSCSI server address specified. If the address is IPv4, IPv6 addresses that might come in the SendTargets response from the iSCSI server are filtered out. When DNS names are used to specify an iSCSI server, or when the SendTargets response from the iSCSI server has DNS names, ESXi relies on the IP family of the first resolved entry from DNS lookup.

When you set up Dynamic Discovery, you can only add a new iSCSI system. You cannot change the IP address, DNS name, or port number of an existing iSCSI system. To modify the parameters, delete the existing system and add a new one.

1. Click **Storage** in the inventory, click **Adapters**, and click **Configure iSCSI**.
 2. Click **Add dynamic target**.
- The new dynamic target appears in the list.

3. To add an address for the new dynamic target, click the target in the list and enter the address .
4. Optional: To change the port number of the new dynamic target, click the target **Port** text box and enter the new port number.
5. Optional: To edit the dynamic target settings, select the new target from the list of available targets, click **Edit settings**, configure the parameters that you want to change, and click **Save**.
6. Optional: To delete a specific target, select the target and click **Remove dynamic target**.
The target no longer appears in the list of existing dynamic targets.
7. Click **Save configuration**.

Set Up a Static Target in the VMware Host Client

With iSCSI initiators, you can use static discovery to manually enter information for the targets.

Required privileges: **Host > Configuration > Storage Partition Configuration**

When you set up Static Discovery, you can only add new iSCSI targets. You cannot change the IP address, DNS name, iSCSI target name, or port number of an existing target. To make changes, remove the existing target and add a new one.

In addition to the dynamic discovery method, you can use static discovery and manually enter information for the targets. The iSCSI adapter uses a list of targets that you provide to contact and communicate with the iSCSI servers.

1. Click **Storage** in the inventory, click **Adapters**, and click **Configure iSCSI**.
2. Click **Add static target**.
The new static target appears in the list.
3. To add a name for the new static target, click the target in the list and enter the name.
4. To add an address for the new static target, click the target in the list and type the address.
5. Optional: To change the port number of the new static target, click the target **Port** text box and type the new port number.
6. Optional: To edit the static target settings, select the new target from the list of available targets, click **Edit settings**, configure the parameters that you want to change, and click **Save**.
7. Optional: To delete a specific target, select the target and click **Remove static target**.
The target no longer appears in the list of existing static targets.
8. Click **Save configuration**.

Edit Advanced Settings for iSCSI in the VMware Host Client

The advanced iSCSI settings control such parameters as header and data digest, ARP redirection, delayed ACK, and so on. Generally, you do not need to change these settings because your host works with the assigned predefined values.

Required privilege: **Host > Configuration > Storage Partition Configuration**

CAUTION: Do not change the advanced iSCSI settings unless you are working with the VMware support team or otherwise have thorough information about the values to provide for the settings modification.

1. Click **Storage** in the inventory, click **Adapters**, and click **Configure iSCSI**.
2. Click **Advanced settings** to display the entire list of settings.
3. Edit the parameters that you want to change and click **Save configuration**.

Set Up CHAP Authentication for an iSCSI Adapter in the VMware Host Client

You can set up all targets to receive the same CHAP name and secret from the iSCSI initiator at the initiator level. By default, all discovery addresses or static targets inherit the CHAP parameters that you set up at the initiator level.

- Before you set up CHAP parameters for software or dependent hardware iSCSI, determine whether to configure one-way, also known as normal, or mutual CHAP. Independent hardware iSCSI adapters do not support mutual CHAP.
 - In one-way CHAP, the target authenticates the initiator.
 - In mutual CHAP, both the target and the initiator authenticate each other. Use different secrets for CHAP and mutual CHAP.

When you configure CHAP parameters, verify that they match the parameters on the storage side.

- Required privileges: **Host > Configuration > Storage Partition Configuration**

The CHAP name must be fewer than 511 alphanumeric characters and the CHAP secret must be fewer than 255 alphanumeric characters. Some adapters, for example the QLogic adapter, might have lower limits, 255 for the CHAP name and 100 for the CHAP secret.

1. Click **Storage** in the inventory, click **Adapters**, and click **Configure iSCSI**.
2. To configure one-way CHAP, expand **CHAP authentication** to display all parameters.
 - a) Select the CHAP security level.
 - b) Enter the CHAP name.

Make sure that the name you enter matches the name configured on the storage side.
 - c) Enter a one-way CHAP secret to use for authentication. Use the same secret that you enter on the storage side.
3. To configure mutual CHAP, select **Use CHAP** as an option for one-way CHAP. Expand **Mutual CHAP authentication** to display all parameters.
 - a) Select **Use CHAP**.
 - b) Enter the mutual CHAP name.
 - c) Enter the mutual CHAP secret.

Use different secrets for the one-way CHAP and the mutual CHAP.

4. Click **Save configuration**.

If you change the authentication settings for an iSCSI adapter, you only use the updated credentials for new iSCSI sessions. Existing sessions persist until either the connection is lost due to some outside factor, such as force re-authentication, or you remove and add the adapter iSCSI targets.

Managing Storage Devices in the VMware Host Client

You can use the VMware Host Client to manage local and networked storage devices that the ESXi host you are managing has access to.

View Storage Devices in the VMware Host Client

View all storage devices available to a host. If you use third-party multipathing plug-ins, the storage devices available through the plug-ins also appear on the list.

The Storage Devices view allows you to list the host storage devices, analyze their information, and modify properties.

1. Click **Storage** in the VMware Host Client inventory and click **Devices**.

All storage devices available to the host are listed under **Devices**.

2. To view details for a specific device, select the device from the list.

Clear a Device Partition Table in the VMware Host Client

When you are logged in to an ESXi host with the VMware Host Client, you can clear the partition table of a disk device that is accessible from the host.

Verify that the device is not in use by ESXi as boot disk, VMFS datastore, or vSAN.

1. Click **Storage** in the VMware Host Client and click **Devices**.
2. Right-click a device from the list, click **Clear partition table** and click **Yes**.

Clearing the partition table might cause data loss.

Edit Individual Device Partitions in the VMware Host Client

When you log in to an ESXi host with the VMware Host Client, you can remove individual partitions of a device by using the partition editor.

Verify that the device is not in use by ESXi as boot disk, VMFS datastore, or vSAN.

1. Click **Storage** in the VMware Host Client and click **Devices**.
2. Right-click a device from the list and click **Edit partitions**.
3. Select a partition and click **Delete partition**.
4. Optional: Click **Reset** to restore the original partitions.
5. Click **Save partitions**.
6. Confirm that you want to alter the partition.

Manage Persistent Memory

ESXi 6.7 and later provides support for the latest computer memory technology, which is called non-volatile memory (NVM) or persistent memory (PMem).

PMem combines the high data transfer rate of volatile computer memory with the persistence and resiliency of traditional storage. PMem devices have low access latency and can retain stored data through reboots or power outages.

Modes of Consumption of the Persistent Memory Resources of the Host

When you add a physical PMem device to a host, ESXi detects the PMem resource and exposes it as a host-local PMem datastore to the virtual machines that run on the host. Depending on the guest operating system, virtual machines can have direct access to the PMem resources.

Each host can have only one local PMem datastore that pools and represents all PMem resources of the host.

Persistent memory combines the properties of both memory and storage. So, virtual machines can consume the PMem resources of the ESXi host as memory (through virtual NVDIMM devices) or as storage (through virtual PMem hard disks).

The host-local PMem datastore stores all direct-accessed NVDIMM devices and virtual PMem hard disks.

Virtual PMem (vPMem)

In this mode, if the guest operating system is PMem-aware, the virtual machine can have direct access to the physical PMem resources of the host and use them as standard, byte-addressable memory.

Virtual machines use virtual non-volatile dual in-line memory modules (NVDIMMs) for direct access to PMem. The NVDIMM is a memory device that sits on an ordinary memory channel, but contains non-volatile memory. In vSphere 7.0, the virtual NVDIMM is a new type of device that represents the physical PMem regions of the host. A single virtual machine can have up to 64 virtual NVDIMM devices. Each NVDIMM device is stored on the host-local PMem datastore.

Note: To add an NVDIMM device to a virtual machine, the virtual machine must be of hardware version 14 and the guest operating system must support persistent memory. If the guest operating system is not PMem-aware, you can still use PMem, but you cannot add an NVDIMM device to the virtual machine.

Virtual PMem Disks (vPMemDisk)


In this mode, the virtual machine does not have direct access to the PMem resources of the host.

You must add a virtual PMem hard disk to the virtual machine. A virtual PMem hard disk is a traditional SCSI disk to which the PMem Storage Policy is applied. The policy automatically places the hard disk on the host-local PMem datastore.

In this mode of usage, there are no requirements for the hardware version of the virtual machine and the guest operating system.

Note: If the guest operating system is not PMem-aware, virtual machines can use PMem only through vPMemDisks.

The following diagram illustrates how the persistent memory components interact.



For information about how to configure and manage VMs with NVDIMMs or virtual persistent memory disks, see the *vSphere Resource Management* documentation.

Structure of the PMem Datastore

The VMware Host Client user interface provides information about the complex structure of the host-local PMem datastore. To analyze this information and use it for troubleshooting and management purposes, you must be familiar with the concepts related to that complex structure.

Modules

In the VMware Host Client user interface, modules represent the physical NVDIMMs that are connected to the motherboard of the host.

In the VMware Host Client, you can check the health status of each module and identify unhealthy NVDIMM modules.

Interleave Sets

Interleave sets are logical groupings of one or multiple modules. Interleave sets reveal how information is spread across the physical DIMMs and how ESXi reads the information from the modules. Because ESXi reads from each interleave set in turns, interleave sets ensure higher memory throughput.

For example, if an interleave set consists of two modules, ESXi reads the information from the two physical DIMMs in parallel and then proceeds with the next interleave set.

Namespaces

The VMware Host Client user interface provides information about the way NVDIMMs are grouped into interleave sets.

Namespaces are regions of contiguously addressed memory ranges in the NVDIMM. Namespaces can go across interleave sets. The PMem datastore is built on top of the namespaces. In the VMware Host Client, you can view the capacity, health status, and the location ID of every namespace.

View Information About Modules, Interleave Sets, and Namespaces in the VMware Host Client

In the VMware Host Client you can view information about the modules, interleave sets, and namespace of the host-local PMem datastore. As a result, you can easily identify an unhealthy module and perform troubleshooting.

Verify that the host has at least one physical NVDIMM device.

You cannot perform most of the traditional datastore management tasks on the host-local PMem datastore. However, you can use the information about modules, interleave sets, and namespaces for troubleshooting purposes.

1. In the **Navigator** pane, click **Storage**.
2. On the **Persistent Memory** tab, view information about the host-local PMem datastore.
 - Click **Modules** to view information about the NVDIMMs that make up the PMem datastore.
 - Click **Namespaces** to view information about namespaces on the NVDIMMs.
 - Click **Interleave sets** to see how the modules, or physical NVDIMMs, are grouped into interleave sets.

Delete a Namespace in the VMware Host Client

In the VMware Host Client, you can delete namespaces that were not created by ESXi, but by an OS that was previously installed on the host machine.

- Put the host in maintenance mode.
 - Back up the content of the namespace if you might need that content at a later time.
1. In the VMware Host Client, click **Storage**.
 2. Under the **Persistent Memory** tab, click **Namespaces**.
 3. Optional: In the list of namespaces, check the State column to determine which namespaces ESXi currently uses. To free up space, you must delete namespaces whose state is In Use.
 4. Select a namespace and click the **Delete** icon.

Important: Deleting a namespace frees up space on the datastore, but you can use the free space only after you reboot the host.
 5. Click the **Reboot host** icon to reboot the host.

The selected namespace is deleted from the PMem datastore. ESXi automatically creates a new namespace that the PMem datastore can use. The new namespace has the same capacity, type, and location ID as the deleted one.

Monitoring Storage in the VMware Host Client

In the VMware Host Client, you can monitor the storage health of the ESXi host. You can also monitor events and tasks associated with different datastores, storage adapters, and storage devices on the host that you are managing.

Monitor Datastores in the VMware Host Client

In the VMware Host Client, you can monitor the health of a datastore, and events and tasks associated with that datastore.

Starting with vSphere 6.5 Update 1 and later you enable the vSAN service in the vSphere Client, you can also monitor your vSAN environment.

1. Click **Storage** in the inventory.
2. Click **Datastores**.
3. Click a datastore from the list.
The datastore expands in the VMware Host Client inventory.
4. Click **Monitor** under the name of the datastore.
5. Optional: Click **Events** to view events associated with the datastore.
6. Optional: Click **vSAN** to view the configuration parameters of the vSAN environment of your host.
7. Optional: Click **Hosts** to view the hosts that reside on this datastore.
8. Optional: Click **Health** to view details about the status of various parameters, such as **Performance service**, **Network**, **Physical disk**, **Data**, **Cluster**, and **Limits**.

Monitoring vSAN in the VMware Host Client

You can use the VMware Host Client to monitor the vSAN environment of your ESXi host.

vSAN Concepts

VMware vSAN uses a software-defined approach that creates shared storage for virtual machines.

It virtualizes the local physical storage resources of ESXi hosts and turns them into pools of storage that can be divided and assigned to virtual machines and applications according to their quality-of-service requirements. vSAN is implemented directly in the ESXi hypervisor.

You can configure vSAN to work as either a hybrid or all-flash cluster. In hybrid clusters, flash devices are used for the cache layer and magnetic disks are used for the storage capacity layer. In all-flash clusters, flash devices are used for both cache and capacity.

You can activate vSAN on existing host clusters, or when you create a new cluster. vSAN aggregates all local capacity devices into a single datastore shared by all hosts in the vSAN cluster. You can expand the datastore by adding capacity devices or hosts with capacity devices to the cluster. vSAN works best when all ESXi hosts in the cluster share similar or identical configurations across all cluster members, including similar or identical storage configurations. This consistent configuration balances virtual machine storage components across all devices and hosts in the cluster. Hosts without any local devices also can participate and run their virtual machines on the vSAN datastore.

In vSAN Original Storage Architecture (OSA), each host that contributes storage devices to the vSAN datastore must provide at least one device for flash cache and at least one device for capacity. The devices on the contributing host form one or more disk groups. Each disk group contains one flash cache device, and one or multiple capacity devices for persistent storage. Each host can be configured to use multiple disk groups.

In vSAN Express Storage Architecture (ESA), all storage devices claimed by vSAN contribute to capacity and performance. Each host's storage devices claimed by vSAN form a storage pool. The storage pool represents the amount of caching and capacity provided by the host to the vSAN datastore.

For best practices, capacity considerations, and general recommendations about designing and sizing a vSAN cluster, see the *VMware vSAN Design and Sizing Guide*.

Characteristics of vSAN

The following characteristics apply to vSAN, its clusters, and datastores.

vSAN includes numerous features to add resiliency and efficiency to your data computing and storage environment.

Table 567: vSAN Features

Supported Features	Description
Shared storage support	vSAN supports VMware features that require shared storage, such as HA, vMotion, and DRS. For example, if a host becomes overloaded, DRS can migrate virtual machines to other hosts in the cluster.
On-disk format	vSAN on-disk virtual file format provides highly scalable snapshot and clone management support per vSAN cluster. For information about the number of virtual machine snapshots and clones supported per vSAN cluster, see the <i>Configuration Maximums</i> documentation.
All-flash and hybrid configurations	vSAN can be configured for all-flash or hybrid cluster.
Fault domains	vSAN supports configuring fault domains to protect hosts from rack or chassis failures when the vSAN cluster spans across multiple racks or blade server chassis in a data center.
File service	vSAN file service enables you to create file shares in the vSAN datastore that client workstations or VMs can access.
iSCSI target service	vSAN iSCSI target service enables hosts and physical workloads that reside outside the vSAN cluster to access the vSAN datastore.
Stretched cluster and Two node cluster	vSAN supports stretched clusters that span across two geographic locations.
Support for Windows Server Failover Clusters (WSFC)	vSAN 6.7 Update 3 and later releases support SCSI-3 Persistent Reservations (SCSI3-PR) on a virtual disk level required by Windows Server Failover Cluster (WSFC) to arbitrate an access to a shared disk between nodes. Support of SCSI-3 PRs enables configuration of WSFC with a disk resource shared between VMs natively on vSAN datastores. Currently the following configurations are supported: - Up to 6 application nodes per cluster. - Up to 64 shared virtual disks per node. Note: Microsoft SQL Server 2012 or later running on Microsoft Windows Server 2012 or later has been qualified on vSAN.
vSAN health service	vSAN health service includes preconfigured health check tests to monitor, troubleshoot, diagnose the cause of cluster component problems, and identify any potential risk.
vSAN performance service	vSAN performance service includes statistical charts used to monitor IOPS, throughput, latency, and congestion. You can monitor performance of a vSAN cluster, host, disk group, disk, and VMs.
Integration with vSphere storage features	vSAN integrates with vSphere data management features traditionally used with VMFS and NFS storage. These features include snapshots, linked clones, and vSphere Replication.
Virtual Machine Storage Policies	vSAN works with VM storage policies to support a VM-centric approach to storage management.

Supported Features	Description
	If you do not assign a storage policy to the virtual machine during deployment, the vSAN Default Storage Policy is automatically assigned to the VM.
Rapid provisioning	vSAN enables rapid provisioning of storage in the vCenter Server® during virtual machine creation and deployment operations.
Deduplication and compression	vSAN performs block-level deduplication and compression to save storage space. When you enable deduplication and compression on a vSAN all-flash cluster, redundant data within each disk group is reduced. Deduplication and compression is a cluster-wide setting, but the functions are applied on a disk group basis. Compression-only vSAN is applied on a per-disk basis.
Data at rest encryption	vSAN provides data at rest encryption. Data is encrypted after all other processing, such as deduplication, is performed. Data at rest encryption protects data on storage devices, in case a device is removed from the cluster.
Data in transit encryption	vSAN can encrypt data in transit across hosts in the cluster. When you enable data-in-transit encryption, vSAN encrypts all data and metadata traffic between hosts.
SDK support	The VMware vSAN SDK is an extension of the VMware vSphere Management SDK. It includes documentation, libraries and code examples that help developers automate installation, configuration, monitoring, and troubleshooting of vSAN.

Monitor vSAN in the VMware Host Client

You can use the VMware Host Client to monitor the vSAN environment of your ESXi host.

vSAN service must be enabled in the vSphere Client before you can view the vSAN related screens for a datastore.

1. Click **Storage** in the inventory.
2. On the **Datastores** tab, click **vSAN Datastore**.
The vSAN Datastore expands in the VMware Host Client navigator.
3. Click **Monitor**.
You are present with the **vSAN**, **Host**, and **Health** tabs in the UI.

Option	Description
vSAN	Displays the configurations for the current host. You can edit the settings for the claiming mode and deduplication. You can also view the settings for: • Encryption – vSAN supports encryption of the information for the whole vSAN datastore. • iSCSI Service – Additional service through the iSCSI service. • Performance Service - Collects data on how the datastore works. For example, the speed of a read/write operation.
Hosts	Displays a list of all the hosts on the vSAN server with their IP and the fault domain they belong to.

Option	Description
Health	The Health tab contains tests organized in groups. You are present with the following groups: • Performance Service • Network • Physical disk • Data • Cluster • Limits Each group is labeled with a status icon for an error, warning, unknown or healthy. The status of the group represents the most severe state of the test belonging to that group. To view the tests and their descriptions, click the expand icon in the top right corner of the group of interest. From the expanded card you can review all the tests belonging to the group, the result of their execution and get more information about what each test examines on the system.

4. Select the vSAN parameter that you want to monitor.


Edit Settings for a vSAN Datastore

You can edit the settings for a vSAN datastore when you must exit from a misconfigured state of the current host.

You can only edit the **Claiming Mode** and **Deduplication** settings for a vSAN datastore. These changes take effect only on the current host. They are not synced to the other hosts participating into the vSAN cluster.

Note: Only use these settings for troubleshooting.

1. Click **Storage** in the inventory.
2. On the **Datastores** tab, click a vSAN datastore from the table.
3. Click **Monitor** and click the **vSAN** tab.



- Click **Edit settings**.

The **Edit Settings** dialog box opens.

- Change the settings. Select **Auto** or **Manual** from the **Claiming Mode**.

Option	Action
Claiming Mode	- Select Auto or Manual from the Claiming Mode. - If you select Auto, it automatically takes all disks and claims them in a group or groups of the same size. Note: The Auto mode is deprecated. It can only claim hybrid disk groups which are not compatible with most of the vSAN features. - If you select Manual, you must manually organize the disks in groups and reclaim them by using the vSphere Web Client. For instance, selecting manual claiming mode is appropriate when the vCenter Server is unavailable.
Deduplication	Select Enabled or Disabled for the Deduplication.

- Click **Save**.

Performing Storage Refresh and Rescan Operations in the VMware Host Client

The refresh operation for datastores, storage devices, and storage adapters updates the lists and storage information that the VMware Host Client displays. It updates information such as the datastore capacity.

When you perform storage management tasks or make changes in the SAN configuration, you might need to rescan your storage.

Perform Adapter Rescan in the VMware Host Client

When you make changes in your SAN configuration and these changes are isolated to storage accessed through a specific adapter, perform rescan for only this adapter. When you rescan an adapter, you discover any new LUNs that are available on that adapter.

- Click **Storage** in the VMware Host Client inventory and click **Adapters**.
- Click **Rescan**.

Perform Device Rescan in the VMware Host Client

When you rescan a device, you discover any new VMFS volumes that are available on the device.

- Click **Storage** in the VMware Host Client inventory and click **Devices**.
- Click **Rescan**.

Change the Number of Scanned Storage Devices in the VMware Host Client

The range of scanned LUN IDs for an ESXi host can be from 0 to 16,383. ESXi ignores LUN IDs greater than 16,383. The configurable `Disk.MaxLUN` parameter controls the range of scanned LUN ID range. The parameter has a default value of 1024.

The `Disk.MaxLUN` parameter also determines how many LUNs the SCSI scan code attempts to discover using individual INQUIRY commands if the SCSI target does not support direct discovery using REPORT_LUNS.

You can modify the `Disk.MaxLUN` parameter depending on your needs. For example, if your environment has a smaller number of storage devices with LUN IDs from 1 through 100, set the value to 101. As a result, you can improve device

discovery speed on targets that do not support REPORT_LUNS. Lowering the value can shorten the rescan time and boot time. However, the time to rescan storage devices might also depend on other factors, including the type of the storage system and the load on the storage system.

In other cases, you might need to increase the value if your environment uses LUN IDs that are greater than 1023.

1. Click **Manage** in the VMware Host Client inventory and click **Advanced Settings**.
2. Scroll down to `Disk.MaxLUN`.
3. Right-click `Disk.MaxLUN`, and click **Edit option**.
4. Enter a new value and click **Save**.

The SCSI scan code does not scan the LUNs with IDs greater than or equal to the value you enter.

For example, to discover LUN IDs from 0 to 100, set `Disk.MaxLUN` to 101.

Networking in the VMware Host Client

When you connect to an ESXi host using the VMware Host Client, you can view and configure vSphere standard switches, port groups, physical NICs, VMkernel NICs, and TCP/IP stacks.

Manage Port Groups in the VMware Host Client

You can manage port group settings to configure traffic management, enhance networking security, and enhance performance.

By using the VMware Host Client, you can add and remove port groups. You can also examine port group information and edit port group settings, such as NIC teaming and traffic shaping.

View Port Group Information in the VMware Host Client

In the VMware Host Client, you can view information about port group configuration, network details, virtual switch topology, NIC teaming policy, offload policy, and security policy.

1. Click **Networking** in the VMware Host Client inventory and click **Port groups**.
2. Click an item from the list of available port groups.
Information about network details, virtual switch topology, NIC teaming policy, offload policy, and security policy is displayed.

Add a Virtual Switch Port Group in the VMware Host Client

You can add a port group to a virtual switch in the VMware Host Client. Port groups provide networking for virtual machines.

1. Right-click **Networking** in the VMware Host Client inventory and click **Add port group** from the pop-up menu.
2. Enter a name for the new port group.
3. Set the VLAN ID to configure VLAN handling in the port group.
The VLAN ID also reflects the VLAN tagging mode in the port group.

VLAN Tagging Mode	VLAN ID	Description
External Switch Tagging (EST)	0	The virtual switch does not pass traffic associated with a VLAN.
Virtual Switch Tagging (VST)	From 1 to 4094	The virtual switch tags traffic with the tag that you entered.

VLAN Tagging Mode	VLAN ID	Description
Virtual Guest Tagging (VGT)	4095	Virtual machines handle VLANs. The virtual switch permits traffic from any VLAN.

4. Select a virtual switch from the drop-down menu.
5. Expand **Security** and select options that you want to enable for promiscuous mode, MAC address changes, and forged transmits.
6. Click **Add**.
Your port group is created.
7. Optional: Click **Refresh** to display the new port group in the list.

Edit Port Group Settings in the VMware Host Client

To enhance networking security and improve networking performance in the VMware Host Client, you can edit various port group settings, such as the port group name, VLAN ID, and virtual switch. You can also configure security, NIC teaming, and traffic shaping components.

1. Click **Networking** in the VMware Host Client inventory and click **Port groups**.
2. Right-click the port group in the list that you want to edit and select **Edit settings**.
3. Optional: Enter a new port group name.
4. Optional: Enter a new value for the VLAN ID.
The VLAN ID reflects the VLAN tagging mode in the port group.

VLAN Tagging Mode	VLAN ID	Description
External Switch Tagging (EST)	0	The virtual switch does not pass traffic associated with a VLAN.
Virtual Switch Tagging (VST)	From 1 to 4094	The virtual switch tags traffic with the tag that you entered.
Virtual Guest Tagging (VGT)	4095	Virtual machines handle VLANs. The virtual switch permits traffic from any VLAN.

5. Optional: Select a virtual switch from the drop-down menu.
6. Optional: Expand **Security** and select whether to reject, accept, or inherit the Security policy exceptions from vSwitch.

Option	Description
Promiscuous Mode	• Reject. Placing a guest adapter in promiscuous mode has no effect on which frames are received by the adapter. • Accept. Placing a guest adapter in promiscuous mode causes it to detect all frames passed on the vSphere distributed switch that are allowed under the VLAN policy for the port group that the adapter is connected to. • Inherit from vSwitch. Placing a guest adapter in promiscuous mode causes it to inherit the configuration from the associated virtual switch.
MAC Address Changes	• Reject. If you set the MAC Address Changes to Reject and the guest operating system changes the MAC address

Option	Description
	of the adapter to anything other than what is in the <code>.vmx</code> configuration file, all inbound frames are dropped. If the guest operating system changes back the MAC address to match the MAC address in the <code>.vmx</code> configuration file, inbound frames are passed again. • Accept. Changing the MAC address from the guest operating system has the intended effect: frames to the new MAC address are received. • Inherit from vSwitch. If you set MAC Address Changes to Inherit from vSwitch, the MAC address changes to one of the associated virtual switches.
Forged Transmits	• Reject. Any outbound frame with a source MAC address that is different from the one set on the adapter are dropped. • Accept. No filtering is performed and all outbound frames are passed. • Inherit from vSwitch. The outbound frame configuration is inherited from the associated virtual switch.

7. Optional: Expand **NIC teaming** and configure the following components.

Option	Description
Load Balancing	Specify how to choose an uplink. • Inherit from vSwitch. Choose the uplink that is selected for the associated virtual switch. • Route based on IP hash. Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash. • Route based on source MAC hash. Choose an uplink based on a hash of the source Ethernet. • Route based on originating port ID. Choose an uplink based on the originating port ID. • Use explicit failover order. Always use the highest order uplink from the list of active adapters which passes failover detection criteria . Note: IP-based teaming requires the physical switch to be configured with EtherChannel. For all other options, EtherChannel must be deactivated.
Network Failover Detection	Specify the method to use for failover detection. • • Inherit from vSwitch. Inherits the respective configuration of the associated virtual switch. • Link Status only. Relies only on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by a spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch. • Beacon only. Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine a link failure. This detects many of the failures that are not detected by link status only.

Option	Description
	Note: Do not use beacon probing with IP-hash load balancing.
Notify Switches	Select Yes, No, or Inherit from vSwitch to notify switches if a failover occurs. If you select Yes, when a virtual NIC is connected to the distributed switch or that virtual NIC's traffic is routed over a different physical NIC in the team because of a failover event, a notification is sent out over the network to update the lookup tables on physical switches. In almost all cases, this process is preferred for the lowest latency of failover occurrences and migrations with vMotion. Note: Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode. No such issue exists with NLB running in multicast mode.
Failback	Select Yes, No, or Inherit from vSwitch to deactivate or activate failback. This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to the default setting of Yes, the adapter returns to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. If failback is set to No, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.
Failover Order	Specify how to distribute the workload for uplinks. If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, set this condition by moving them into different groups: • Active Uplinks. Continue to use the uplink when the network adapter connectivity is up and active. • Standby Uplinks. Use this uplink if one of the active adapter's connectivities is down. Note: When using IP-hash load balancing, do not configure standby uplinks. You cannot configure failover order if any of the port group components are configured to inherit the configuration from the associated virtual switch.

8. Optional: To configure traffic shaping, expand **Traffic shaping**, click **Enabled**, and specify the following parameters.

Option	Description
Average Bandwidth	Establishes the number of bits per second to limit across a port, averaged over time—the allowed average load.
Peak Bandwidth	The maximum number of bits per second to limit across a port when it is sending/receiving a burst of traffic. This is the maximum bandwidth used by a port whenever it is using its burst bonus.
Burst Size	The maximum number of bytes to limit in a burst. If this parameter is set, a port might gain a burst bonus when it does not use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by Average Bandwidth , it might be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter represents the

Option	Description
	maximum number of bytes that might be accumulated in the burst bonus and so transferred at a higher speed.

Traffic shaping policy is applied to the traffic of each virtual network adapter attached to the virtual switch.

- Click **Save** to apply your changes.

Remove a Virtual Switch Port Group in the VMware Host Client

You can remove port groups from virtual switches in case you no longer need the associated labeled networks.

Verify that there are no VMkernel NICs and no powered-on virtual machines connected to the port group that you want to remove.

- Click **Networking** in the VMware Host Client inventory and click the **Port groups** tab.
- Right-click the port group that you want to remove and select **Remove** from the pop-up menu.
- To remove the port group, click **Remove**.
- Optional: Click **Refresh** to verify that you have removed the port group.

Manage Virtual Switches in the VMware Host Client

In the VMware Host Client, you can configure various virtual switch settings, such as link discovery, NIC teaming, and traffic shaping.

View Virtual Switch Information in the VMware Host Client

In the VMware Host Client, you can view information about virtual switches, such as configuration, network details, virtual switch topology, and so on.

- Click **Networking** in the VMware Host Client inventory and click **Virtual switches**.
- Click on a switch from the list of available virtual switches.

Information about virtual switch configuration, network details, and virtual switch topology is displayed.

Add a Standard Virtual Switch in the VMware Host Client

In the VMware Host Client, you can add a standard virtual switch to provide network connectivity for the host that you are managing and for the virtual machines on that host, and to handle VMkernel traffic. Depending on the type of connection that you want to create, you can create vSphere standard switch with a VMkernel adapter, connect an existing physical network adapter to the new switch, or create the switch with a virtual machine port group.

- Right-click **Networking** in the VMware Host Client inventory and click **Add standard vSwitch** in the pop-up menu.
- Optional: Click **Add uplink** to add a new physical uplink to a virtual switch.
- Enter a name for the virtual switch and click **Create virtual switch**.
- Select an uplink for the virtual switch.
- Expand **Link discovery** and select an option for the virtual switch mode.

Operation	Description
Listen	ESXi detects and displays information about the associated physical switch port, but information about the vSphere Standard Switch is not available to the switch administrator.

Operation	Description
Advertise	ESXi makes information about the vSphere Standard Switch available to the switch administrator, but does not detect and display information about the physical switch.
Both	ESXi detects and displays information about the associated physical switch and makes information about the vSphere Standard Switch available to the switch administrator.
None	ESXi does not detect or display information about the associated physical switch port, and information about the vSphere Standard Switch is not available to the switch administrator.

6. In the Protocol section, select **Cisco Discovery Protocol** from the drop-down menu.
7. Expand **Security** and accept or reject promiscuous mode, MAC address changes, and forged transmits of the virtual machines attached to the standard switch.

Option	Description
Promiscuous mode	• Reject. The VM network adapter receives only frames that are addressed to the virtual machine. • Accept. The virtual switch forwards all frames to the virtual machine in compliance with the active VLAN policy for the port to which the VM network adapter is connected. Note: Promiscuous mode is insecure mode of operation. Firewalls, port scanners, intrusion detection systems, must run in promiscuous mode.
MAC address changes	• Reject. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter (set in the <code>.vmx</code> configuration file), the switch drops all inbound frames to the adapter. If the guest OS changes the effective MAC address of the virtual machine back to the MAC address of the VM network adapter, the virtual machine receives frames again. • Accept. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter, the switch allows frames to the new address to pass.
Forged transmits	• Reject. The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the <code>.vmx</code> configuration file. • Accept. The switch does not perform filtering, and permits all outbound frames.

- Click **Add**.

Remove a Standard Virtual Switch in the VMware Host Client

You can remove the virtual standard switch if you no longer need it.

- Click **Networking** in the VMware Host Client inventory and click the **Virtual switches** tab.
- Right-click the virtual switch that you want to remove from the list and click **Remove**.
- Click **Yes**.

Add a Physical Uplink to a Virtual Switch in the VMware Host Client

You can connect multiple adapters to a single vSphere standard switch to provide NIC teaming. The team can share traffic and provide failover.

- Click **Networking** in the VMware Host Client inventory and click **Virtual switches**.
- Click a virtual switch from the list and click **Add uplink**.
- Select a physical NIC from the available options.
- Click **Save**.

Edit Virtual Switch Settings in the VMware Host Client

In the VMware Host Client, you can edit virtual switch settings, such as the virtual switch uplinks.

- Click **Networking** in the VMware Host Client inventory and click **Virtual switches**.
- Right-click the virtual switch that you want to edit and click **Edit Settings**.
- Optional: Click **Add uplink** to add a new physical uplink to the virtual switch.
- Change the maximum transmission unit (MTU).
The MTU improves the networking efficiency by increasing the amount of payload data transmitted with a single packet, that is, enabling jumbo frames.
- Optional: Click the **Remove** icon (✖) to remove the old uplink from the virtual switch.
- Expand **Link discovery** and select an option for the virtual switch mode.

Operation	Description
Listen	ESXi detects and displays information about the associated physical switch port, but information about the vSphere Standard Switch is not available to the switch administrator.
Advertise	ESXi makes information about the vSphere Standard Switch available to the switch administrator, but does not detect and display information about the physical switch.
Both	ESXi detects and displays information about the associated physical switch and makes information about the vSphere Standard Switch available to the switch administrator.
None	ESXi does not detect or display information about the associated physical switch port, and information about the vSphere Standard Switch is not available to the switch administrator.

- In the Protocol section, select **Cisco Discovery Protocol** from the drop-down menu.
- Expand **Security** and accept or reject promiscuous mode, MAC address changes, and forged transmits of the virtual machines attached to the standard switch.

Option	Description
Promiscuous mode	Reject. The VM network adapter receives only frames that are addressed to the virtual machine.

Option	Description
	• Accept. The virtual switch forwards all frames to the virtual machine in compliance with the active VLAN policy for the port to which the VM network adapter is connected. Note: Promiscuous mode is insecure mode of operation. Firewalls, port scanners, intrusion detection systems, must run in promiscuous mode.
MAC address changes	• Reject. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter (set in the <code>.vmx</code> configuration file), the switch drops all inbound frames to the adapter. If the guest OS changes the effective MAC address of the virtual machine back to the MAC address of the VM network adapter, the virtual machine receives frames again. • Accept. If the guest OS changes the effective MAC address of the virtual machine to a value that is different from the MAC address of the VM network adapter, the switch allows frames to the new address to pass.
Forged transmits	• Reject. The switch drops any outbound frame from a virtual machine adapter with a source MAC address that is different from the one in the <code>.vmx</code> configuration file. • Accept. The switch does not perform filtering, and permits all outbound frames.

9. Optional: Expand **NIC teaming** and configure the following components.

Option	Description
Load Balancing	Specify how to choose an uplink. • Route based on IP hash. Choose an uplink based on a hash of the source and destination IP addresses of each packet. For non-IP packets, whatever is at those offsets is used to compute the hash. • Route based on source MAC hash. Choose an uplink based on a hash of the source Ethernet. • Route based on originating port ID. Choose an uplink based on the originating port ID. • Use explicit failover order. Always use the highest order uplink from the list of Active adapters which passes failover detection criteria. Note: IP-based teaming requires the physical switch to be configured with EtherChannel. For all other options, EtherChannel must be deactivated.
Network Failover Detection	Specify the method to use for failover detection. • Link Status only. Relies only on the link status that the network adapter provides. This option detects failures, such as cable pulls and physical switch power failures, but not configuration errors, such as a physical switch port being blocked by spanning tree or that is misconfigured to the wrong VLAN or cable pulls on the other side of a physical switch. • Beacon only. Sends out and listens for beacon probes on all NICs in the team and uses this information, in addition to link status, to determine link failure. This detects many of the failures previously mentioned that are not detected by link status alone. Note: Do not use beacon probing with IP-hash load balancing.
Notify Switches	Select Yes, No, or Inherit from vSwitch to notify switches in the case of failover.

Option	Description
	If you select Yes, whenever a virtual NIC is connected to the distributed switch or whenever that virtual NIC's traffic might be routed over a different physical NIC in the team because of a failover event, a notification is sent out over the network to update the lookup tables on physical switches. In almost all cases, this process is desirable for the lowest latency of failover occurrences and migrations with vMotion. Note: Do not use this option when the virtual machines using the port group are using Microsoft Network Load Balancing in unicast mode. No such issue exists with NLB running in multicast mode.
Failback	Select Yes, No, or Inherit from vSwitch to deactivate or activate failback. This option determines how a physical adapter is returned to active duty after recovering from a failure. If failback is set to Yes (default), the adapter is returned to active duty immediately upon recovery, displacing the standby adapter that took over its slot, if any. If failback is set to No, a failed adapter is left inactive even after recovery until another currently active adapter fails, requiring its replacement.
Failover Order	Specify how to distribute the work load for uplinks. If you want to use some uplinks but reserve others for emergencies in case the uplinks in use fail, set this condition by moving them into different groups: • Active Uplinks. Continue to use the uplink when the network adapter connectivity is up and active. • Standby Uplinks. Use this uplink if one of the active adapter's connectivities is down. Note: When using IP-hash load balancing, do not configure standby uplinks.

10. Optional: To configure traffic shaping, expand **Traffic shaping**, click **Enabled**, and specify the following parameters.

Option	Description
Average Bandwidth	Establishes the number of bits per second to allow across a port, averaged over time—the allowed average load.
Peak Bandwidth	The maximum number of bits per second to allow across a port when it is sending or receiving a burst of traffic. This tops the bandwidth used by a port whenever it is using its burst bonus.
Burst Size	The maximum number of bytes to allow in a burst. If this parameter is set, a port may gain a burst bonus when it doesn't use all its allocated bandwidth. Whenever the port needs more bandwidth than specified by Average Bandwidth , it may be allowed to temporarily transmit data at a higher speed if a burst bonus is available. This parameter tops the number of bytes that may be accumulated in the burst bonus and then transferred at a higher speed.

Traffic shaping policy is applied to the traffic of each virtual network adapter attached to the virtual switch.

11. Click **Save**.

Manage Physical Network Adapters in the VMware Host Client

Assign a physical adapter to a standard switch to provide connectivity to virtual machines and VMkernel adapters on the host that you are managing.

View Physical Network Adapter Information in the VMware Host Client

In the VMware Host Client, you can view various information about physical network adapter (NIC) configuration and settings.

1. Click **Networking** in the VMware Host Client inventory and click **Physical NICs**.
2. Click on the network adapter that you want to view information about.

Edit Physical NICs in the VMware Host Client

You can edit physical NIC speed by using the VMware Host Client.

1. Click **Networking** in the VMware Host Client inventory and click **Physical NICs**.
2. Select the NIC in the table that you want to edit.
3. Click **Edit settings** and select a speed from the drop-down menu.
4. Click **Save**.

Manage VMkernel Network Adapters in the VMware Host Client

In the VMware Host Client, you can add and remove VMkernel network adapters (NICs), as well as view and modify the VMkernel NIC settings.

View VMkernel Network Adapter Information in the VMware Host Client

In the VMware Host Client, you can view information about VMkernel network adapters (NICs), such as TCP/IP configuration, network details, virtual switch topology, and so on.

1. Click **Networking** in the VMware Host Client inventory and click **VMkernel NICs**.
2. Click a NIC in the list to display configuration and topology details.

Add a VMkernel Network Adapter in the VMware Host Client

You can add a VMkernel network adapter (NIC) on a VMware vSphere® Standard Edition™ switch to provide network connectivity for hosts. The VMkernel NIC also handles the system traffic for VMware vSphere® vMotion®, IP storage, Fault Tolerance, logging, vSAN, and so on.

1. Right-click **Networking** in the VMware Host Client inventory and click **Add VMkernel NIC**.
2. In the **Add VMkernel NIC** dialog box, configure the settings for the VMkernel adapter.

Option	Description
New port group Label	Adding a VMkernel NIC also adds a port group. Specify a name for that port group.
VLAN ID	Enter a VLAN ID to determine the VLAN for the network traffic of the VMkernel adapter to use.
IP Version	Select IPv4, IPv6, or both. Note: The IPv6 option does not appear on hosts that do not have IPv6 enabled.

3. Select a virtual switch from the drop-down menu.

4. Optional: Expand the IPv4 settings section to select an option for obtaining IP addresses.

Option	Description
Use DHCP to obtain IP settings	IP settings are obtained automatically. A DHCP server must be present on the network.
Use static IP settings	Enter the IPv4 address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack.

5. Optional: Expand the IPv6 settings section to select an option for obtaining IPv6 addresses.

Option	Description
DHCPv6	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Auto Configuration	Use router advertisement to obtain IPv6 addresses.
Static IPv6 Addresses	1. Click Add address to add a new IPv6 address. 2. Enter the IPv6 address and subnet prefix length.

6. Select a TCP/IP stack from the drop-down menu.

After you set a TCP/IP stack for the VMkernel adapter, you cannot change it. If you select the vMotion or the Provisioning TCP/IP stack, you can use only this stack to handle vMotion or for Provisioning traffic on the host. All VMkernel adapters for vMotion on the default TCP/IP stack are deactivated for future vMotion sessions. If you use the Provisioning TCP/IP stack, VMkernel adapters on the default TCP/IP stack are deactivated, and you cannot perform some operations. Such operations include traffic Provisioning, such as virtual machine cold migration, cloning, and snapshot migration.

7. Optional: Select the services to enable for the default TCP/IP stack on the host.

vMotion enables the VMkernel adapter to advertise itself to another host as the network connection where vMotion traffic is sent. You cannot use vMotion to perform migrations to selected hosts if the vMotion service is not enabled for any VMkernel adapter on the default TCP/IP stack, or if no adapters use the vMotion TCP/IP stack.

8. Review your setting selections and click **Create**.

Edit VMkernel Network Adapter Settings in the VMware Host Client

You might have to change the supported traffic type for a VMkernel network adapter, or the way IPv4 or IPv6 addresses are obtained.

1. Click **Networking** in the VMware Host Client inventory and click **VMkernel NICs**.

2. Select the VMkernel adapter that resides on the target standard switch, click **Actions**, and select **Edit settings** from the drop-down menu.

3. Optional: Edit the VLAN ID.

The VLAN ID determines the VLAN that the network traffic of the VMkernel adapter uses.

4. Optional: To edit the IP version, select IPv4, IPv6, or both from the drop-down menu.

Note: The IPv6 option does not appear on hosts that do not have IPv6 enabled.

5. Optional: Expand the IPv4 settings section to select an option for obtaining IP addresses.

Option	Description
Use DHCP to obtain IP settings	IP settings are obtained automatically. A DHCP server must be present on the network.

Option	Description
Use static IP settings	Enter the IPv4 address and subnet mask for the VMkernel adapter. The VMkernel Default Gateway and DNS server addresses for IPv4 are obtained from the selected TCP/IP stack.

6. Optional: Expand the IPv6 settings section to select an option for obtaining IPv6 addresses.

Option	Description
DHCPv6	Use DHCP to obtain IPv6 addresses. A DHCPv6 server must be present on the network.
Auto Configuration	Use a router advertisement to obtain IPv6 addresses.
Static IPv6 Addresses	1. Click Add address to add an IPv6 address. 2. Enter the IPv6 address and subnet prefix length.

7. Optional: Select the service to activate or deactivate for the default TCP/IP stack on the host.

vMotion enables the VMkernel adapter to advertise itself to another host as the network connection where the vMotion traffic is sent. It is not possible to use vMotion to perform migrations to selected hosts if the vMotion service is not enabled for any VMkernel adapter on the default TCP/IP stack, or if no adapters use the vMotion TCP/IP stack.

8. Review your setting modifications and click **Save** to apply your changes.

Remove a VMkernel Network Adapter in the VMware Host Client

In the VMware Host Client, you can remove a VMkernel network adapter if you no longer need it.

1. Click **Networking** in the VMware Host Client inventory and click **VMkernel NICs**.
2. Right-click the VMkernel adapter that you want to remove and click **Remove**.
3. Click **Confirm** to remove the network adapter.

View TCP/IP Stack Configuration on a Host in the VMware Host Client

You can view the DNS and routing configuration of a TCP/IP stack on a host. You can also view the IPv4 and IPv6 routing tables, the congestion control algorithm, and the maximum number of allowed connections.

1. Click **Networking** in the host inventory and click **TCP/IP stacks**.
2. Click a stack from the list.
The configuration settings of the stack you selected are displayed.

Change the Configuration of a TCP/IP Stack on a Host in the VMware Host Client

You can change the DNS and default gateway configuration of a TCP/IP stack on a host. You can also change the congestion control algorithm, the maximum number of connections, and the name of custom TCP/IP stacks.

1. Click **Networking** in the VMware Host Client inventory and click **TCP/IP stacks**.
2. Right-click a stack from the list and select **Edit settings**.
The Edit TCP/IP configuration - Provisioning stack dialog box opens.

3. Specify how the host obtains settings for this TCP/IP stack.

- Select the **Use DHCP services from the following adapter** radio button, and select an adapter from which to receive the default settings configuration for the TCP/IP stack.
- Select the **Manually configure the settings for this TCP/IP stack** to change the settings configuration.

Option	Description
Basic configuration	Host name Edit the name of your local host. Domain name Edit the domain name. Primary DNS server Enter a preferred DNS server IP address. Secondary DNS server Type an alternate DNS server IP address. Search domains Specify DNS suffixes to use in DNS search when resolving unqualified domain names.
Routing	Edit the IPv4 and IPv6 gateway information. Note: Removing the default gateway might cause you to lose your connection to the host.
Advanced Settings	Edit the congestion control algorithm and the maximum number of connections.

4. Click **Save**.

Configure ESXi Firewall in the VMware Host Client

ESXi includes a firewall that is enabled by default. During installation, the ESXi firewall is configured to block incoming and outgoing traffic, except traffic for services that are enabled in the host security profile.

As you open ports on the firewall, consider that unrestricted access to services running on an ESXi host may expose a host to outside attacks and unauthorized access. Reduce the risk by configuring the ESXi firewall to allow access only from authorized networks.

Note: The firewall also allows Internet Control Message Protocol, or ICMP, pings and communication with DHCP and DNS (UDP only) clients.

Manage ESXi Firewall Settings by Using the VMware Host Client

When you are logged in to an ESXi host with the VMware Host Client, you can configure incoming and outgoing firewall connections for a service or a management agent.

Note: If different services have overlapping port rules, enabling one service might implicitly enable other services. You can specify which IP addresses are allowed to access each service on the host to avoid this problem.

1. Click **Networking** in the inventory.

2. Click **Firewall rules**.

The VMware Host Client displays a list of active incoming and outgoing connections with the corresponding firewall ports.

3. For some services you can manage service details. Right-click a service and select an option from the pop-up menu.

- Use the Start, Stop, or Restart buttons to change the status of a service temporarily.
- Change the Startup Policy to configure the service to start and stop with the host, the firewall ports, or manually.

Add Allowed IP Addresses for an ESXi Host by Using the VMware Host Client

By default, the firewall for each service allows access to all IP addresses. To restrict traffic, configure each service to allow traffic only from your management subnet. You can also deselect some services if your environment does not use them.

1. Click **Networking** in the inventory and click **Firewall rules**.
2. Click a service from the list and click **Edit settings**.
3. In the Allowed IP Addresses section, click **Only allow connections from the following networks** and enter the IP addresses of networks that you want to connect to the host.
Separate IP addresses with commas. You can use the following address formats:
 - 192.168.0.0/24
 - 192.168.1.2, 2001::1/64
 - fd3e:29a6:0a81:e478::/64
4. Click **OK**.

Monitor Networking Events and Tasks in the VMware Host Client

You can view details about the events and tasks associated with the port groups, virtual switches, physical network adapters, VMkernel network adapters, and TCP/IP stacks on the ESXi host that you are managing.

Monitor Port Groups in the VMware Host Client

In the VMware Host Client, you can monitor port group performance by viewing the events and tasks of the port groups on the host.

1. Click **Networking** in the inventory.
2. Click **Port groups**.
3. Click a port group from the list.
The port group expands in the VMware Host Client inventory.
4. Click **Monitor** under the port group name in the VMware Host Client inventory.
5. Optional: Click **Events** to view the events associated with the port group.

Monitor Virtual Switches in the VMware Host Client

In the VMware Host Client, you can monitor virtual switch performance by viewing the events and tasks of the virtual switches on the host.

1. Click **Networking** in the inventory.
2. Click **Virtual switches**.
3. Click a virtual switch from the list.
The virtual switch expands in the VMware Host Client inventory.

4. Click **Monitor** under the virtual switch name in the VMware Host Client inventory.
5. Optional: Click **Events** to view the events associated with the virtual switch.

Monitor Physical Network Adapters in the VMware Host Client

In the VMware Host Client, you can monitor physical network adapter (NIC) performance by viewing the events and tasks of the physical NICs on the host.

1. Click **Networking** in the inventory.
2. Click **Physical NICs**.
3. Click a physical network adapter from the list.
The physical network adapter expands in the VMware Host Client inventory.
4. Click **Monitor** under the physical network adapter name in the VMware Host Client inventory.
5. Optional: Click **Events** to view the events associated with the physical network adapter.

Monitor VMkernel Network Adapters in the VMware Host Client

In the VMware Host Client, you can monitor VMkernel network adapter performance by viewing the events and tasks of the VMkernel network adapters on the host.

1. Click **Networking** in the inventory.
2. Click **VMkernel NICs**.
3. Click a VMkernel network adapter from the list.
The VMkernel network adapter expands in the VMware Host Client inventory.
4. Click **Monitor** under the VMkernel network adapter name in the VMware Host Client inventory.
5. Optional: Click **Events** to view the events associated with the VMkernel network adapter.

Monitor TCP/IP Stacks in the VMware Host Client

In the VMware Host Client, you can monitor TCP/IP stacks performance by viewing the events and tasks of the TCP/IP stacks on the host.

1. Click **Networking** in the inventory.
2. Click **TCP/IP stacks**.
3. Click a TCP/IP stack from the list.
The TCP/IP stack expands in the VMware Host Client inventory.
4. Click **Monitor** under the TCP/IP stack name in the VMware Host Client inventory.
5. Optional: Click **Events** to view the events associated with the TCP/IP stack.
6. Optional: Click **Tasks** to view the tasks associated with the TCP/IP stack.

Setup for Windows Server Failover Clustering on VMware vSphere

Setup for Windows Server Failover Clustering describes the supported configurations for a WSFC with shared disk resources you can implement using virtual machines with Failover Clustering for Windows Server 2012 and above releases. You get step-by-step instructions for each configuration and a checklist of clustering requirements and recommendations.

Unless stated otherwise, the term Windows Server Failover Clustering (WSFC) applies to Failover Clustering with Windows Server 2012 and above releases.

Setup for Windows Server Failover Clustering covers ESXi and VMware® vCenter® Server.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

Intended Audience

This information is for system administrators who are familiar with VMware technology and Failover Clustering for Windows.

Note: This is not a guide to using Failover Clustering for Windows. Use your Microsoft documentation for information about installation and configuration of Microsoft Cluster Service or Failover Clustering.

Note:

All instances of "WSFC" in this document and elsewhere is intended to refer to all Windows OS and Microsoft application clustering configurations which utilize one or more shared disks among the clustered nodes, namely:

- The Microsoft Cluster Service (MSCS)
- The Windows Server Failover Clustering (WSFC)
- Microsoft SQL Server Always On Failover Cluster Instances (FCI)

Note: Other WSFC-based solutions not accessing shared storage such as SQL Server Always On Availability Groups or Exchange Database Availability Group, require no special storage configurations on the vSphere side using VMFS or NFS. This guide should not be used for such configurations.

Getting Started with WSFC on VMware® vSphere®

VMware® vSphere® supports Windows clustering using WSFC across virtual machines. Clustering virtual machines can reduce the hardware costs of traditional high-availability Windows clusters.

Note: vSphere High Availability (vSphere HA) supports a Windows clustering solution. *vSphere Availability* describes vSphere HA functionality.

Clustering Configuration on VMware® vSphere®

Several applications use clustering, including stateless applications such as Web servers, and applications with built-in recovery features such as database servers. You can set up WSFC clusters in several configurations, depending on your environment.

A typical clustering setup includes disks resources that are shared between nodes. A shared disk is required as a quorum disk. In a cluster of virtual machines across physical hosts, shared disks can be RDMS, vVol or VMFS VMDKs.

Note: **VMDKs** in this guide refers to Shared disk files from VMFS datastore shared in physical compatibility mode (Attached to SCSI/NVMe controller with bus sharing mode set to **Physical**) for cluster across physical machines. This is also known as Clustered VMDKs.

If you are using RDMS or vVol as shared disks, they can be accessed using Fibre Channel (FC) SAN, FCoE, or iSCSI. In vSphere 8.0 U3 or later vVol also supports shared disk from NVMe storage arrays (FC, TCP).

If you are using VMFS VMDKs as shared disks, only FC SAN connected storage devices are supported. FC SAN can be from SCSI and NVMe storage arrays. VMFS VMDKs are not supported by other SANs such as iSCSI or FCoE.

In vSphere 8.0 U3 or later VMFS VMDKs also support storage from NVMe TCP arrays.

In ESXi, clustered VMDKs on VMFS datastores are supported in configurations when the VMs hosting the nodes of a cluster are located on different ESXi hosts, known as cluster-across-boxes (CAB) configuration.

Hosting virtual nodes of a WSFC on a Single Host

A cluster of WSFC virtual machines on a single host (also known as a cluster in a box or CIB) consists of clustered virtual machines on the same ESXi host. The virtual machines are connected to the same storage, either local or remote. This configuration protects against failures at the operating system and application level, but it does not protect against hardware failures.


Note: Do not host virtual nodes of a WSFC on the same ESXi host for a production deployment.

Note: Windows Server 2012 and above releases support up to five nodes (virtual machines).

The following figure shows a cluster in a box setup.

- Two virtual machines on the same ESXi host run WSFC.
- The virtual machines share a private network connection for the private heartbeat and a public network connection.
- Each virtual machine is connected to shared storage, which can be local or on a SAN.
- Each shared disk should be attached to SCSI controller with bus sharing mode set to **Virtual**.

Figure 141: Virtual Machines Clustered on a Single Host



Clustering Virtual Machines Across Physical Hosts


A cluster of virtual machines across physical ESXi hosts (also known as a cluster across boxes or CAB) protects against software failures and hardware failures on the physical ESXi hosts by placing the cluster nodes on separate ESXi hosts. This configuration requires shared storage for the clustered disk resources.

The following figure shows a cluster-across-boxes setup.

- Two virtual machines on two different ESXi hosts run WSFC.
- The virtual machines share a private network connection for the private heartbeat and a public network connection.
- Each virtual machine is connected to shared storage.
- ESXi host supports up to 16 WSFC clusters (multi-cluster) running on the same ESXi host.

Note: In a cluster of virtual machines across physical hosts, shared disks can be physical mode or vVol with vSphere 6.7 or later. You can use VMFS VMDKs from SCSI SAN with vSphere 7.0 or later and from NVMe SAN with vSphere 8.0 or later. If you are using RDMs or vVol as shared disks, they can be on Fibre Channel (FC) SAN, FCoE or iSCSI. If you are using VMFS VMDKs as shared disks, they can be on FC SAN. vSphere 8.0 U3 or later vVol will also support shared disk from NVMe storage arrays (FC, TCP). vSphere 8.0 U3 or later VMFS VMDKs will also support storage from NVMe TCP array.

Figure 142: Virtual Machines Clustered Across Hosts



Note: Windows Server 2012 and above systems support up to five nodes (virtual machines). For supported guest operating systems see [Table 4](#).

Note: Private and Public network interconnect can share a single virtual NIC in a VM.

You can expand the cluster-across-boxes model and place multiple virtual machines on multiple ESXi hosts.


Clustering Physical Machines with Virtual Machines

For a simple WSFC clustering solution with low hardware requirements with RDMs, you might choose to have one standby host.

Set up your system to have a virtual machine corresponding to each physical machine on the standby host, and create clusters, one each for each physical machine and its corresponding virtual machine. In case of hardware failure in one of the physical machines, the virtual machine on the standby host can take over for that physical host.

The following figure shows a standby host using three virtual machines on a single ESXi host. Each virtual machine is running WSFC.

Note: This solution is not applicable for vVol and clustered VMDKs.

Figure 143: Clustering Physical and Virtual Machines

Hardware and Software Requirements for WSFC on VMware®vSphere®

All vSphere WSFC configurations require certain hardware and software components.

The following table lists hardware and software requirements that apply to all vSphere WSFC configurations.

Table 568: Clustering Requirements

Component	Requirement
Virtual SCSI adapter	LSI Logic SAS for Windows Server 2012 or later. VMware Paravirtual for Windows Server 2012, 2012 R2. VMware Paravirtual for Windows Server 2016 or later. (recommended) NVMe Controller for Windows Server 2022 with OS Build 20348.1547 and later, virtual machine hardware version must be 21 or later.
Operating system	Windows Server 2012 or later releases. For supported guest operating systems see Table 4 .
Virtual NIC	VMXNET3 is required. You must activate in-guest Windows Receive Side Scaling (RSS) on the virtual NIC.

Component	Requirement
	Note: Drivers, improvements and updates to the VMXNET3 virtual NIC are delivered through VMware Tools. Upgrade to the latest version of VMware Tools on any VM using the VMXNET3 virtual NIC.
I/O timeout	Set to 60 seconds or more. Modify <code>HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue</code> . The system might reset this I/O timeout value if you re-create a cluster. You must reset the value in that case.
Disk format	pRDM (physical compatibility mode) vVol disks can be thin or thick format VMFS VMDKs with vSphere 7.0 <code>eagerzeroedthick</code> format
Disk and networking setup	Add networking before disks.
Number of nodes	Windows Server 2012 or later: up to five-node clustering For supported guest operating systems see Table 4 .
NTP server	Synchronize Windows Active Directory domain controllers and WSFC nodes with a common NTP server, and deactivate host-based time synchronization. See KB 1189 for more information.

The following table lists requirements that apply to Clustered VMDK as a disk resource for a WSFC.

Table 569: Additional Requirements for Clustered VMDK WSFC

Component	Requirement
Windows Cluster Parameters	Set the Windows Cluster Parameter <code>QuorumArbitrationTimeMax</code> to 60
Physical disk /Storage Array LUN	- The physical disk that virtual disks are stored on must support ATS SCSI - The physical disk must also support SCSI-3 Persistent Reservations, specifically Write Exclusive All Registrants (WEAR) type reservations
VMFS Datastore	- VMFS version must be VMFS6 - Fibre Channel (FC) only - The datastore must have attribute ClusteredVMDK set. See Activating Clustered VMDK Support.

What Shared Storage Configurations are Supported with WSFC?

Different WSFC cluster setups support different types of shared storage configurations. Some setups support more than one type. Select the recommended type of shared storage for best results.

Differences between Physical compatibility RDMs and Virtual compatibility RDMs

An RDM (Raw Device Mapped) is a special mapping file in a VMFS volume that manages metadata for its mapped device. The mapping file is presented to the management software as an ordinary disk file, available for file-system operations. To the virtual machine, the storage virtualization layer presents the mapped device as a virtual SCSI device. There are two types of compatibility modes for RDMs.

- A Physical compatibility RDM (pRDM), specifies minimal SCSI virtualization of the mapped device.
 - Virtual Machine Snapshots are not available when the RDM is used in physical compatibility mode.
 - RDMs in physical compatibility mode are required to allow SCSI commands be directly passed to a LUN to satisfy the requirements of SCSI-3 Persistent Reservations used by WSFC.

- A Virtual compatibility RDM (vRDM), specifies full virtualization of the mapped device.
 - VMkernel sends only READ and WRITE to the mapped device. The mapped device appears to the guest operating system the same as a virtual disk file in a VMFS volume.
 - If you are using a raw disk in virtual mode, you can realize the benefits of VMFS, such as advanced file locking for data protection and snapshots for streamlining development processes.
 - Virtual mode is more portable across storage hardware than physical mode, presenting the same behavior as a virtual disk file.
 - Supported for CIB configuration only.
- Note:** Do not use cluster-in-a-box configuration for a production deployment.

For more information, see [KB 2147661](#).

Table 570: Shared Storage Requirements

Storage Type	Clusters on One Physical Machine (Cluster in a Box)	Clusters Across Physical Machines (Cluster Across Boxes)	Clusters of Physical and Virtual Machines (Standby Host Clustering)
Clustered VMDK	No	Yes	No
VMDK in Virtual compatibility mode (SCSi controller sharing mode set to virtual)	Yes	No	No
pRDM (physical compatibility mode)	No	Yes (recommended)	Yes
vRDM (virtual compatibility mode)	Yes (not recommended)	No	No

Note: The multi-writer feature must not be used for a clustered disk resource for WSFC.

Use of in guest options for storage such as iSCSI or SMB shares within guest operating systems configured with WSFC, in any configuration supported by Microsoft, is transparent to ESXi hosts and there is no need for explicit support statements from VMware.

ESXi Supports Round Robin Path Selection Policy (PSP_RR) for disk resources used by WSFC

ESXi supports PSP_RR for WSFC.

- ESXi supports PSP_RR for Windows Server 2012 and above releases. Windows Server 2008 SP2 and earlier are not supported.
- PSPs configured in mixed mode is supported. In a 2 node WSFC cluster in CAB configuration one ESXi host can be configured to use PSP_FIXED and the other ESXi host can use PSP_RR.
- All hosts must be running ESXi 6.0 or later.
- Mixed mode configurations of ESXi 6.0 or later with previous ESXi releases are not supported.
- If the number of paths to a storage device exceeds 5, WSFC cluster storage validation may fail if Round Robin path policy (PSP_RR) with iops =1 is used. In such cases, we recommend setting the number of iops to 5 or more when Round Robin path policy is used. This is not a VMware issue.

ESXi Supports iSCSI for WSFC

ESXi supports iSCSI storage and up to 5 node WSFC clusters.

- ESXi supports iSCSI for Windows Server 2012 and above releases. Windows Server 2008 SP2 and earlier are not supported.
- Cluster-across-box (CAB) and cluster-in-a-box (CIB) are supported. A mixture of CAB and CIB is not supported.
- Clustered VMDKs on VMFS datastore in CAB are not supported over iSCSI.
- No qualification is needed for software iSCSI initiator in a guest operating system.
- N+1 cluster configuration comprising of a cluster between "N" virtual machines on separate ESXi hosts and one physical machine running Windows natively is supported.
- All hosts must be running ESXi 6.0 or later.
- Cluster nodes registered to hosts with different ESXi versions are not supported.
- Mixed mode of iSCSI config is supported.

ESXi Supports FCoE for WSFC

ESXi supports FCoE storage and up to 5 node WSFC clusters.

- ESXi supports FCoE for Windows Server 2012 and above releases. Windows Server 2008 SP2 and earlier are not supported.
- Cluster-across-box (CAB) and cluster-in-a-box (CIB) are supported. A mixture of CAB and CIB is not supported.
- CAB configurations are supported with some cluster nodes on physical hosts. In a CAB configuration, a max of one virtual machine in a host can see a LUN.
- Clustered VMDK is not supported on a datastore connected over FCoE.
- N+1 cluster configuration, in which one ESXi host has virtual machines which are secondary nodes and one primary node is a physical box are supported.
- Required DRS affinity (CIB) or Anti-affinity (CAB) rules for WSFC virtual machines.
- All hosts must be running FCoE initiators. Mixed cluster nodes running FC and FCoE are not supported.
- Mixed mode FCoE configuration is supported.

ESXi Supports vMotion for WSFC

vSphere supports vMotion of a VM hosting a node of a WSFC.

Pre-requisites for vMotion support:

- vMotion is supported only for a cluster of virtual machines across physical hosts (CAB).
- Do not migrate more than 8 WSFC virtual machines at the same time, for VMs with cluster shared resources. This may cause failover of cluster roles to other VMs.
- The vMotion network must be a 10Gbps Ethernet link. 1Gbps Ethernet link for vMotion of WSFC virtual machines is not supported.
- vMotion is supported for Windows Server 2012 and above releases. Windows Server 2008 SP2 and earlier are not supported.
- The WSFC cluster heartbeat time-out must be modified at least to the values listed below:
 - `(get-cluster -name <cluster-name>).SameSubnetThreshold = 10`
 - `(get-cluster -name <cluster-name>).CrossSubnetThreshold = 20`
 - `(get-cluster -name <cluster-name>).RouteHistoryLength = 40`
- The virtual hardware version for the WSFC virtual machine must be version 11 and later.

Modifying the WSFC heartbeat time-out:

WSFC nodes use the network to send heartbeat packets to other nodes of the cluster. If a node does not receive a response from another node for a specified period of time, the cluster removes the node from cluster membership. By default, a guest cluster node is considered down if it does not respond within 5 seconds in Windows 2012, 2012 R2. Other nodes that are members of the cluster will take over any clustered roles that were running on the removed node.

An WSFC virtual machine can stall for a few seconds during vMotion. If the stall time exceeds the heartbeat time-out interval, then the guest cluster considers the node down and this can lead to unnecessary failover. To allow leeway and make the guest cluster more tolerant, the heartbeat time-out interval needs to be modified to allow at least 10 missed heartbeats. The property that controls the number of allowed heart misses is **SameSubnetThreshold**. You will need to modify this from its default value to at least 10. From any one of the participating WSFC cluster nodes run the following command:

```
(get-cluster -name <cluster-name>).SameSubnetThreshold = 10
```

You can also adjust other properties to control the workload tolerance for failover. Adjusting delay controls how often heartbeats are sent between the clustered node. The default setting is 1 second and the maximum setting is 2 seconds. Set the **SameSubnetDelay** value to 1. Threshold controls how many consecutive heartbeats can be missed before the node considers its partner to be unavailable and triggers the failover process. The default threshold is 5 heartbeats and the maximum is 120 heartbeats. It is the combination of delay and threshold that determines the total elapsed time during which clustered Windows nodes can lose communication before triggering a failover. When the clustered nodes are in different subnets, they are called **CrossSubnetDelay** and **CrossSubnetThreshold**. Set the **CrossSubnetDelay** value to 2 and the **CrossSubnetThreshold** value to 20.

Note: Recommended values for WSFC hearbeats settings are now defaults in Windows Server 2016 and above.

VMware vSphere® Virtual Volumes™ Support for WSFC

ESXi supports VMware vSphere® Virtual Volumes™ storage and up to 5 node WSFC clusters. The storage array must support SCSI-3 PRs at the subsidiary LUN level. For NVMe, persistent reservation should be supported at namespace level.

- ESXi supports vVols Storage for Windows Server 2012 and above releases.
- Only Cluster-across-box (CAB) is supported.
- WSFC on vVols can work with any type of disk, "Thin" as well as "Thick"-provisioned disks.
- This feature enables customers to move away from using pRDM.
- WSFC on vVols supports vSphere HA, DRS and vMotion.
- The underlying transport protocol can be FC, iSCSI, FCoE or NVMe Fabric storage (FC, TCP).
- Virtual adapter NVMe support is added as part of vSphere 8.0 U3 along with PVSCSI.
- Cluster-in-a-box (CIB) and a mixture of CAB and CIB is not supported.
- N+1 cluster configuration, in which one ESXi host has virtual machines which are secondary nodes and one primary node is a physical box is not supported.

You Can Hot Extend a Shared vVol Disk

Hot extension of a shared vVol disk is supported with ESXi 8.0 Update 2 or later. This allows you to increase the size of a shared disk without deactivating the cluster.

Extending the size of shared vVol disk is a two step process. First you must increase the size of shared vVol disk using vCenter. Then extend the partition size of the disk from the guest OS.

Best Practices for Hot Extending Shared vVol Disk:

- You can only extend one shared vVol disk at a time from a single VM. You cannot extend multiple shared vVol disks simultaneously.
- Do not reconfigure any other VM parameter while you are changing the size of a shared vVol disk. Any other change you make in VM configuration will not be performed during this operation.
- All VMs sharing the disk must be accessible and registered with vCenter. VMs cannot be in a suspended state or in maintenance mode. VMs cannot be in APD or PDL state.
- Do not perform shared disk extension when a cluster upgrade is in progress and there are hosts in the cluster that do not support this feature.

- Hot extend of a shared vVol disk is not supported for disks that have IOFilters attached.

1. Select any one VM that is sharing the vVol disk you want to extend and click **Edit Settings**.
2. Select the disk you want to extend and enter the new

Edit Settings |



> Memory	16	GB	
▼ Hard disks 5 total 367 GB			
> Hard disk 1	52	GB	
> Hard disk 2	121	GB	
> Hard disk 3	5	GB	
> Hard disk 4	90	GB	
> Hard disk 5	99	GB	
> SCSI controller 0	VMware Paravirtual		
> SCSI controller 1	VMware Paravirtual		
> Network adapter 1	Public	☒ Connected	
> Network adapter 2	Private	☒ Connected	
> Network adapter 3	VM Network	☒ Connected	
> CD/DVD drive 1	Client Device	☐ Connected	
> USB xHCI controller	USB 3.1		

CANCEL

OK

size.

3. Select **OK**.

All VM operations on all VMs that are sharing the vVol disk that is being extended are blocked until the extend operation is complete.

After the operation is complete, each VM that is sharing the vVol disk shows the new disk

Edit Settings | ×

> Memory	16	GB	⌵
▼ Hard disks * 5 total 373 GB			
> Hard disk 1	52	GB	⌵
> Hard disk 2	121	GB	⌵
> Hard disk 3	5	GB	⌵
> Hard disk 4	90	GB	⌵
> Hard disk 5 *	105	GB	⌵
> SCSI controller 0	VMware Paravirtual		⋮
> SCSI controller 1	VMware Paravirtual		⋮
> Network adapter 1	Public	☒ Connected	⋮
> Network adapter 2	Private	☒ Connected	⋮
> Network adapter 3	VM Network	☒ Connected	⋮
> CD/DVD drive 1	Client Device	☐ Connected	⋮
> USB xHCI controller	USB 3.1		⋮

CANCEL OK

size.

Note: If the operation fails, vCenter displays an error message including the list of hosts where the disk extension failed. Run the extend operation again from one of the failed hosts.

After the size of the vVol disk is extended you must also extend the disk partition size from the guest OS. This must be done from the VM which owns the disk resource. Use the appropriate utility depending on your guest OS. Before you extend the size, verify all VMs that are sharing the disk display the new extended size. You can verify using the "disk mgmt tool" where the amount of increased size is displayed as unallocated. The Windows Fail over clustering tool shows the increased size of the shared vVol disk.

To extend the size of shared vVol disk for Windows Guest OS:

1. Open the Disk Management snap-in and right click the data volume that you want to extend.
2. Select **Extend Volume** and follow the instructions in the **Extend Volume Wizard**.
3. After the Extend Volume operation is complete, the disk displays the increased size.

vSphere WSFC Setup Limitations

Before you set up WSFC in a virtual environment, review the list of functions that are not supported for this release, and requirements and recommendations that apply to your configuration.

The following environments and functions are not supported for WSFC setups with this release of vSphere:

- Using VMDKs on NFS datastore as a shared disk resource for WSFC.
- Mixed environments, such as configurations where one cluster node is running a different version of ESXi than another cluster node.
- vSphere Fault Tolerance (FT).
- Migration with vMotion® of clustered virtual machines using vRDMs.
- N-Port ID Virtualization (NPIV)
- ESXi hosts or clusters that have memory overcommitment are not suitable for deploying WSFC virtual machines. Memory overcommitment can cause virtual machines to stall for short durations. This can be significantly disruptive as the WSFC clustering mechanism is time-sensitive and timing delays can cause the virtual machines to behave incorrectly.
- Suspend or resume of a VM, node of a WSFC.
- Storage spaces are not supported with Failover clustering on Windows 2012 and above.
- VM configurations changes:
 - Hot adding memory.
 - Hot adding CPU.
 - Increasing the size of a shared disk is supported only for the disks created from vSphere Virtual Volumes storage with vSphere 8.0 U2 or later.
 - Pausing or resuming the virtual machine state. This I/O intensive operation is disruptive of the timing sensitive WSFC clustering software.
 - Hot adding storage controllers LSI Logic SAS, VMware Paravirtual or NVMe Controller.
 - Hot adding network adapters.
 - Any other hardware changes while cluster VMs are in powered on state except hot add disk or hot share disk.
 - Using snapshots.
 - Cloning a VM.
- Sharing disks between virtual machines without a clustering solution may lead to data corruptions
- NVMe controller for WSFC is only supported for Windows Server 2022 with OS Build 20348.1547 and the shared disks should be created only from Clustered VMDK activated datastores (SCSI or NVMe). To use NVMe Controller, virtual machine hardware version must be 21 or later.

WSFC and Booting from a SAN

You can boot a VM from a SAN, node of a WSFC.

Booting from a SAN is complex. Problems that you encounter in physical environments extend to virtual environments. For general information about booting from a SAN, see the *vSphere Storage* documentation.

Follow these guidelines when you place the boot disk of a virtual machine on a SAN-based VMFS volume:

- Consider the best practices for boot-from-SAN that Microsoft publishes in the following knowledge base article: <http://support.microsoft.com/kb/305547/en-us>.
- Test clustered configurations in different failover scenarios before you put them into production environments.

Cluster Virtual Machines Across Physical ESXi Hosts

You can create a WSFC cluster that consists of two or more virtual machines on two ESXi or more hosts. Use this method for production deployments.

A cluster across physical hosts requires specific hardware and software.

- Up to 5 nodes in a single WSFC is supported

- Supported shared storage configuration, such as RDM in physical compatibility mode, vVol, clustered VMDK. For more information, see [What Shared Storage Configurations are Supported with WSFC?](#).
- If the cluster validation wizard completes with the warning: `Validate Storage Spaces Persistent Reservation` you can safely ignore the warning. This check applies to the Microsoft Storage Spaces feature, which does not apply to VMware vSphere.
- ESXi host now supports up to 16 WSFC clusters (multi-cluster) running on the same ESXi host.
- With vSphere 8.0 U2 or later, you can use NVMe controller in addition to existing Paravirtual controller for WSFC with Clustered VMDK for Windows Server 2022 (OS Build 20348.1547) and later. To use NVMe Controller virtual machine hardware version must be 21 or later.

Clustered VMDK support for WSFC

vSphere supports the use of VMDKs on clustered datastore as shared disk resources for a WSFC. Using VMDKs reduces the extra overhead to manage the virtual disks compared to pRDMs.

Microsoft Clustering Service uses SCSI-3 PRs commands to coordinate access to a clustered disk resource. These commands (PR-IN and PR-Out) are emulated at VSCSI layer on a datastore. The feature requires support from the datastore perspective. A datastore configured to host clustered VMDKs is referred to as a clustered VMDK datastore in this document.

With vSphere 8.0 U2, a new configuration using WSFC with Clustered VMDK and Windows Server 2022 or later supports NVMe Virtual Adapters.

With vSphere 8.0 U3, VMFS VMDKs also support storage from NVMe TCP array using PVSCSI/NVME controllers.

Activating Clustered VMDK Support

You can activate clustered VMDK support when you create a new VMFS6 datastore, or activate clustered VMDK on an existing VMFS6 datastore.

Before activating clustered VMDK support, ensure all hosts connected to the datastore are using ESXi 7.0 or later and are managed by vCenter 7.0 or later. All hosts connected to the datastore must be managed by the same vCenter while deactivating or activating the clustered VMDK flag on the datastore. Once the clustered VMDK flag is activated or deactivated, hosts can be managed by any vCenter with version 7.0 or later.

NVMe FC Datastores are supported for activating clustered VMDK from vSphere 8.0 onwards. With vSphere 8.0 U2, a new configuration using WSFC with Clustered VMDK and Windows Server 2022 or later supports NVMe Virtual Adapters. In vSphere 8.0 U3 or later NVMe TCP Datastores are supported.

1. Navigate to any inventory object that is a valid parent object of a datastore, such as a host, a cluster, or a data center, and click the **Datastores** tab.
2. Right-click the parent object and select **Create a New Datastore**.
3. On the **New Datastore** page choose a device with the attribute **Clustered VMDK Supported** set to `Yes`. Create a VMFS6 version Datastore.
4. Once the Datastore is created go to the **Configure Datastore** tab.
5. Under **General** go to the **Datastore Capabilities** section. Click the `Enable` button for **Clustered VMDK**.
6. Once activated, you can use the **Disable** button to deactivate clustered VMDK. Disable will not work if there are any WSFC VMs using clustered VMDKs on the datastore. All virtual machines with shared disks on the datastore must be powered off.

Configuration Maximums for Clustered VMDK Support

Configuration Maximums for Clustered VMDK support.

Configuration	Maximum
Max number of VMs in a single WSFC cluster	5
Max number of WSFC clusters (multi-cluster) running on the same set of ESXi hosts	16
Max number of clustered VMDKs per ESXi host	192

Recommendations for using Clustered VMDKs with WSFC

Follow these recommendations when using Clustered VMDKs with WSFC.

1. Do not present LUNs used for clustered VMDKs to ESXi hosts unless the host is configured with ESXi 7.0 or later. This might cause slow boot times, host to become unresponsive and other issues. A host with a version lower than ESXi 7.0 cannot mount a clustered VMDK datastore. This is because the ESXi hosts on which WSFC VMs run must have physical Persistent reservation (SCSI/NVMe) of type WEAR on the device. A host must have ESXi 8.0 or later to mount a clustered VMDK datastore, if the backend LUN is created from NVMe FC SAN.
With vSphere 8.0 U2, a new configuration using WSFC with Clustered VMDK and Windows Server 2022 or later supports NVMe Virtual Adapters.
With vSphere 8.0 U3, Clustered VMDKs also support storage from NVMe TCP array using PVSCSI/NVMe controllers for Windows Server 2022 or later.
2. Make sure that all VMs hosting nodes of WSFC are migrated off or powered off properly before removing it from a clustered VMDK datastore to ensure the resources, like heartbeat (HB) slots, are freed. If a VM fails or there is an APD on the clustered VMDK datastore during power-off, always power on and power off before removing the VM from a cluster.
3. Do not combine clustered and non-clustered VMDKs on the same clustered datastore. Although VMs that are using non-shared disks on a clustered datastore will continue to work normally and support all operations such as snapshot and clone.
4. Do not keep clustered VMDKs for different clusters on the same shared datastore. Use a different clustered datastore for different WSFC clusters.
5. Set vHardware (virtual compatibility) to vSphere 7.0 or later while using the clustered VMDK feature.

Requirements for using Clustered VMDKs with WSFC

You must follow these requirements when using Clustered VMDKs with WSFC.

1. You must set the Windows Cluster Parameter **QuorumArbitrationTimeMax** to 60.
2. The physical disk that virtual disks are stored on must support ATS SCSI commands.
3. The physical disk must support SCSI-3 Persistent Reservations, specifically Write Exclusive All Registrants (WEAR) type reservations.
4. VMDKs must be Eager Zeroed Thick Provisioned.
5. Storage devices can be claimed by VMware Native Multipathing Plugin (NMP) or by any other third-party MultiPathing Plugin (MPP). Check with vendors regarding the support for Clustered VMDK before using a vendor's third-party MultiPathing Plugin (MPP).

Limitations of Clustered VMDK support for WSFC

Clustered VMDK support is only possible with the following limitations:

1. Cluster in a Box (CIB) configuration is not supported.
2. Spanned VMFS datastores cannot store clustered VMDKs.
3. A physical LUN with a clustered datastore cannot have any additional VMFS datastores.

4. On-line Storage vMotion and snapshots are not supported.
5. Only disks with 512 sector size are supported.
6. VMFS5 and earlier are not supported for storing clustered VMDKs.
7. Only datastores connected through Fibre Channel (FC) are supported in this release.
8. Clustered VMDK datastore cannot be used as an ESXi installation disk.
9. Diagnostic partition is not supported on disks that are marked clustered.
10. SCSI-2 reservations are not supported on clustered VMDKs.
11. Hot expansion of a VMDK that is associated with a clustered VM is not supported.
12. Moving more than one node of WSFC to the same ESXi host is not allowed and causes vMotion to fail. You must use anti-affinity DRS Rules to separate VMs and nodes of the cluster on different ESXi hosts.
13. Cloning and Fault Tolerance are not supported.

Create the First Node for WSFC Clusters Across Physical Hosts

To create the first node, you create and configure a virtual machine and install a guest operating system on the virtual machine.

1. Open the vSphere Client and connect to the vCenter Server system.
Use the user name and password of the user who will have administrator permissions on the virtual machine.
2. In the vSphere Client, right-click the host and select **New Virtual Machine**.
3. Proceed through the wizard to create the virtual machine.

Page	Action
Creation type	Select Create a new virtual machine .
Name and folder	Enter a name and select a location.
Compute resource	Select a cluster, host, vApp or resource pool to run this virtual machine.
Storage	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Compatibility	The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.
Guest operating system	Select the guest operating system that you intend to install.
Customize hardware	Select virtual hardware, advanced virtual machine options, and SDRS rules if required or applicable.
Ready to Complete	Review your selections.

4. Click **Finish** to complete creating the virtual machine.
Note: Do not add shared cluster disks at this time.
5. Install a Windows Server operating system on the virtual machine.

Create Additional Nodes for a WSFC Across Physical Hosts

To create additional nodes in a cluster of virtual machines across physical hosts, you can create a template from the first virtual machine and use it to deploy additional virtual machines or use the procedure outlined previously to manually create additional VMs onto another ESXi host.

CAUTION: If you clone a virtual machine with an RDM setup, the cloning process converts the RDMs to virtual disks. Unmap all RDMs before cloning, and remap them after cloning is complete.

CAUTION: If you clone a virtual machine with clustered VMDKs, unmap all clustered VMDKs before cloning, and remap them after cloning is complete.

CAUTION: Use Windows sysprep operations before cloning to a template.

1. In the vSphere Client, right-click the first virtual machine you created and select **Clone > Clone to Template**.
2. Proceed through the wizard to create the virtual machine template.

Page	Action
Name and folder	Enter a name (for example, Node2_Template) and select a location.
Compute Resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .
Storage	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Ready to Complete	Click Finish to create the virtual machine template.

3. Right-click the virtual machine template and select **Deploy VM from this Template**.
4. Proceed through the deployment wizard to deploy the virtual machine.

Page	Action
Name and folder	Enter a name (for example, Node2) and select a location.
Compute resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .
Datastore	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Clone options	Select Customize the operating system .

5. Select a new guest operating system from the list.
 - a) Click the **Create a new specification** button to add a new guest operating system. Proceed through the **New VM Guest Customization Spec** wizard.
 - b) Click **Finish** to exit the wizard.
6. Click **Finish** to deploy the virtual machine.

Add Hard Disks to the First VM, first node of a WSFC, for Clusters Across Physical Hosts with pRDMs

In an WSFC cluster with shared disk, storage disks are shared between nodes.

Before you add hard disks to the first node, complete the following tasks:

- If you are using physical compatibility mode RDMs, ask your SAN administrator for the location of unformatted SAN LUNs. The hard disks you create in this task must point to SAN LUNs.
- If you are using vVols storage for shared disk, please see *Add Hard Disks to the First Node for Clusters Across Physical Hosts with VMware vSphere® Virtual Volumes™*.
- If you are using Clustered VMDKs on VMFS datastores, please see *Add Hard Disks to the First Node for Cluster Across Physical Hosts with Clustered VMDKs on VMFS datastores*.

Note: Use RDMs in physical compatibility mode.

1. In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
2. Click the **New device** drop-down menu, select **SCSI Controller**.
3. In new SCSI Controller, select **VMware Paravirtual** and set **SCSI Bus Sharing** to **Physical**. Click **OK**.

Note: LSI SAS is supported, but VMware Paravirtual is recommended.

4. Select the newly created virtual machine in step 1, right-click and select **Edit Settings**.
5. Click the **New device** drop-down menu, select **RDM Disk**.
6. Select an unformatted LUN and click **OK**.
7. Expand **New Hard disk** and select a datastore for **Location**.
This datastore must be on shared storage because you need a single shared RDM file for each shared disk.
8. Select **Physical** as the compatibility mode.
9. Select the newly created SCSI controller in step 2, for virtual device node (for example, select **SCSI (1:0)**).
Note: This must be a new SCSI controller. You cannot use SCSI 0. SCSI controller 0 is reserved for boot disks and disks which should not be shared
10. Click **OK**.
The wizard creates a new hard disk.

Add Hard Disks to the First Node of a WSFC Across Physical Hosts with vVols

In an WSFC cluster, storage disks are shared between nodes.

Before you add hard disks to the first node, complete the following tasks:

- Use HW Version 13 or later. See [Virtual Volumes \(vVols\) now supports WSFC](#).
 - Prepare vVol deployment. See *vSphere Storage* documentation for more information.
1. In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
 2. Click the **New device** drop-down menu, select **SCSI Controller**.
 3. In new SCSI Controller, select **VMware Paravirtual** and set **SCSI Bus Sharing** to **Physical**. Click **OK**.
Note: LSI SAS is supported, but VMware Paravirtual is recommended.
 4. Select the newly created virtual machine in step 1, right-click and select **Edit Settings**.
 5. Click the **New device** drop-down menu, select **New Hard Disk**.
 6. Expand **New Hard disk** and select the disk size required.
 7. Select a datastore for **Location**.
This datastore must be on shared vVols storage.
 8. Under Disk Provisioning, select either Thick or Thin Provision.
 9. Expand the New Hard Disk.
 10. Select the newly created SCSI controller in step 2, for virtual device node (for example, select **SCSI (1:0)**).
Note: This must be a new SCSI controller. You cannot use SCSI 0.
 11. Click **OK**.
The wizard creates a new hard disk.

Add Hard Disks to First Node for Cluster Across Physical Hosts with Clustered VMDKs on VMFS

Add hard disks to the first node for cluster across physical hosts with clustered VMDKs on VMFS datastores.

1. In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
2. Click the **New device** drop-down menu, select **SCSI Controller** or **NVMe Controller**.
 - If a new SCSI Controller is created, select **VMware Paravirtual** and set **SCSI Bus Sharing** to **Physical**. Click **OK**.

- If a new NVMe Controller is created, set **NVME Bus Sharing** to **Physical**. Click

> CPU	2	
> Memory	4	
> New Hard disk *	90	
> New SCSI controller	VMware	
> New Network	VM Net	
> New CD/DVD Drive	Client D	
> New USB Controller	USB 3.2	
> Video card	Specify	
> New SATA Controller	New SAT	

3. Select the newly created virtual machine in step 1, right-click and select **Edit Settings**
4. Click the **New device** drop-down menu, select **New Hard Disk**.
5. Expand **New Hard disk** and select the disk size required.
6. Select a datastore for **Location**.
This datastore must be the datastore with clustered VMDK flag enabled.
7. Under Disk Provisioning, select **Thick Provision with Eager zeroed**.
8. Select the newly created SCSI controller or NVMe controller in step 2, for virtual device node (for example, select **SCSI (1:0)** or **NVMe (1:0)**).
Note: This must be a new SCSI controller. You cannot use SCSI 0.
9. Click **OK**.
The wizard creates a new hard disk.

Add Hard Disks to Additional Nodes for Clusters Across Physical Hosts

To allow shared access to disk resources, point to existing disks on the VM, the first node of a WSFC cluster. Use the same SCSI IDs while assigning disks to all additional nodes.

Before you begin, obtain the following information:

- Obtain SCSI IDs for all virtual disks to be shared.
- Obtain disk file path on datastore for all shared disks.

Note: The procedure for clustered VMDK, pRDMs and vVol is the same.

1. In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
2. Click the **New device** drop-down menu, select **SCSI Controller** or **NVMe Controller**.
 - If a new SCSI Controller is created, select **VMware Paravirtual** and set **SCSI Bus Sharing** to **Physical**. Click **OK**.

- If a new NVMe Controller is created, set **NVMe Bus Sharing** to **Physical**. Click

> CPU	2 	
> Memory	4	
> New Hard disk *	90	
> New SCSI controller	VMware	
> New Network	VM Net	
> New CD/DVD Drive	Client D	
> New USB Controller	USB 3.2	
> Video card	Specify	
> New SATA Controller	New SAT	

You must select same SCSI controller or NVMe controller type.

3. Select the newly created virtual machine in step 1, right-click and select **Edit Settings**
 4. Click the **New device** drop-down menu, select **Existing Hard Disk**.
 5. In **Disk File Path**, browse to the location of the disk to be shared specified for the first node.
 6. Expand **New Hard disk**.
 7. Select the same SCSI or NVMe ID you chose for the first virtual machine's shared storage disks, (for example, select **SCSI (1:0)** or **NVMe (1:0)**).
- Note:** The disk SCSI or NVMe ID for this virtual machine's shared storage must match the corresponding SCSI or NVMe ID for the first virtual machine.
8. Click **OK**.

Use WSFC in an vSphere HA and vSphere DRS Environment

When you use WSFC in a vSphere High Availability (vSphere HA) or vSphere Distributed Resource Scheduler (DRS) environment, you must configure your hosts and virtual machines to use certain settings. All hosts that run WSFC virtual machines must be managed by a vCenter Server system.

VMware vSphere provides support for vMotion of a VM, node of a WSFC.

Enable vSphere HA and vSphere DRS in ESXi

ESXi hosts that are running WSFC virtual machines can be part of a vSphere cluster with both vSphere HA and vSphere DRS enabled. You can enable vSphere HA and vSphere DRS in the Cluster Settings dialog box.

1. Browse to the cluster in the vSphere Client.
2. Click **Configure**.
3. Under **Services**, select **vSphere DRS** and click **Edit**.
4. Select the **Turn ON vSphere DRS** and **Turn on vSphere HA** check boxes.
5. Click **OK**.

Create VM-VM Affinity Rules for WSFC Virtual Machines

For WSFC virtual machines in a cluster, you must create VM-VM affinity or anti-affinity rules. VM-VM affinity rules specify which virtual machines should be kept together on the same host (for example, a cluster of WSFC virtual machines on one physical host). VM-VM anti-affinity rules specify which virtual machines should be kept apart on different physical hosts (for example, a cluster of WSFC virtual machines across physical hosts).

For a cluster of virtual machines on one physical host, use affinity rules.

Note: Do not use a cluster on one physical host for production deployment.

For a cluster of virtual machines across physical hosts, use anti-affinity rules.

Note: vMotion is supported for a virtual machines, nodes of a WSFC, across physical ESXi hosts.

1. Browse to the cluster in the vSphere Client.
2. Click **Configure**.
3. Under **Configuration**, click **VM/Host Rules**.
4. Click **Add**.
5. In the **Create VM/Host Rule** dialog box, type a name for the rule.
6. From the **Type** drop-down menu, select a rule.
 - For a cluster of virtual machines on one physical ESXi host, select **Keep Virtual Machines Together**.
 - For a cluster of virtual machines across physical ESXi hosts, select **Separate Virtual Machines**.
7. Click **Add**.
8. Select the two virtual machines to which the rule applies and click **OK**.
9. Click **OK**.

Set DRS Automation Level for WSFC Virtual Machines

Set the automation level of all virtual machines in an WSFC cluster to Partially Automated. When you set the vSphere DRS automation level for the virtual machine to Partially Automated, vCenter Server will perform initial placement of virtual machines when they are powered on and will provide migration recommendations for them.

1. Browse to the cluster in the vSphere Client.
2. Click **Configure**.
3. Under **Services**, select **vSphere DRS** and click **Edit**.
4. Expand DRS Automation, under Virtual Machine Automation select the **Enable individual virtual machine automation levels** check box and click **OK**.
5. Under **Configuration**, select **VM Overrides** and click **Add**.
6. Click the **+** button, select the WSFC virtual machines in the cluster and click **OK**.
7. Click the Automation level drop-down menu, and select **Partially Automated**.
8. Click **OK**.

Note: VMware recommends partially automated mode for WSFC virtual machines, but there is no technical restriction which prevents the setting for WSFC virtual machines to be fully automated. If the fully automated setting is used please tune WSFC advanced settings to prevent WSFC failover during vMotion.

Using vSphere DRS Groups and VM-Host Affinity Rules with WSFC Virtual Machines

You can set up two types of DRS groups: virtual machine DRS groups, which contain at least one virtual machine, and host DRS groups, which contain at least one host. A VM-Host rule establishes an relationship between a virtual machine DRS group and a host DRS group.

You must use VM-Host affinity rules because vSphere HA does not obey VM-VM affinity rules. This means that if a host fails, vSphere HA might separate clustered virtual machines that are meant to stay together, or vSphere HA might put clustered virtual machines that are meant to stay apart on the same host. You can avoid this problem by setting up DRS groups and using VM-Host affinity rules, which are obeyed by vSphere HA.

For a cluster of virtual machines on one physical host, all WSFC virtual machines must be in the same virtual machine DRS group, linked to the same host DRS group with the affinity rule "Must run on hosts in group."

For a cluster of virtual machines across physical hosts, all WSFC virtual machines must be in the same virtual machine DRS group and all hosts must be in the same host DRS group. The virtual machine and host DRS groups must be linked the affinity rule "Must run on hosts in group."

CAUTION: Limit the number of hosts to two when you define host DRS group rules for a cluster of virtual machines on one physical ESXi host. (This does not apply to clusters of virtual machines across physical ESXi hosts.) Since vSphere HA does not obey VM-VM affinity rules, virtual machines in the configuration could be spread across hosts during a vSphere HA recovery from host failure if more than two hosts are included in a host DRS group rule.

Create a Virtual Machine DRS Group (WSFC)

Before you can create a VM-Host affinity rule, you must create the host DRS group and the virtual machine DRS group that the rule applies to.

For both, a cluster of virtual machines on one physical host(CIB) and a cluster of virtual machines across physical hosts(CAB), create one virtual machine DRS group that contains all WSFC virtual machines. For example, VMGroup_1 contains WSFC_VM_1, WSFC_VM_2 ... WSFC_VM_5.

1. Browse to the cluster in the vSphere Client.
2. Click **Configure**.
3. Under **Configuration**, select **VM/Host Groups** and click **Add**.
4. In the **Create VM/Host Group** dialog box, type a name for the group.
5. Select **VM Group** from the **Type** drop down box and click **Add**.
6. Click the check box next to a virtual machine to add it. Continue this process until all desired virtual machines have been added.
 - For a cluster of virtual machines on one physical host, add all WSFC virtual machines to one group.
 - For a cluster of virtual machines across physical hosts, add all WSFC virtual machines to one group.
7. Click **OK**.

Create a Host DRS Group (WSFC)

Before you can create a VM-Host affinity rule, you must create the host DRS group and the virtual machine DRS group that the rule applies to.

For a cluster of virtual machines on one physical host, create one host DRS group that contains both the ESXi hosts. For example, HostGroup_1 contains ESXi_HOST_1, ESXi_HOST_2.

For a cluster of virtual machines across physical hosts, create one host DRS group that contains all the ESXi hosts. For example, HostGroup_1 contains ESXi_HOST_1 running WSFC_VM_1, ESXi_HOST_2 running WSFC_VM_2 ESXi_HOST_5 running WSFC_VM_5 and ESXi_HOST_6 as a standby host.

1. Browse to the cluster in the vSphere Client.
2. Click **Configure**.
3. Under **Configuration**, select **VM/Host Groups** and click **Add**.
4. In the **Create VM/Host Group** dialog box, type a name for the group (for example, HostGroup_1).
5. Select **Host Group** from the **Type** drop down box and click **Add**.
6. Click the check box next to a host to add it. Continue this process until all desired hosts have been added.
7. Click **OK**.

Set up VM-Host Affinity Rules for DRS Groups (WSFC)

Create VM-Host affinity rules to specify whether the members of a selected virtual machine DRS group can run on the members of a specific host DRS group.

Create virtual machine DRS groups that contain one or more WSFC virtual machines as described in [Create a Virtual Machine DRS Group \(WSFC\)](#).

Create host DRS groups that contain one or more ESXi hosts, as described in [Create a Host DRS Group \(WSFC\)](#).

1. Browse to the cluster in the vSphere Client.
2. Click **Configure**.
3. Under **Configuration**, select **VM/Host Rules** and click **Add**.
4. In the **Create VM/Host Rule** dialog box, type a name for the rule.
5. From the **Type** menu, select **Virtual Machines to Hosts**.
6. Select the virtual machine DRS group and the host DRS group to which the rule applies.
For example, select VMGroup_1 and HostGroup_1.
7. Select **Must run on hosts in group**.
8. Click **OK**.

vSphere WSFC Setup Checklist

When you set up a WSFC on VMware vSphere, see the checklists to configure your environment according to the requirements. You can also use the checklists to verify that your setup meets the requirements if you need technical support.

Requirements for Clustered Disks

Each type of clustered disk has its own requirements, depending on whether it is in a single-host cluster or multihost cluster.

Table 571: Requirements for Clustered Disks

Component	Single-Host Clustering	Multihost Clustering
Clustered VMDK	SCSI bus sharing mode set to virtual is not supported.	SCSI bus sharing or NVMe Bus Sharing mode set to physical is supported. NVMe is only supported on vSphere 8.0 U2 or later with Windows Server 2022 or later.
VMDK (In virtual compatibility mode)	SCSI bus sharing mode set to virtual is supported. (Not Recommended)	SCSI bus sharing mode set to Physical is not supported.
vRDM	Device type must be set to virtual compatibility mode. SCSI bus sharing mode must be set to virtual mode. A single, shared RDM mapping file for each clustered disk is required.	Not supported.
pRDM	Not supported.	RDM device type must be set to Physical compatibility mode during hard disk creation. SCSI bus sharing mode of vSCSI Controller hosting pRDM must be set to physical. A single, shared RDM mapping file for each clustered disk is required.
vVols storage	Not Supported.	Bus sharing mode of PVSCSI/NVMe controller must be set to physical.
All types	All clustered nodes must use the same target SCSI ID (on the virtual SCSI adapter) or same target NVMe ID (on the virtual NVMe adapter) for the same clustered disk. A separate virtual SCSI/NVMe adapter must be used for clustered disks. Multiple virtual SCSI/NVMe adapters for clustered disks (up to supported maximum) can be used to enhance performance. In this setup at least one controller should be dedicated to non-shared disk.	

Other Requirements and Recommendations

The following table lists the components in your environment that have requirements for options or settings.

Table 572: Other Clustering Requirements and Recommendations

Component	Requirement
Disk	Select Thick Provision with Eager zeroed during disk provisioning for VMDKs (for clustered VMDKs or VMDKs in virtual compatibility mode). For vVols select either thick or thin. Modify the .vmx file for the WSFC VMs to add these lines: • <code>scsiX.returnNoConnectDuringAPD = "TRUE"</code> • <code>scsiX.returnBusyOnNoConnectStatus = "FALSE"</code> • <code>nvmeX.returnNoConnectDuringAPD = "TRUE"</code> (if you use NVMe controller for boot device) • <code>nvmeX.returnBusyOnNoConnectStatus = "FALSE"</code> (if you use NVMe controller for boot device) • Where X is the boot device SCSI bus controller ID number. By default, X is set to 0. The only disks that you should not create with the <code>Thick Provision</code> option are RDM files (both physical and virtual compatibility mode). Disk I/O timeout is 60 seconds or more (<code>HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeoutValue</code>).

Component	Requirement
Windows	Use: - Windows Server 2012 - Windows Server 2012 R2 - Windows Server 2016 - Windows Server 2019 with vSphere 6.7 Update 3 - Windows Server 2022 with vSphere 7.0 or later For Windows Server 2012 and above, you can use up to five cluster nodes. Note: If you recreate the cluster, this value might be reset to its default, so you must change it again. The cluster service must restart automatically on failure (first, second, and subsequent times). Note: Windows Server 2008 and 2008 R2 have reached the end of extended support (no regular security updates). Consider the vendor support ability while hosting a WSFC on VMware vSphere.
ESXi	Do not overcommit memory. Set the Memory Reservation (minimum memory) option to the same as the amount of memory assigned to the virtual machine. ESXi uses a different technique to determine if Raw Device Mapped (RDM) LUNs are used for WSFC cluster devices, by introducing a configuration flag to mark each device as "perennially reserved" that is participating in an WSFC cluster. For ESXi hosts hosting passive WSFC nodes with RDM LUNs, use the <code>esxcli storage core device setconfig -d <naa.id> --perennially-reserved=true</code>. See KB 1016106 for more information.
Multipathing	Contact your multipathing software vendor for information and support of non-VMware multipathing software in vSphere.

Required Information for Technical Support

The following table lists the files and settings to collect if you require technical support. Technical support uses these files and settings to analyze issues with a WSFC setup.

Note: Verify that your setup complies with the checklists in [Table 1](#) and [Table 2](#) before you contact technical support.

Table 573: Information Required by Technical Support

File or Information	Description or Location
<code>vm-support tarball</code>	Contains the vmkernel log, virtual machine configuration files and logs, and so on.
Application and system event logs of all virtual machines with the problem	
Cluster log of all virtual machines with the problem	%ClusterLog%, which is usually set to %SystemRoot%\cluster\cluster.log or use Get-ClusterLog Powershell Cmdlet
Disk I/O timeout	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue
vSphere Client display names and Windows NETBIOS names of the virtual machines experiencing the problem	
Date and time that the problem occurred	
SAN configuration of the ESXi system	Details about LUNs, paths, and adapters.
(Optional) Memory dump of the virtual machine	Required if a guest virtual machine fails (a blue screen appears with an error message).

Cluster Virtual Machines on One Physical Host

You can create an WSFC cluster with up to five nodes on a single ESXi host.

Note: For production deployments, do not set up a WSFC on VMs residing on a single ESXi host.

Note: Windows Server 2012 and above systems support up to five nodes (virtual machines). For supported guest operating systems see [Table 4](#).

Create the First Node of a WSFC on One Physical ESXi Host

To create the first node, you create and configure a virtual machine and install a guest operating system on the virtual machine.

1. Open the vSphere Client and connect to the vCenter Server system.

Use the user name and password of the user who will have administrator permissions on the virtual machine.

2. In the vSphere Client, right-click the host and select **New Virtual Machine**.

3. Proceed through the wizard to create the virtual machine.

Page	Action
Creation type	Select Create a new virtual machine .
Name and folder	Enter a name and select a location.
Compute resource	Select a cluster, host, vApp or resource pool to run this virtual machine.
Storage	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Compatibility	The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.
Guest operating system	Select the guest operating system that you intend to install.
Customize hardware	Select virtual hardware, advanced virtual machine options, and SDRS rules.
Ready to Complete	Review your selections.

4. Click **Finish** to complete creating the virtual machine.

Note: Do not add shared cluster disks at this time.

5. Install a Windows Server operating system on the virtual machine.

Create Additional Nodes for Clusters on One Physical Host

Create a template from the first virtual machine and deploy the second node from that template. You can have up to five nodes on Windows Server 2012 or later.

Note: For a clean template, use Windows Sysprep.

CAUTION: If you clone a virtual machine with an RDM setup, the cloning process converts the RDMs to virtual disks. Unmap all RDMs before cloning, and remap them after cloning is complete.

1. In the vSphere Client, right-click the first virtual machine you created and select **Clone > Clone to Template**.

2. Proceed through the wizard to create the virtual machine template.

Page	Action
Name and folder	Enter a name (for example, Node2_Template) and select a location.
Compute Resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .
Storage	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.

Page	Action
Ready to Complete	Click Finish to create the virtual machine template.

3. Right-click the virtual machine template and select **Deploy VM from this Template**.

4. Proceed through the deployment wizard to deploy the virtual machine.

Page	Action
Name and folder	Enter a name (for example, Node2) and select a location.
Compute resource	Select the host or cluster where you will run the virtual machine.
Disk Format	Select Same format as source .
Datastore	Select a datastore as the location for the virtual machine configuration file and the .vmdk file.
Clone options	Select Customize the operating system .

5. Select a new guest operating system from the list.

- Click the **Create a new specification** button to add a new guest operating system. Proceed through the **New VM Guest Customization Spec** wizard.
- Click **Finish** to exit the wizard.

6. Click **Finish** to deploy the virtual machine.

Add Hard Disks to the First Node for Clusters on One Physical Host

In an WSFC cluster, storage disks are shared between nodes. You set up a quorum disk and at least one shared storage disk.

1. In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.

2. Click the **New device** drop-down menu, select **SCSI Controller**.

3. In new SCSI Controller, select **VMware Paravirtual** and set **SCSI Bus Sharing** to **Virtual**. Click **OK**.

Note: LSI SAS is supported, but VMware Paravirtual is recommended.

4. Select the newly created virtual machine in step 1, right-click and select **Edit Settings**

5. Click the **New device** drop-down menu, select **Hard Disk**.

6. Expand **New Hard disk** and select the disk size required.

7. Select a datastore for **Location**.

8. Under Disk Provisioning, select **Thick Provision with Eager zeroed**.

9. Select the newly created SCSI controller in step 2, for virtual device node (for example, select **SCSI (1:0)**).

Note: This must be a new SCSI controller. You cannot use SCSI 0.

10. Click **OK**.

The wizard creates a new hard disk.

Add Hard Disks to Additional Nodes for Clusters on One Physical Host

To allow shared access to clustered services and data, use **Add an Existing disk** on the second VM to point to the disks from the first VM that should be shared. Use the same SCSI IDs.

Before you begin, obtain the following information:

- Obtain SCSI IDs for all virtual disks to be shared.
1. In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
 2. Click the **New device** drop-down menu, select **SCSI Controller**.
 3. In new SCSI Controller, select **VMware Paravirtual** and set **SCSI Bus Sharing** to **Virtual**. Click **OK**.

Note: LSI SAS is supported, but VMware Paravirtual is recommended.
 4. Select the newly created virtual machine in step 1, right-click and select **Edit Settings**
 5. Click the **New device** drop-down menu, select **Existing Hard Disk**.
 6. In **Disk File Path**, browse to the location of the disk to be shared specified for the first node.
 7. Expand **New Hard disk**.
 8. Select the same SCSI ID you chose for the first virtual machine's shared storage disks, (for example, select **SCSI (1:0)**).

Note: The disk SCSI ID for this virtual machine's shared storage must match the corresponding SCSI ID for the first virtual machine.
 9. Click **OK**.

Cluster Physical and Virtual Machines

You can create an WSFC cluster in which each physical server has a corresponding virtual machine. This type of configuration is known as a standby host cluster.

A standby host cluster has specific hardware and software requirements.

- Network configuration should be consistent between a physical and virtual node of a WSFC. Ensure that all required VLANs are available to ESXi hosts, hosting a virtual node of a WSFC cluster.
- Use RDMs in physical compatibility mode. You cannot use virtual disks or RDMs in virtual compatibility mode for shared storage.
- Use the STORport Miniport driver for the Fibre Channel (FC) HBA (QLogic or Emulex) in the physical Windows machine.
- Do not run multipathing software in the physical or virtual machines.
- Use only a single physical path from the host to the storage arrays in standby host configurations.

Create the First Node for a Cluster of Physical and Virtual Machines

The first node in a standby host setup is a physical machine.

For information about setting up a physical machine that participates in an WSFC cluster, see the Failover Cluster for Windows documentation.

- Set up the physical machine using the settings listed in the table.

Component	Requirement
Windows Cluster Administrator application	Advanced minimum configuration if Windows 2012 is used.
Storage	Access to the same storage on a SAN as the ESXi host on which the corresponding virtual machine will run.
Operating system	Installed on each physical machine.

Create the Second Node for a Cluster of Physical and Virtual Machines

To create the second node, you set up a virtual machine, a second node of a WSFC cluster, where the first node is a physical server.

Before you begin, make sure that the shared storage that is visible from the physical machine that you configured in [Create the First Node for a Cluster of Physical and Virtual Machines](#) is also visible from the ESXi host.

1. Open the vSphere Client and connect to the vCenter Server system.
Use the user name and password of the user who will have administrator permissions on the virtual machine.
2. In the vSphere Client, right-click the host and select **New Virtual Machine**.
3. Proceed through the wizard to create the virtual machine.

Page	Action
Creation type	Select Create a new virtual machine .
Name and folder	Enter a name and select a location.
Compute resource	Select a cluster, host, vApp or resource pool to run this virtual machine.
Storage	Select a datastore as the location for the virtual machine configuration file and the virtual machine disk (.vmdk) file.
Compatibility	The host or cluster supports more than one VMware virtual machine version. Select a compatibility for the virtual machine.
Guest operating system	Select the guest operating system that you intend to install.
Customize hardware	Select virtual hardware, advanced virtual machine options, and SDRS rules.
Ready to Complete	Review your selections.

4. Click **Finish** to complete creating the virtual machine.
5. Install a Windows Server operating system on the virtual machine.

Add Hard Disks to the Second Node for a Cluster of Physical and Virtual Machines

When you add hard disks to the second node, you set up the disks to point to the shared storage disks for the first node. The setup allows shared access to clustered disks.

1. In the vSphere Client, select the newly created virtual machine, right-click and select **Edit Settings**.
2. Click the **New device** drop-down menu, select **SCSI Controller**.
3. In new SCSI Controller, select **VMware Paravirtual** and set **SCSI Bus Sharing** to **Physical**. Click **OK**.

Note: LSI SAS is supported, but VMware Paravirtual is recommended.

4. Select the newly created virtual machine in step 1, right-click and select **Edit Settings**.
5. Click the **New device** drop-down menu, select **RDM Disk**.
6. Select an unformatted LUN and click **OK**.
7. Expand **New Hard disk**.
8. Select a datastore for **Location**.
This datastore must be on shared storage because you need a single shared RDM file for each shared disk.
9. Select compatibility mode **Physical**.
10. Select the newly created SCSI controller in step 2, for virtual device node (for example, select **SCSI (1:0)**).
Note: This must be a new SCSI controller. You cannot use SCSI 0.
11. Click **OK**.
The wizard creates a new hard disk.

Install Microsoft Cluster Service

After you set up the first and second nodes, you must configure Microsoft Failover Clustering.

See Microsoft's documentation for creating and configuring server clusters on the Microsoft Web site.

In complex storage solutions, such as an FC switched fabric, a storage unit might have a different identity (target ID or raw disk ID) on each computer in the cluster. Although this is a valid storage configuration, it might cause a problem when you add a node to the cluster.

Create Additional Physical-Virtual Pairs

If you have additional physical machines, you can create additional clusters for each.

1. On the ESXi host, set up an additional virtual machine for the physical machine.
2. Cluster the physical machine with the new virtual machine.

VMware Aria Operations Plugin in vCenter Server

The *VMware Aria Operations Plugin in vCenter Server* documentation provides information about installing and using VMware Aria Operations .

Intended Audience

This information is intended for anyone who wants to install and use VMware Aria Operations in vSphere. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and data center operations.

Introduction

The earlier release of VMware Aria Operations with vCenter Server was shipped with the vSphere Client plugin.

The new VMware Aria Operations plugin in vCenter Server, provides a mechanism to provide specific metrics and high-level information about data centers, datastores, VMs, and hosts, for the vCenter Server and vSAN. The plugin is supported only in the HTML5 version of the vSphere Client.

Installing VMware Aria Operations

The VMware Aria Operations plugin in vCenter Server is integrated with vCenter Server .

The menu **VMware Aria Operations** appears in the vCenter UI when the vCenter is added as an endpoint in Aria Operations.

You can navigate to the VMware Aria Operations plugin from the vSphere Client user interface by clicking **Menu > VMware Aria Operations Manager**.

To view the widgets in the vSphere Client user interface, you need to install a new instance of VMware Aria Operations.

Install a New Instance of VMware Aria Operations in vCenter Server

You can use this option to install a new instance of VMware Aria Operations.

Virtual CPU and Memory Requirements

- 4 vCPU
- 16 GB vRAM

Storage Requirements

- 274 GB (size on disk for thick provisioning)
- Datastore space > 200 GB

Network Requirements

- Internet connection for online install.

1. From the vSphere Client user interface, click **Menu > VMware Aria Operations Configuration**.
2. Click **Install**.

The new Install VMware Aria Operations window appears.

3. Provide the logged in vCenter Server IP and FQDN credentials and click **Test Connection**.
4. Click **Next**.
5. Enter the data center, cluster, host, datastore, and network information.
6. For a static IP configuration, enter the network details. (Optional)
For a DHCP instance, if the fields are left empty, the credentials are assigned by the DHCP server.
7. Read the **Summary** and click **Install**.
The default user name of an installed instance of VMware Aria Operations is `admin` and the default password is `Vmware@123`.

Using the VMware Aria Operations Plugin in vCenter Server

You can use the plugin to view specific vCenter Server details. You can also navigate to VMware Aria Operations from the plugin.

Note:

vCenters in linked mode do not display data in the VMware Aria Operations plugin page.

Monitor Cluster Capacity

VMware Aria Operations plugin in vCenter Server enables you to monitor and analyze details such as hosts, VMs, time remaining, and capacity remaining on each cluster.

View and Analyse Cluster Usage

You can view all the clusters associated with vCenter that is added as an endpoint in Aria Operations. You can also view details such as hosts, VMs, time remaining, and capacity remaining of the resources (CPU, Memory, and Storage) for each cluster.

1. Log in to the vSphere Client.
2. Select **VMware Aria Operations** from the main menu.
3. Monitor and analyze the vCenter cluster usage
 - a) In the **Capacity Overview** panel, you can view the overall information of inventory objects at the vCenter Server level. You can view the number of data centers, clusters, hosts, virtual machines, and resource pools for the vCenter Server.
 - b) In the **Time Remaining** panel, you can view the number of days remaining until the cluster reaches its usable capacity.
 - c) In the **Capacity Remaining** panel, you can keep track of the operational capacity remaining on each cluster. VMware Aria Operations computes the capacity remaining based on the resources (CPU, Memory, and Storage) utilization trend.
 - d) In the **All Clusters** panel, click the >> icon to view the detailed information for each cluster.

Note:

- In the **Overview** tab, click on the **VIEW SUMMARY IN VMWARE ARIA OPERATIONS** link for further information.
- In the **Time Remaining** tab, click on the **VIEW DETAILS IN VMWARE ARIA OPERATIONS** link for further information.
- In the **Capacity Remaining** tab, click on the **VIEW DETAILS IN VMWARE ARIA OPERATIONS** link for further information.

Access VMware Aria Operations

You can manage multiple vCenters using the VMware Aria Operations UI. For more information and troubleshooting, you can open and access VMware Aria Operations by clicking on the **VMWARE ARIA OPERATIONS LINK**.

Documentation Legal Notice

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the “Documentation”) is for your informational purposes only and is subject to change or withdrawal by Broadcom at any time. This Documentation is proprietary information of Broadcom and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Broadcom.

If you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all Broadcom copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to Broadcom that all copies and partial copies of the Documentation have been returned to Broadcom or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, BROADCOM PROVIDES THIS DOCUMENTATION “AS IS” WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL BROADCOM BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF BROADCOM IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice

The manufacturer of this Documentation is Broadcom Inc.

Provided with “Restricted Rights.” Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b) (3), as applicable, or their successors.

Copyright © 2005–2025 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

